

Musings on the Social Discount Rate¹

Arnold C. Harberger and Glenn P. Jenkins

Development Discussion Paper: 2015-01

Abstract

This paper is mainly concerned with weighted average measures of the social discount rate, where the components of the average are the marginal productivity of investment (measured by its gross-of-tax rate of return), and the marginal rate of time preference (measured by the net-of-tax yield of capital). We believe that these components should best be measured using data (the national accounts) that span the whole economy and reflect the product actually produced and the rewards actually perceived. We use a methodology based on just four familiar parameters to generate productivity estimates applicable to a wide range of countries. In the process, we make an adjustment for infrastructure investment, also excluding income from land, monopoly markups, supra-marginal returns due to TFP increases, and returns to capital in financial intermediation. The end products are estimates of social discount rates averaging around 8 percent for the advanced countries, and 10 percent for healthy developing countries and Asian Tigers.

Revised paper published as: Harberger, A. C., & Jenkins, G. P. (2015). Musings on the social discount rate. *Journal of Benefit-Cost Analysis*, 6(01), 6-32.

Keywords: Social discount rate, benefit-cost analysis, opportunity cost, time preference

JEL Classification: H43, D61

¹ This paper is an outgrowth of Harberger's keynote address at the 2013 meetings of the SBCA. Like that presentation, this one is an essay of persuasion, by which we hope to narrow the range of differences that we observe among professionals on the size of social discount rate. Consistent with its origin and its persuasive purpose, we try to maintain a relatively informal tone, taking our readers, as it were, on a sort of guided tour through the maze of problems that arise as we search for an appropriate discount rate. To help maintain this tone, we decided not to make references to specific authors or pieces of literature. In trying to persuade, we want neither to pick fights with some nor to offend others by failing to mention them, labeling this effort as "musings" in the hope of setting our desired tone.

This paper will explore some of the main ideas that arise in today's debates about the social discount rate. We hope that, by emphasizing the links between theory and real world observations, it will help to narrow the wide range of disagreement that one now observes on this topic.

To set the stage, we will focus on the choice of the real discount rate to be used in the analysis of actual public projects, programs and policies. Any viable candidate for this choice has to be able to survive scrutiny by those who make and influence policies and by interested members of the general public, as well as by the economics profession itself. Thus we must be able to defend our choice and support our methodology in ways that can be easily communicated. We will soon see, however, that meeting this condition is not easy. The road that leads to a given final choice of a discount rate has many forks along the way, at any of which a different choice would set us on a different path.

Starting with something simple, we must recognize the need for a numeraire in any system of quantitative benefit-cost analysis (BCA). We are adding up the benefits and the costs of a project in each period, taking the difference between them, and finally transiting time by the use of the real social discount rate. Obviously, we need a common unit of measurement covering all the different types of benefits and costs. We believe that there are only two plausible candidates for this - the

GDP deflator and the consumer price index. That is, we should measure our benefits and costs either in production baskets or in consumption baskets.

A debate of the 1960s provides a convenient starting point for our journey. On the one side of this debate were advocates of the opportunity cost of investment, which they took to be the marginal productivity of capital (gross of taxes but net of depreciation) in the private sector. On the other side were partisans of a time preference rate of discount, which they took to be the after-tax rate of return to capital in the private sector.

This debate had what we consider to be quite a happy outcome, the outgrowth of a sort of reinterpretation of the concept of opportunity cost. Rather than thinking of the opportunity cost of public funds as their “best” alternative yield, this reinterpretation looked upon it as their “likely” alternative yield. Thus, whereas the prior argument had been between the marginal productivity of private capital (ρ) and its rate of return (r) after all relevant taxes, this new consensus took the opportunity cost to be a weighted average of the two $\omega = f_1\rho + f_2r$. Here f_1 and f_2 represent “sourcing coefficients”, f_1 being the fraction of new investible funds coming at the expense of displaced (potential) private investment, and f_2 being the fraction of such funds coming at the expense of private consumption.

There were initially a few who argued for deriving f_1 and f_2 from tax-based scenarios, but such efforts foundered on the facts that each tax change involved a different set of sourcing coefficients, and that changes in tax laws were quite infrequent and rarely similar from event to event. Rather quickly it was recognized that the “natural” source of funds for incremental investment or other spending was the capital market. Government borrowing was like a sponge—it released funds when they were demanded and absorbed them when they were supplied. With capital market sourcing, f_1 represented investment demand that was squeezed out, and f_2 reflected incremental saving that was stimulated by an incremental demand for funds. We defer to a later point the discussion of f_3 , the fraction of increased borrowing met by inflows from abroad. This was not discussed in the early debate.

Capital market sourcing fits quite neatly with the way that supply and demand are treated in other compartments of benefit-cost analysis. Any new demand represents a rightward shift of demand for the item in question. This stimulates a rise in price and a movement forward (f_2) along the supply curve of the product, as well as a movement backward (f_1) along the demand curve representing other (non-project) users of the product. If there are no distortions in the market for the product, the market price (P) prevails for both suppliers and other demanders, so weighted sourcing yields $P = f_1P + f_2P$, and the market price is the appropriate measure of opportunity cost.

But when taxes and/or subsidies are present, the price paid by demanders (P^d) will differ from the price received by suppliers (P^s), and we have economic opportunity cost $EOC = f_1 P^d + f_2 P^s$. This formula can be (and is in practice) applied to any non-tradable good or service, to tradable goods generally (as an aggregate), and to the demand and supply of any well-defined class of labor in a given labor market area. Thus the idea of capital market sourcing is simply an extension of the way other markets are handled in benefit-cost analysis.

When we contemplate the coefficients f_1 and f_2 , we have to bear in mind that in benefit-cost analysis we are not considering “our” project or program as necessarily “new”, being introduced into an “old”, preceding situation. Something of this flavor may indeed be true at the very outset, but many BCAs deal with project or program lives of 30 or 40 years or more. So, the profile of net benefits and costs that we analyze is really the difference between two moving pictures—one showing how the economy would evolve “with” our project or program, the other tracing a similar evolution “without” it. Thus, the flows of net benefits $F_t = (B_t - C_t)$ attributable to our project are typically all in the future. BCA usually precedes the first major investments by one to several years, and the end of the project’s life might be a further 50 years ahead, or else might be artificially compressed by assigning an arbitrary terminal date, and allocating to it a residual present value representing the sum of all subsequent F_t ’s, discounted to that date.

The standard BCA (and applied welfare economics generally), in dealing with future flows, assumes that the economy follows an equilibrium path with the project and an alternative equilibrium path without it. Thus, certainly for major projects and programs, we are concerned with full-equilibrium adjustments on both (with/without) sides.

