INTRODUCTION TO COST-BENEFIT ANALYSIS

P A R T I I

LABOR MARKET ISSUES

Arnold C. Harberger

University of California, Los Angeles

December 2008

Background

There is probably no part of cost-benefit analysis that is the source of more grievous mistakes than the labor market. And understandably -- people respond to claims like "job creation" and "employment impact" in a way that has no counterpart concerning benefits linked to the generation of foreign exchange, or to high rates of real return on capital invested in a project. More's the pity, because most of labor-oriented claims are specious and ill-founded -- they represent sloppy thinking at best.

If there is one simple lesson that is more important than all others in the labor area it is that, like the use of capital, materials, and any other factor of production, the employment of labor represents a cost to a project. It may be that this cost is lower than the wage or salary paid by the project, thus generating an external benefit, but taking both direct costs and external benefits into account the labor used in a project is virtually certain to entail a net cost.

How can one best dispel the false notions evoked by the idea of "job creation" by a project. I think the best single response is that we are not comparing our project with "nothing". Instead we are comparing it with the alternative use of the resources it occupies. Thus, if one

thinks of our project's use of capital as generating a number of new jobs, we have to ask what the plausible alternative use of such capital funds would have generated. Typically, a given public-sector project will be one of a number of alternative uses of public funds that are being considered; in this case each project has direct competitors under scrutiny at the same time. But even when there are no direct competitors the project will be drawing, from the capital market, funds which would otherwise be used for investment or consumption.

The preceding paragraphs look upon jobs being generated by capital investment, which is the way those who use the term "job creation" typically think about the issue. However, we must recognize that many new jobs arise that are not connected to new capital investment, and that firms often lay off workers without there being any corresponding disinvestment in buildings, machines or other capital assets.

This is why economists greatly prefer to view the project's use of labor as a labor market phenomenon. If a project hires new workers, the costs involved in that particular action will be essentially the same, regardless of whether the project is at the same time acquiring capital goods, or not.

When we take this view, and see the project entering the market with a new demand for labor, we see it as similar to other markets. Just as an added demand for foreign exchange is in the final analysis met by some combination of displaced imports and newly-stimulated exports; just as an increased demand for capital is ultimately filled by some combination of displaced other investments and newly-stimulated savings; just as a new demand for wheat will in the end come from a combination of increased world production and the displacement of the demand of other users, so it is with labor. An added demand for labor should be seen as being filled in part by displacing the demands of others and in part by stimulating an increase in labor supply.

Measuring the Economic Opportunity Cost of Labor -- The Standard Case

For the labor that the project gets, the normal assumption is that the project pays these workers a wage equal to their supply price. This in turn is based on the idea that these workers are getting a "market wage" for the particular type of work and in the particular area of their employment.

It is not true that the market wage is the same for all truck drivers or for all barbers or for all secretaries, throughout the whole labor market of a country. For years I have told my classes of seeing ads in the Chicago papers, offering $200 per day for truck drivers to come to northern Alaska to work on the Alaska oil pipeline, which was then being constructed. This wage seemed surprisingly high, for at that time the typical wage for truck drivers in Chicago was around $50 per day. Yes, the wage in Alaska was a lot higher, but why did these ads continue month after month? The answer is that the ads did not precipitate a mad rush of applicants, leading to a long waiting list of willing workers. This is what would have resulted if the Alaska offer were really at a significantly "above-market" (for Alaska) wage. The fact is that it required a huge premium to attract drivers from a pleasant, "normal" Midwestern routine to a place where below-zero temperatures and near-total darkness reigned for half the year, where mosquitoes swarmed in the outside air for the other half, where separation from family and friends was understood to be part of the job, where most customary forms of entertainment had no place. Economists conclude in such a case that the huge premium embodied in the $200-a-day wage was an "equalizing differential" -- a compensation that was required in order to stimulate a supply of truck drivers that was sufficient to meet the pipeline project's demand. Thus, if we had a trucking project in Chicago, we would consider the market wage to be $50 per day, but if we had one in northern Alaska, we would take the market wage to be $200 per day. And if our project were located

along the beaches of Hawaii, maybe we would find the market wage there to be only $30 per day.

What, then, is the economic cost of labor in such cases? Following our earlier line, we look to taxes as the main source of divergence. If workers move from Chicago, they generate a reduction in the taxes the government receives on the basis of their earnings. These include the payroll taxes linked to their earnings, the personal income taxes they pay, plus (where applicable) the sales and consumption taxes paid as the workers spend their income. The amount of these lost taxes would be the same, regardless of whether the workers moved from Chicago to Alaska (where they would earn $200 a day) or to Hawaii (where they would earn $30 a day). Suppose these lost Chicago taxes were $5 for payroll taxes, $6 for income taxes and $4 for sales taxes, for a total of $15. On this basis, it would seem appropriate for us to take the economic opportunity cost of truck driver labor to be $215 (= $200 + $15) in Alaska and $45 (= $30 + $15) in Hawaii.

However, this is not the recommended procedure. The question is, what to do about the taxes that these workers or their employers will pay, on the basis of their wages, as they take up their project employments? In the case of the economic opportunity cost of foreign exchange, we built into the opportunity cost the tariff revenue that was lost as imports were displaced, but counted the tariff or excise taxes paid by the project on its tradable inputs as external benefits of the project itself. We did this because of the great variability of possible tax treatments among the many different tradable items the project might buy. The lost tariff revenue from the "sourcing" of the foreign exchange was standard, regardless of what the money was then spent on. But the spending of this money carried very different tax effects depending on the specific items being bought. In the case of workers of a given type being hired at a given place (e.g.,

truck drivers in Alaska), the amount of taxes that will likely be paid in that place, on the basis of their wages, is just as predictable as the taxes lost at the source. Hence it is quite appropriate to include these "destination taxes" as benefits, operating to offset the cost of the taxes lost at the source.

Thus, suppose the estimated taxes in Alaska were $20 for payroll taxes, $40 for income taxes and $10 for sales and excise taxes, for a total of $70. Then our figure for the economic opportunity cost (EOCL) for truck drivers in Alaska would be $200 + 15 - $70, or $145, significantly less than the $200 prevailing market wage.

In contrast, the EOCL of truck drivers in Hawaii would tend to be greater than the $30 market wage being paid. This is because the taxes newly paid in Hawaii in the basis of a $30 wage are likely to be significantly less than the $15 of taxes lost at the point from which the labor was sourced. Taking the Hawaii taxes to be $2 for payroll taxes, $3 for income taxes, and $2 for excise taxes, for a total of $7 a day, we have EOCL = $30 + $15 - $7 = $38 per day.

On The "Sourcing" of Project Labor

In the above examples, we have taken Chicago to be the source of the truck drivers drawn to either Alaska or Hawaii. In actuality, workers are drawn from many different sources, the identification of which is much more subtle than it initially appears.

One's initial tendency concerning the sourcing of project labor is to think like a pollster -- if the project already exists, just go and ask the workers what their last job was, where it was located, how much they earned there, what taxes they paid, etc. This, however, is not what is needed. When project workers leave their last job to come to this one, they are very likely going to be replaced by other workers, who in turn will leave other jobs where they again might be

replaced, etc., etc. When we seek the