Weighted-Average Methods

The development of the weighted-average approach rendered it unnecessary to choose between ρ and r as the relevant discount rate, and many (we believe most) analysts were content with this solution. Some, however, making bows to the theory of consumer choice that we teach in our classrooms, held to the position that the actual discounting should be done using the time preference rate, r . Much to their credit, however, they at the same time recognized the need to account for the loss in production from displaced investments. They made allowance for this via a shadow price of investible funds. A numerical example is probably the easiest way to show this. Suppose the gross-of-tax rate of return on investment, ρ , equals 12 percent, and the net-of-tax return on savings, r , is 4 percent. Suppose, too, that when \$1 million of funds are raised in the capital market, this displaces \$750,000 of investment and \$250,000 of consumption. Weighted-average people would derive from this information an economic opportunity cost $\omega = 10\% = (0.75)(12\%) + (0.25)(4\%)$. The time-preference people, on the other hand,

would still want to use 4% as their discount rate, but would recognize that in displacing 750 of investment, the economy was forgoing a steady future flow of 90 ($= 750 \times 12\%$) per year, and was incurring an annual cost of 10 ($= 250 \times 4\%$) on payments covering the supply price of the new saving that helped finance the project. The present value of these future flows, discounted at 4%, is 2500 [$=(90+10)/0.04$], and the shadow price of investible funds, (SPIF) is 2.5 ($=2500/1000$).²

While the debate will probably never end, the fact is that most projects that would be rejected by applying a straight 10% discount rate would also be rejected by applying a 4% rate and then punishing the project with a shadow price of 2.5 for investible funds. In our specific case of a steady perpetuity, net present value is zero using an investment cost of 1000 and a discount rate of 10% applied to a perpetual net benefit of 100. But NPV is also zero using a 4% discount rate for the perpetuity and a shadow price of 2.5 for the investment of 1000.

² This assumes that either the displaced investment was a perpetuity, or (much more plausibly) that the assets in question would be replaced as they depreciated by new ones of similar yields (i.e. using depreciation allowances to pay for them rather than tapping the market). Thus a sequence of one-year assets costing 1000 in year T and generating 1100 in year T+1 would yield the same net profile (-1000, +100, +100, ...) as a single investment of 1000 with a perpetual yield of 100. Obviously, any pattern of investment flows over time can be simulated by a sequence of one-year flows, so our conclusion applies so long as depreciation allowances are reinvested. Note that this reinvestment issue applies not to project's own investment profile but rather to the investments that are displaced by new demands in the capital market.

We accordingly want to consider as close allies those who opt for a weighted average discount rate of ω with a SPIF of one, and those who opt for a discount rate of r with a SPIF of ω/r . The real enemies of sound economics are those who press for the use of low discount rates like r , without due recognition of the costs entailed when forgone investments would have had rates of marginal productivity much higher than r .

Measuring the Marginal Productivity of Investment (ρ)

The discount rates actually used in official and semi-official BCAs span an extremely wide range, so we are well advised to inquire into how they ought to be estimated. To begin this exercise, we return to the classroom example of the previous section. That case can be made to look almost magical. Suppose the private rate of return (ρ) of 12% is subject to corporation and property taxes taking half, leaving a private rate of return (i) of 6%, which in turn is subject to a personal income tax of 1/3, leaving a time preference rate (r) of 4%. In this idealized picture the market interest rate (i) equals the rate of return on equity and also that on bonds, including government bonds (i_g). Here we have $i = i_g + D_b$, and $r = i + D_p$, with the business distortion D_b equal to 6% and the personal distortion D_p equal to minus 2%. We can then show that the weighted average discount rate $\omega = f_1\rho + f_2r$ is equal to $f_1 \cdot (i + D_b) + f_2 \cdot (i + D_p)$ which in turn, under our assumptions, equals $i_g + f_1D_b + f_2D_p$. That is, we can get to the “right” discount rate ω by two

routes. One takes the weighted average of the productivity rate ρ and the time preference rate r . The other takes the (real) government bond rate i_g and adds to it the weighted average distortion. The first route gives us $(0.75)(12\%)+(0.25)(4\%)=10\%$; the second yields $6\%+(0.75)(6\%)+0.25(-2\%)=10\%$. Two routes, same answer! But real-world cases look nothing like this magical example. The real rate on government bonds is for most countries quite low. Adjusting for tax distortions surely raises it, but not nearly by enough to match the weighted average of ρ and r .

So we have to struggle with the discrepancy and try to devise a sensible way to deal with it. We believe the most meaningful way to approach this problem is to go directly to the task of measuring ρ and r . We would like to reach results that have relevance for many individual countries as well as international organizations and national foreign aid agencies. So, we are going to build our analysis on four parameters—the country's average rate of GDP growth (g) in recent years, its average ratio (s) of investment to GDP, the share (a) of return to capital in the country's GDP, and the percentage (δ) that true economic depreciation bears to the country's stock of reproducible capital. With these four parameters and a little work, we can derive a plausible range of estimates for ρ , the rate of return on the country's reproducible capital.

The starting point for this exercise is the concept of equilibrium growth, in which the stock of reproducible capital is growing at the same rate as GDP (Y). We will later probe the sensitivity of our results to deviations from this assumption. But staying with it for now, we have:

$$\text{Gross Investment} = I_t = sY_t = (g + \delta)K_{t-1}$$

$$\text{Gross Rate of Return} = (\rho + \delta) = \frac{\text{return to capital}}{\text{stock of reproducible capital}}$$

$$(\rho + \delta) = \frac{aY_t}{(sY_t)/(g + \delta)} = a(g + \delta)/s$$

$$\rho = \left[\frac{a(g + \delta)}{s} \right] - \delta$$

Table 1: Plausible Parameter Ranges by Country Type

Country Type	Advanced	Healthy Developing	Asian Tigers
Gross Investment/GDP	0.15 to 0.20	0.20 to 0.25	0.25 to 0.35
GDP growth rate (g)	0.02 to 0.04	0.04 to 0.06	0.06 to 0.08
Capital's share in GDP (a)	0.33 to 0.40	0.40 to 0.50	0.50
Average depreciation rate	0.04	0.04	0.04

Readers should note that for these three groups, we run the gamut from 0.33 to 0.5 for a, from 0.02 to 0.08 for g, and from 0.15 to 0.35 for s. Also, we would stress

that the fraction “a” reflects gross-of-depreciation income from capital, and that “s” reflects gross investment, both as a fraction of GDP.³

One question not yet addressed is, how do our results change when a country deviates from equilibrium growth? The answer is that the formula for ρ changes to

$$\rho = a \left[\frac{(g+e+\delta)}{s} \right] - \delta.$$

This derives from a growth rate of reproducible capital equal to $g+e$, compared with GDP growth of g . Thus, capital growing 1 percent faster than output would add 0.02 to ρ if a was 0.50 and s was 0.25.

Table 2 presents the estimated real rate of return for the three types of countries. The middle level is simply the midpoint of the resulting range. The last cell in each row gives the adjustment to the rate of return that would result if the capital stock was growing one percentage point faster or slower than the GDP.

Table 2: Real Rate of Return to Capital (ρ_0), as measured from national accounts data

Country Type	Low	Middle	High	$\pm 0.01(ae/s)$ low, high
Advanced	0.09	0.105	0.12	$\pm 0.022, 0.2$
Healthy Developing	0.12	0.14	0.16	$\pm 0.02, 0.02$
Asian Tigers	0.13	0.145	0.16	$\pm 0.014, 0.02$

³ The oft-cited one quarter share of capital in some advanced countries usually refers to net income as a fraction of national income, the same figure for depreciation being excluded from both numerator and denominator.

Table 2 presents the results one gets working, respectively, with the low and high ends of the ranges presented for each country type. For the Asian Tigers, the higher return of 16% comes from the lower initial figures ($g=0.06$ and $s=0.25$).

Readers are warned not to treat these numbers as giving the rate of productivity we need for calculating the social discount rate. Rather, they should be taken as the basis from which we will work, as we subject them to a series of adjustments. In our next calculations, we make two types of adjustments, one affecting the size of the relevant capital stock, the other affecting the size of the income stream attributable to that stock.

Adjusting the Capital Stock

First, let us note that estimates of the capital stock flowing from the national accounts include improvements to land, but not the land itself, so to bring capital income and capital stock into accord, we must either add the value of land to the denominator of ρ , or subtract the income from land from the numerator. The latter approach is clearly to be preferred, because (a) it is incredibly difficult to arrive at good estimates of the total market value of land (net of improvements) and (b) the annual GDP figures capture only part of the return to land, the other part being generated by the real appreciation stemming from economic growth. So our present exercise simply starts with the stock of reproducible capital (fixed capital plus inventories) as the denominator of our rate of return.

The principal adjustment to be made to this stock concerns infrastructure investment of various types. Items such as roads, school buildings, parks, public offices and museums certainly yield benefits, but these are rarely, and if so only very partially, reflected in our measures of GDP. They are not part of the capital whose return we are trying to measure. An adjustment is therefore necessary in order to exclude these items from the base of the return-to-capital ratio. Across nations, the fraction of GDP devoted to public investment varies greatly, but it is also true that countries with high ratios of public investment to GDP tend to do more infrastructure investment than those with low ratios. We feel we are if anything overestimating the size of the stock of “remunerative capital” when we take it to be 80 percent of the stock of reproducible capital that we get from the national accounts. With this adjustment, our estimate of ρ increases by a factor of 1.25, yielding the figures shown in Table 3.

Table 3: Real Rates of Return (ρ') to Remunerative Investment
(as measured from national accounts data)

Country Type	Low	Middle	High
Advanced	0.1125	0.1312	0.15
Healthy Developing	0.15	0.175	0.20
Asian Tigers	0.1625	0.1812	0.20

These will be seen as high numbers by laymen as well as professionals. But we must insist that they represent flows of real money (or its equivalent in imputed rents). Real people (or real governments) are enjoying these real benefits, on the basis of the productive assets that they own. In the next section, we will peel off of this return several components which we feel should not contribute to our estimates of ρ and r . But here we want to stress the importance of “starting from the top” -- i.e., from the total returns to capital that are measured in the GDP.

Our exclusion of 1/5 of investment as non-remunerative has equivalences in a whole gamut of patterns of public investment. We have estimated that ρ' (the measured rate of return on private investment) is equal to our previously measured ρ_0 (Table 2) times 1.25. Thus is true when government invests 0.2I in infrastructure, with a yield of zero, and the private sector invests 0.8I, with a yield of $1.25 \rho_0$ (I = gross investment from the national accounts). Obviously, the same result emerges if state-owned enterprises in any amount participate alongside the private sector, and also earn $1.25 \rho_0$. But in the real world, infrastructure investments often generate some low rate of yield ρ_i , and state-owned enterprises usually generate a rate of return $\rho_s < 1.25 \rho_0$. Our estimate, that the private sector rate of return is $1.25 \rho_0$, can be generated by any combination of h_i, h_s, h_p and $\rho_i, \rho_s, 1.25 \rho_0$ that fits the equation

$$\rho_0 = h_i \rho_i + h_s \rho_s + h_p (1.25) \rho_0.$$

Here h_i, h_s, h_p are the fractions that infrastructure, state-owned enterprises and private investment occupy in total I.

The high figures of Table 3 come in principle from national accounts data via the coefficients g, a, s and δ . More precise estimates can be obtained by generating capital stocks of buildings, machinery and equipment, vehicles and inventories, using perpetual inventory procedures with separate depreciation rates for each class (zero for inventories). Harberger participated in such exercises in Mexico, El Salvador and Honduras, which yielded rates of return around or above 20 percent, even without making any allowance for infrastructure investment. These results only add to our confidence in the ranges presented in Table 3.

But we are not suggesting that a country's actual discount rate should be based on the numbers in Table 3. They are subject to further adjustments, to which we now turn. These will help to answer Harberger's plaint. "Why can't I get 20 percent returns on the dollars I might send to Mexico?"

Adjusting the Returns to Capital

In this section we explore four components of the measured return to capital that should be eliminated from the numerator of our rate-of-return calculation, if our

objective is to reach measures of ρ and r to be used in calculating the social discount rate. They include:

- a) Returns to land included in the GDP
- b) The share of labor in monopoly markups.
- c) Above-normal returns on supra-marginal investments
- d) Some portion of returns to capital in financial intermediation

a) Returns on land included in the GDP

Here we suggest assigning to land something like 1/3 of the value added in agriculture and something like 1/10 of the rental income (actual or imputed) from residences and commercial buildings. For agricultural value added we suggest 3% of GDP in the advanced countries, 9% of GDP in the Asian Tigers and 15% of GDP in the healthy developing countries. For the land component of residential and commercial rents, we suggest 1 percent of GDP.

b) The share of labor in monopoly markups

It seems that only economists are aware that monopoly markups are not in principle part of the return to capital but instead are the return to the “monopoly right” conferred by market power. The value of the marginal product of each factor exceeds that factor’s cost by the percentage monopoly markup, μ . μ can be regarded as a privately imposed and privately collected tax on all elements of

costs. This gives rise to complications via input-output connections. Luckily, these can be avoided using the simple assumption that material inputs enter their respective final products in fixed proportions. Then, for any product i , applying to value added (V) a markup μ'_i defined as $C_i\mu_i/V_i$ produces a result that is economically equivalent to what we get from applying μ_i to total costs C_i . Obviously, the fixed proportions assumption is an approximation to reality, but we believe it to be a quite good (and an extremely useful) one.

We are pretty much in limbo trying to estimate the size of monopoly markups. We should realize, however, that they are by definition something over and above the normal return to capital, and something that in principle does not vary greatly from year to year. Monopoly enterprises should gain the advantage conferred by their market power year in and year out. We suggest an allowance for monopoly markups equal to one fifth of the total return to capital (after adjustment a). This, however, is not all to be deducted from capital's return. Under our assumption the markup μ' raises the marginal productivity of labor by μ' percent over its cost, and similarly for the marginal product of capital. The marginal product of capital is what the economy loses when investment is displaced, so we should deduct from capital's gross earnings only labor's share of one fifth of the stated return.

c) Supra-marginal Investments

When we build models using production functions like $Y_i = \phi_i(L_i, K_i)$, any increment to labor or capital is assumed to itself be marginal and is thus assigned the marginal product generated by the production function. We know, however, that this does not describe what happens in the real world. In actuality, new investments appear each year which generate returns that are well above average. In our classrooms, we sometimes illustrate this with a bar chart, with invested amounts on the horizontal axis and expected real yield on the vertical axis. Individual investment projects are then arrayed in descending order of their prospective internal rates of return. This, on the assumption of independence among their respective investment profiles, gives us a priority ordering for the projects. Normally, each firm will carry its investment to the point where the projects it rejects all have prospective yields below that firm's relevant cost of capital. Thus, the "triangle" representing supra-marginal yields really does not belong as a component of the opportunity cost of investment. We would emphasize, however, that not all high-yielding investments are safe from being squeezed out when funds are tight. Small businesses are particularly vulnerable at such times. They often have substantially higher cost of capital even in good times as well as being squeezed disproportionately when funds get tighter.

So how do we go about estimating the amount by which capital's return should be reduced (for our purposes) due to supra-marginal investments? Not, in this case, by charging directly through the line, but instead by making an end run around it, to which we now turn.

Economists are well acquainted with the growth equation

$$g_y = S_L g_L + S_k g_k + \lambda$$

g_y, g_L, g_k are the growth rates of output, labor and capital, respectively, and λ is the rate of real cost reduction (=TFP improvement =“technical advance”). But this equation has a twin, a “dual”, about which many are quite unaware:

$$\lambda = S_L g_w + S_k g_{(\rho+\delta)} - g_P$$

This dual says that the fruits of real cost reduction have to go somewhere, and the allocation is that labor gets its share (S_L) times g_w , the growth rate of real wages, consumers benefit from any fall in prices (or suffer from any rise) and capital's return grows to absorb whatever benefit of real cost reduction has not been gobbled up by labor and consumers. This explanation makes it clear that capital is the residual claimant in this story. For a competitive producer, the market determines

what wages have to be paid and what the price of the product will be. So the capital factor gains only to the extent that λ exceeds $S_L g_w - g_P$.⁴

For the economy as a whole, with the GDP deflator as the numeraire, g_P is obviously zero in the growth equation. This enables us to calculate rates of aggregate real cost reduction implied by the various parameter packages that we have assumed. The only further datum that we need is the rate of growth of the labor factor. The correct measure of the change in labor inputs is $\sum_i w_i \Delta L_i$, where increments to labor are weighted by their corresponding initial wage levels. This automatically incorporates the contribution of human capital accumulation by weighting the higher-paid components of the labor force more heavily as, over time, they include an ever higher fraction of the labor force. Our estimates of λ are based on the assumption that raw labor (N) grows at one percent per year in the advanced countries and the Asian Tigers, and 1 ½ percent per year in the healthy developing countries. To these figures we add ½ percent per year for improvement in labor quality in the advanced countries, and one percent per year in the other two categories. See Table 4.

⁴ In a competitive industry the return to capital is driven to its long run competitive level. Thus $S_k g_{(\rho+\delta)}$ tends to zero, and the equation becomes $\lambda = S_L g_w - g_P$. Thus in the “average” industry $\lambda = S_L g_w$ (i.e. wage growth is determined by the average rate of TFP improvement), and its relative price does not change, while relative prices will tend to rise in industries with lower than average real cost reduction, and to fall in industries with λ above average.

Table 4: Implied Rates of Real Cost Reduction

$$[(\lambda = g_y - S_L g_L - S_K g_K = S_L(g - g_L) \text{ when } g_y = g_K = g)]$$

	g	g_L	S_L	λ	(g_N)
Advanced Countries	0.02	0.015	0.67	0.0033	(0.01)
	0.04	0.015	0.6	0.015	(0.01)
Healthy Developing Countries	0.04	0.025	0.6	0.009	(0.015)
	0.06	0.025	0.5	0.0175	(0.015)
Asian Tigers	0.06	0.020	0.5	0.020	(0.01)
	0.08	0.020	0.5	0.030	(0.01)

Note: g_N is the rate of growth of the number of workers, g_L is the rate of growth of labor hours worked, adjusted for labor quality

Our view on this subject is that when real cost reductions occur, the enterprises involved enjoy supra-marginal returns for a while. These, however, are eroded with time, and rates of return go back to the competitive level (recall that we have already dealt with monopoly profits (which do not go back, and which may in part result from real cost reductions). We feel that allocation of supra-marginal returns to capital equal to $\lambda_t + 0.75 \lambda_{t-1} + 0.5 \lambda_{t-2} + 0.25 \lambda_{t-3}$ is likely to err on the generous side.

To calculate supra-marginal returns, we assume that capital receives the full real cost reduction (λ_t) of time t, plus $\frac{3}{4}$ of λ_{t-1} , plus $\frac{1}{2}$ of λ_{t-2} , plus $\frac{1}{4}$ of λ_{t-3} . A

given innovation would receive its reward in the year it was introduced, but this would erode in the following 3 years. The added return generated by the innovation might be 4 points in t , 3 points in $t+1$, 2 points in $t+2$, 1 point in $t+3$, and zero thereafter. $g_{(\rho+\delta)}$ is therefore positive only in the first year for this particular innovation, and becomes negative in the next 4 years. Comparing year $t+4$ with year $t-1$, $\Delta(\rho + \delta) = 0$. What can we make of this? The innovation has in fact produced extra returns, accumulating 10 percentage points on the firm's value added over 4 years. But these extra earnings are part of what we have allocated as the return on capital. This is the part we want to exclude from our calculation of capital's marginal rate of return.

At the aggregate level in a real-world exercise, we would do growth accounting and calculate $\lambda_t Y_{t-1} + 0.75 \lambda_{t-1} Y_{t-2} + 0.5 \lambda_{t-2} Y_{t-3} + 0.25 \lambda_{t-3} Y_{t-4}$, and end up excluding that result as we adjust our original measure of the return to capital. In our present exercise, we set λ equal to the calculated level shown in Table 4, which is implied by equilibrium growth under the specified assumptions. The adjustment we make for each year excludes from our measured return to capital an amount equal to $2.5 \lambda Y_t^5$.

⁵ The growth model tells us that wage rates will rise by λ/S_L per year. This is fully compatible with the innovators of each year gaining λY_{t-1} , because in that same year the innovators of the previous four years are losing $\frac{1}{4} \lambda Y$. So a sort of churning takes place among cohorts of innovators, the amount of $2.5 \lambda Y$ being passed each year from earlier cohorts to their successors. We have chosen a very simple allocation

d) Allowance for investment-related costs of financial intermediation

When new demands for funds lead to displacement of other investments, they automatically save the economy the intermediation costs that would normally be linked to those investments. We are not here concerned with the labor and material inputs that are thus freed up; they are presumed to render their marginal product elsewhere in the economy. However, in measuring the returns to capital for the economy as a whole, we obviously have included such returns that would be received by capital in the financial sector. What we would like to do here is exclude that part of these returns that are linked to the investments of each period. This certainly poses a real challenge to the estimator, and we are reluctant to claim that we have a genuine answer. Rather we would consider two possible ways of making an allocation for the saving of intermediation costs. The first is based on the estimated amount of investment of each year, and the second is tied to capital's share in the value added of the financial sector. Our suggestion is that the allowance should be the lower of a) 6 percent of gross investment of the year when funds are taken from the capital market, b) half of the return to capital involved in

of the fruits of innovation. The same general conclusions would follow, however, from more complicated ones, as long as the extra rewards to innovations follow a humped shape. For example, the extra rewards from innovations might grow over three years, and only then drift downward to zero. That would make our calculation more complicated, but the main conclusion would be the same. The innovators of this and possibly prior years would together reap a total sum which would be a constant fraction of GDP in the growth model. But this sum would be churned each year between those newer innovators whose p 's would be rising and those earlier innovators whose p 's would be declining back to their normal levels.

financial intermediation. In real-world cases both of these would be derived from actual data from the period in question. In the case of our equilibrium growth examples, only the first can readily be applied. So the allowance in this case would be $0.06sY$. We want, of course, to exclude only that part (=capital's share of this) which has been counted as part of capital's return. But we recognize that the financial sector is more capital intensive than the rest of the economy, so in calculating the exclusion we set capital's share equal to 0.67 in the financial sector.

Table 5: Adjusting capital returns to derive estimates of ρ

	Advanced		Healthy Developing		Asian Tigers	
1) Capital's Return	0.33Y	0.40Y	0.40Y	0.50Y	0.50Y	0.50Y
2) Adjustment for Land	0.020Y	0.020Y	0.06Y	0.06Y	0.04Y	0.04Y
3) Difference [(1)-(2)]	0.310Y	0.380Y	0.340Y	0.440Y	0.460Y	0.460Y
4) Monopoly markup [$0.2*(3)$]	0.062Y	0.076Y	0.068Y	0.088Y	0.092Y	0.092Y
5) Labor share * (4)	0.042Y	0.046Y	0.040Y	0.044Y	0.046Y	0.046Y
6) Difference [(3)-(5)]	0.268Y	0.334Y	0.300Y	0.396Y	0.414Y	0.414Y
7) Supra-marginal return [$2.5*\lambda*Y$]	0.0083Y	0.0375Y	0.0225Y	0.0438Y	0.05Y	0.075Y
8) Difference [(6)-(7)]	0.260Y	0.297Y	0.278Y	0.352Y	0.364Y	0.339Y
9) Investment	0.15Y	0.20Y	0.20Y	0.25Y	0.25Y	0.35Y
10) Intermediation Cost [$0.06*(9)$]	0.009Y	0.012Y	0.012Y	0.015Y	0.015Y	0.021Y
11) Capital's Share [$2/3*(10)$]	0.006Y	0.008Y	0.008Y	0.010Y	0.010Y	0.014Y
12) Capital's adjusted return [(8)-(11)]	0.254Y	0.289Y	0.270Y	0.342Y	0.354Y	0.325Y
13) Unadjusted rate of return to remunerated capital (from Table 3)	0.1125	0.15	0.15	0.2	0.1625	0.20
14) Estimate of ρ [$((13)*(12))/(1)$]	0.086	0.108	0.101	0.137	0.115	0.130
15) Mid-point of ρ by country type	0.097		0.119		0.122	

Table 5 details the calculations needed in order to reach our estimates of ρ . We certainly do not want to claim accuracy of these estimates, but we do want to

defend the methodology by which they were derived. We have made these estimates on the very rough basis of assumed high and low parameters for different classes of countries. We believe these parameters very likely span the true values in most cases. But we definitely would advise that anyone seeking a good estimate of ρ for any given country should try to get specific values of key parameters like a , g , δ , s in that country, and should also explore the possibility of disequilibrium growth via the parameter e . Where possible, direct estimates of the stock of reproducible capital should be made, accumulating gross investment in different classes of capital assets by the perpetual inventory method. One important detail—when building capital stocks with the aim of calculating real rates of return, each year's product and gross investment flows should be deflated by the numeraire (the GDP deflator or the CPI) rather than by two different indexes as is done in the national accounts. If separate deflators were used, say, for the textile industry, the ratio of income to capital would really turn out to be yards of cloth per loom. Only by having the same deflator on both sides of the calculation can we come to a proper estimate of the rate of return. By the same token, it is advisable to derive the ratios a and s from the nominal side of the national income accounts rather than those appearing on the real side.

Moving from ρ to r

In many treatments, the time preference rate is taken to be that on government bonds or time and saving deposits. We feel that this is a mistake, and that the rate of return to savers should be estimated from the average adjusted rate of return to reproducible capital (ρ), minus all income and property taxes that end up being paid out of that income. Indirect taxes should not be deducted, because they do not distort intertemporal choices and thus do not impact the marginal rate of time preference.

We struggled with the idea of approaching the estimation of the tax wedge through the application of marginal tax rates, but decided against it. Capital income accrues at the margin to entities that are exempt from certain taxes; it also accrues to tax evaders. Moreover, many special tax dispensations apply at the margin. Finally, it was also much more convenient to work with average tax rates, which neatly fit into our national accounts framework. We proceed by working with estimates of the share of income and property taxes in a country's GDP. Tax structures vary widely among countries, even within our three groupings. But one generalization can be made-- among the advanced economies, those with higher ratios of total tax to GDP tend to rely more heavily on value added and other indirect taxes, while

some of those with lower total tax ratios get 50 or even 60 percent of their total from income and property taxes.

Luckily, the International Monetary Fund provides data on “Taxes on Income, Property and Capital Gains” as a fraction of GDP for many countries. For the industrial countries, these figures cluster quite tightly around 15 percent. The ratios in the developing countries have wider variation, but their average is quite stable around 6 percent. The Asian Tigers present a more complicated picture, with rates that can be as low as 5 percent and as high as 15 percent. We will use 10 percent of GDP as our figure for the Asian Tigers.

We have up to now talked about income and property taxes paid out of the income from capital. The information from the IMF includes capital gains as well. Some of these capital gains are reflections of corporate earnings which, of course, appear in the income from capital as presented in the national accounts. But some capital gains (e.g. those on land, residential housing, works of art, etc.) are not counted as part of the GDP. Hence, our estimates based on the IMF information probably overstate the true taxes actually paid on our measured aY . But any overstatement is unlikely to change the order of magnitude of the rough approximations we use.

Because corporation and property taxes strike mainly the income from capital, it is reasonable to assign $1/3$ of these figures to labor income, and $2/3$ to capital income.

Note that the personal income tax is by far the largest component of the figures on which we based our entries, and that in most countries the bulk of the income tax on salaries is paid by the upper deciles at high marginal rates. So it is reasonable to assume that at least a third of income and property taxes do in fact fall on labor income.

Table 6: Transiting from ρ to r

	Advanced		Healthy Developing		Asian Tigers	
16) ρ [from (14)]	0.086	0.108	0.101	0.137	0.115	0.130
17) Capital's gross return [(1)]	0.333Y	0.4Y	0.4Y	0.50Y	0.50Y	0.50Y
18) Capital adjusted return [(12)]	0.254Y	0.289Y	0.270Y	0.342Y	0.354Y	0.325Y
19) Income and property taxes	0.15Y	0.15Y	0.06Y	0.06Y	0.10Y	0.10Y
20) Capital's share (2/3) of (19)	0.10Y	0.10Y	0.04Y	0.04Y	0.067Y	0.067Y
21) Taxes on adjusted capital [(20)*(18/17)]	0.076Y	0.072Y	0.027Y	0.027Y	0.047Y	0.043Y
22) Adjusted capital income after tax [(18)-(21)]	0.178Y	0.217Y	0.243Y	0.315Y	0.307Y	0.282Y
23) Adjusted net rate of return [(16)*(22/18)] = r	0.060	0.081	0.091	0.126	0.098	0.113
24) Midpoint of r by class	0.070		0.108		0.105	

Treatment of Risk

Readers will probably be quick to notice that the rates that emerge for r from the preceding calculations are significantly above the real rates of return on government bonds, savings accounts, etc. One important source of the difference is

the fact that the rates of return that prevail in the capital market reflect the attitude of its participants with respect to risk. Please note that we are not dealing here with default risk, which concerns the gap between promised returns and those that actually end up being paid. Our rates of return are based exclusively on the payments that savers receive as recorded in the national accounts. An important component of their difference from government bond rates is the so called equity premium dealing with rates of return in the stock market. But our figures also include the returns on non-corporate investments of all types.

Our defense of our calculated rates for r -- or, perhaps better put, of the general methodology underlying our calculations -- is not based on the subtleties of portfolio analysis, but rather on the design and use of benefit-cost analysis itself. As we see it, if we were to use a discount rate (such as that on US TIPS) which is close to riskless, we would be obliged to make separate allowances for the riskiness of every project or program that we deal with. Such a prospect utterly terrifies us. How does one deal with the risk associated with a new school, or road, or airport, or irrigation dam? What about programs of urban renewal or medical research or pollution abatement? Between us we've accumulated more than a century of practical experience in BCA work, yet we are still at a loss when it comes to dealing with issues of risk in such cases. We certainly have no objection to people carrying forward the search for a robust and convincing methodology for

incorporating risk into our project-by-project and program-by-program work, but we're pretty sure that it will be a long time before the profession will be widely implementing such a methodology.

But we do have a way of reducing the size of the problem – namely, use a rate of discount that incorporates an allowance for average risk. That is exactly what we get when we base our analysis on national accounts data. Using such a rate, we surely reduce the size of the challenges of dealing with risk, and end up feeling a lot less abashed by our inadequacy at dealing with them. We can take comfort in the fact that many of the projects and programs that we deal with probably have close to average risk attributes. That is certainly better than we get working with a riskless rate, which fits for no practical application whatever, and which thus cries out for us to take specific steps to incorporate risk into every single BCA. When we embed average risk into our approach, we can still worry about cases that have quite obviously above-average or below-average risks. But we feel less troubled by trying to deal with such cases on an ad hoc basis, than we would be if we were using a riskless rate, and were thus called upon to make specific risk adjustments in each and every case.

We have flirted over the years with an alternative way of dealing with risk—the simple assumption that “society” is neutral with respect to risk. The argument here is quite strong. For a family, the incentive to insure a home is much stronger than

the incentive to insure a television set. That is because the loss of the home causes a big percentage drop in a family's wealth, while the loss of a television set causes barely a ripple. Along a utility-of-wealth function the loss of a home generates a chord between the "with" and "without" points, generating a true willingness to pay for insurance. In the case of the television set, the chord generates practically no willingness to pay. Its slope is close to that of the tangent at the "with" position. We believe that most projects, even large ones like Hoover Dam or the TVA or a major airport, represent a tiny fraction of society's wealth, and accordingly do not generate a serious cost linked to society's risk.

Following this line, society would not want to consider any part of the difference between p and r as a necessary compensation for risk. That difference would simply be a social benefit that is forgone when investment is displaced. Such an attitude would help explain why many countries choose to have society bear the costs of catastrophic medical events rather than leave them to the victims' families.

Neutrality with respect to risk makes the choice of the social discount rate quite easy, but we prefer the approach of incorporating average risk into the discount rate. It carries fewer overtones and implications, while leading to a quite similar result.

Readers should also bear in mind that the excess of our measured r over the real government bond rate is not solely, and perhaps not even mainly a reflection of

risk. If it were just that, savers would be indifferent between higher yielding investments and government bonds. That may be true for some (who hold both classes of assets). But the majority of savers probably hold no government bonds at all. Some of these are truly neutral with respect to risk, and therefore the expected yield should reflect their pure marginal rate of time preference. Many other savers are somewhat risk averse, but significantly less so than the market as a whole. For them, risk would account for only a fraction of the difference between their expected yields and the government bond rate. Research concerning the equity premium suggests that this is the case for a wide range of market participants—its main finding is that provision for systemic risk accounts for well under half of the “equity premium”. For us, the size of this fraction doesn’t really matter – we want to work with the supply price of savings, inclusive of whatever risk premia are we incorporated in it.

Issues of Time Preference

We have pointed out that many early advocates of the use of r as the basic discount rate (combined with a shadow price of investible funds) reached that decision on the basis of the theory of consumer behavior with respect to intertemporal decisions. The use of the weighted average ω takes time preference into account via the rate of return received by the savers. There are many categories of savers, receiving many different rates of return. So our f_2r is really $\sum_i f_i r_i$ where i varies

over all types of providers of funds, with different attitudes to risk, different supply prices of funds, all making choices so that the rate they expect to receive on each dollop of funds equals or exceeds their time preference rate for that application, and where the marginal rate they receive is very close to their time preference rate (i.e., the rate at which they discount after-tax benefit flows, measured in numeraire units).

We obviously take the approach we have outlined, and believe it to be the best choice for practical use. But we want to pay respect to an alternative line of thinking. This line would say that our $\sum_i f_i r_i$ pays attention to the time preference of those who supply funds to the capital market—but what about the others? Should not the whole population's time preference rates be represented in the social discount rate?

We see a certain logic in this way of thinking, and will therefore walk briefly down that trail. What we want to point out is that the trail does not lead to an especially low discount rate, as some might claim (or hope). The relevant marginal rate for many is a weighted average of the rates they receive on the various pieces of their investment portfolio. They are indeed the savers and the lenders. But a large fraction of the population in many countries consists of borrowers. In the advanced countries the marginal rate for a large fraction of the population is the rate on their

credit cards, which for them actually represents both the marginal source and the marginal use of funds. Another large fraction has the mortgage rate as its effective marginal rate of time preference. For developing countries, poor farmers and small business people pay extremely high rates to local money lenders and only occasionally free themselves from the burden of outstanding debts.

We do not know what would be the average of the marginal time preference rates of all the households in a given country, but we are quite confident that it would be well in excess of 2 or 3 percent in real terms, and we are happy to encourage any researchers who elect to explore this interesting path. But note that if one were to use such a broad average time preference rate (call it τ) for discounting, it should properly be accompanied by a shadow price of investible funds [e.g. $(f_1\rho + f_2r)/\tau$], to take into account the displacement of investment and stimulation of savings that occurs when a project is financed in the capital market.

This is a good place to emphasize that for most governments, all over the world, the capital market is indeed the marginal source or destination of funds. Governments make budgets once a year, projecting in some detail their expected outlays and receipts. But rarely does the outcome match the expectations. The actual deficit or surplus ends up larger or smaller, often quite a bit larger or smaller than what was originally budgeted. And nearly all governments, it seems, are

perennially in debt. Budget surpluses sometimes arise, so that outstanding debts are reduced, but only rarely are they reduced to the point where the government as a whole passes from debtor to creditor status. So we feel that capital market sourcing must be incorporated into standard cost-benefit analysis, either through the discount rate itself, as we prefer and suggest, or through a shadow price of investible funds.

A Time Preference Rate Linked to the Rate of Economic Growth

Students of economic growth are familiar with the Ramsey equation,

$$(1) \quad r = \gamma + eg .$$

Here r is taken to be the appropriate social rate of discount, γ is the rate at which future utility is discounted so as to be compared with present utility, g is the rate of economic growth, and e (defined as > 0) is the elasticity of utility with respect to consumption. When the Ramsey equation appears in the growth literature, γ is typically an arbitrarily assumed number, and e is typically assumed to be in a range something like $\frac{1}{2}$ to 2. A good example would take $\gamma = 0.02$, $e = 1$, and $g = 0.03$, leading to a social discount rate of 0.05. And if g happened to go to 0.04 the discount rate would shift upward to 0.06.

We feel that such exercises have it all backwards. The interest rates that are relevant for the marginal decisions of individuals are beyond their control, and they

modify their behavior accordingly. Consumption and saving respond to observed changes in real interest rates in a manner very different from that predicted by the Ramsey equation with a constant γ and a single utility function $U_t = F(C_t)$, which remains unchanged as t changes. This double constancy is what gives rise to the present value maximand of $\int F(C_t)e^{-\gamma t} dt$.

We are completely confident that the marginal rate of time preference varies greatly among people simply because each household faces its own interest rate (credit cards, mortgages, stock market, bank deposits, government bonds, etc.).

Moreover, the relevant market rate for any given household will move from one category to another as that household's fortunes change over the years. All of this is compatible with a marginal rate of time preference which varies (e.g. along a supply function for savings at any point in time) and a utility of consumption function that also can often be shifting over time. Economics can easily explain how, faced with a rise in interest rates, one family would decide to save more (in order to enjoy luxury living in retirement) while another would save less, in order to enhance present consumption while maintaining the same projected lifestyle during retirement. These different reactions would stem from different evaluations of the marginal utility of future consumption.

To convince oneself of the rigidity of the assumption of a constant utility discount rate γ and an unchanging function representing the marginal utility of consumption, we invite readers to simply explore some obvious examples. For example, consider a person who inherits a fortune. With a finite horizon and a constant $U_t = F(C_t)$ function, if $r > \gamma$, that person will choose to have the lowest consumption this year and will have ever higher consumption in each subsequent year. In general, g in the Ramsey equation is not an economic growth variable. Rather it represents the instrument by which maximizing individuals spread their consumption over time. In the same case with $r < \gamma$, the reverse occurs and individuals choose to have ever declining consumption over time. It is something of a mystery how g got to be taken as an economic growth variable rather than a solution to each family's allocation of its consumption over time. Actually in most cases, even without a windfall like an inheritance, if the relevant r exceeds a family's constant γ , that family will choose increasing consumption over time, while other families with $r < \gamma$ (either because their r is low or because their γ is high) would opt for decreasing consumption over time. Such powerful implications would divide the world into these two classes of families. Why should we buy into such a thought when simple, natural economics tell us that of course our marginal rate of time preference will vary with the relevant interest rate? Consumer maximization should always lead to an equilibrium in which the consumer is

indifferent between giving up 1 unit of consumption today and expecting $(1 + r)^t$ units of consumption t years from now. This is true if the function $U_t = F(C_t)$ is the same as $U_0 = G(C_0)$, or different from it, and it should also be true whether the relevant market interest rate is high or low. Using the Ramsey equation to derive an r that is different from the relevant market rate is, we believe, a mistake.

Incorporating the Sourcing of Funds from Abroad

We deferred to this point the discussion of the world capital market as a relevant source of funds. Our modified equation for the social discount rate is

$$\omega = f_1\rho + f_2r + f_3MCFF$$

where MCFF stands for marginal cost of foreign funds. The above equation can be tweaked so as to give us an easy solution to the social discount rate problem, simply by assuming that each country faces a flat supply curve of foreign funds -- e.g., an interest rate that remains unchanged no matter how much they borrow. If that were the case (an infinite elasticity of supply), then any added demand for funds would be met by new inflows from abroad, and the interest rate paid on such funds would be the relevant marginal cost of capital for the economy. No investment would be displaced, no additional saving stimulated.

We feel that there is a clear tendency toward equalization of short-term rates, especially among major financial centers where country risk is not a serious issue. But smaller countries do face rising country risk premia as they borrow more, and major players like the U.S. and the European Union actually have some degree of influence over the level of rates around the world. All this does not even consider the capital movements taking place in market segments other than short-term funds. In short, we do not believe that any serious empirical support can be mustered for the idea that new investments in a country are fully reflected in new capital inflows and new savings in new capital outflows.

Sadly, however, the capital flow picture is too complicated for us to derive from empirical evidence the supply curve of foreign funds that each (or even any) country faces. If we had that, combined with the nation's supply curve of savings and demand curve for investment funds, we could extract estimates of f_1, f_2, f_3 that would have solid econometric support. As things stand, we have to build on much rougher evidence and impressions. We know that investment is significantly sensitive to the availability of funds (except in crisis times), and that the rate of saving responds only modestly to interest rate changes. So f_1 should be significantly larger than f_2 . The size of f_3 is more of a problem. Clearly it will vary significantly from country to country, being close to zero for those with the poorest links to world capital markets. For some countries, clues might be gathered from a

study of movements in the net foreign balances of their private banks. For others one would have to factor in the behavior of their central banks, which often build up international reserves to sterilize some part of private inflows, and draw on those reserves to offset some part of capital outflows. For major financial and industrial centers, one might gain insights from variations in foreign ownership of balances and other assets.

The above should make it quite clear that we are not ready to come up, even with well-informed guesses as to the size of f_3 . The best we can do is to build on what we don't see. Except for genuinely outlier countries, we don't see an utter absence or near-absence of capital flows. So f_3 is not likely to be close to zero in most cases. Similarly, we don't see large and palpable sensitivity of net capital outflows and inflows to changes in interest rates and general levels of pressure or ease in their capital markets. So we feel it would be quite a stretch to assume that foreign sources provided as much as half of the funds for the typical new investment and absorbed as much as half of any autonomous increase in savings. So we will bite the bullet and choose 0.3 the figure that we will use for f_3 , making note that we would feel quite uncomfortable with an f_3 as low as 0.2 or as high as 0.4.

There remains the issue of the marginal cost of foreign funds. Here the underlying vision of an upward rising supply curve of such funds gives us a clear signal that

marginal cost should be above average cost, by a factor of $[1+1/(1+\epsilon_f)]$, where ϵ_f is the elasticity of supply of foreign funds facing the country. The latter is quite likely to vary substantially from country to country and from situation to situation for a given country. The figure we use for MCFF should also obey our maxim of being sensitive to what we do not see. We believe that if countries and their inhabitants could readily borrow large sums at a marginal cost of 2 or 3 percent in real terms, there would be a lot more international capital flows than what we actually observe. Similarly, if the marginal cost of foreign funds were something like 10 or 12 percent, we would see much less international movement of funds than what we actually witness. Figures like 4 to 8 percent seem to us to be more plausible. To lean to the conservative side, we will use a real rate of 6 percent for MCFF, partly for the reasons stated above, and partly because we would be uneasy if our final values for the social discount rate ended up being mainly governed by our assumed MCFF.

Calculation of the Social Discount Rate

The figures in Table 7 are based on the numbers we found for ρ in Table 5 and for r in Table 6, plus our assumed value of 6 percent for MCFF. We present two alternative weightings for f_1 and f_2 -- our preferred weighting of (0.6, 0.1), and an alternative weighting of (0.5, 0.2).

Table 7: Weighted Average Estimates of the Social Discount Rate (ω)

$$(\omega = f_1\rho + f_2r + f_3MCFF)$$

Country Type	Advanced		Healthy Developing		Asian Tigers	
ρ [from (14)]	0.086	0.108	0.101	0.137	0.115	0.130
r [from (23)]	0.060	0.081	0.091	0.126	0.098	0.113
MCFF	0.06	0.06	0.06	0.06	0.06	0.06
ω ($f_1=0.6, f_2=0.1, f_3=0.3$)	0.076	0.091	0.088	0.112	0.096	0.107
Mid-point of ω by class	0.083		0.100		0.101	
ω ($f_1=0.5, f_2=0.2, f_3=0.3$)	0.073	0.088	0.087	0.102	0.095	0.105
Mid-point of ω by class	0.080		0.094		0.100	

The figures resulting from our exercise are within the range of real discount rates that have been widely used. The World Bank and some of the regional development banks used a 10 percent cutoff rate for many years, and many developing countries have followed this path. The point to be made here is that many projects have passed the test of rates in this range. Many business firms, even in advanced countries, ask for expected internal rates of return of 15 and even 20 percent on the investment projects they undertake. So we should not be frightened by rates in this range. Recall that we derived them from national accounts data, with plausible adjustments that had the effect of lowering the calculated rate.

We also would call attention to the minimal sensitivity that we found to changes in the weighting scheme. This stems, of course, from our building allowances for

average risk into our estimate of r , but we feel extremely confident that we would only invite troubles, both conceptual and practical, if we tried to incorporate a riskless r into our estimate of the social discount rate.

In closing, we want to emphasize that we are much more interested in convincing others of the broad strategy and conceptual framework that we advocate. We hope that others will join in the search for more accurate values of the parameters involved. Sensitivity tests, too, should be more widely used to convey the level of confidence that we should place in our estimates.

APPENDIX

We would like to comment separately on the influence of cyclical fluctuations in applications of benefit-cost analysis. On the labor side, when labor is drawn from the ranks of the unemployed, it should give rise to an externality in terms of unemployment compensation saved, where such compensation exists in a country. Also, allowance should be made for a supply price of such labor which is lower than the wage to be paid, giving rise to a supplier surplus benefit as an unemployed worker gets a job. But such benefits only prevail for a relatively short time. Neither will survive once normal conditions are restored. More importantly, it should be recognized that even in a deep recession, most labor hired is ultimately sourced from alternative employments rather than from the ranks of the unemployed. Coming out of recession, gross job creation far exceeds the contemporaneous fall in the ranks of the unemployed. To dramatize this, we would like to offer the following comments.

Take it as given that when the overall unemployment rate is 5 percent, the creation of 100 new jobs does not change that rate, hence net absorption is zero. At the other extreme, take it as given that at a general unemployment rate of 25 percent, the creation of 100 jobs will absorb 100 unemployed. We think that benefit-cost analysts should plot some sort of continuous curve between these two points. We

cannot expect to see zero percent absorption at a general unemployment rate of 5 percent, and find as much as 50 percent absorption when the general unemployment rate is 6, 7, 8 or even 10 percent. The ratio of net reduction in unemployment to gross job creation, ideally by occupational category and labor market area, is probably our best guide to the likely magnitude of the absorption coefficient.

The analysis of cyclical effects in the capital market seems to us to be murkier than what we find for labor. By analogy, we might have both f_1 and f_2 equal to zero in a great depression with 25 percent unemployment, and equal to their normal values in normal times. On this basis some continuous curve between full value and zero would seem appropriate to trace what happens as the economy moves from normal times to deep depression. Also, it would seem appropriate that the time preference rate r would be lower in recession times, reflecting cyclically lower market interest rates.

As mentioned in the text, we have been piqued by these issues but have not found them vital, simply because of the forward-looking nature of most practical benefit-cost analyses. We think a good start to dealing with cyclical issues on the capital side would be to undertake extensive studies of the social internal rate of return to capital in private companies (e.g., the Ford Motor Company, in IBM, in Dow

Chemical, General Foods, etc.). The internal rate of return should be on the basis of all capital inflows coming from bank borrowing and bond sales as well as stock issues, and outflows including interest and amortization payments as well as dividends and stock repurchases. All should be in real terms, and low interest rates in cyclical dips should automatically be reflected. If the exercise starts later than the founding of the firm, an estimated real stock of initial capital should be inserted. Terminal capital at the cutoff date should be the real value of equity plus debt at that time. Sensible real depreciation rates or patterns should be used, rather than the nominal rates on the company's books, which are skewed by tax laws and unadjusted for inflation. The social cost of labor could be adjusted cyclically for those hired in recession or depression periods. Such an exercise would span both good and bad years and would give us evidence on the effects of including cyclical considerations in our analyses. We feel it is unfortunate that such exercises, which we feel have a high professional and social value, are hard to sell to our dissertation-writing students. How would such a mundane topic, regardless of its importance, help them in the job market?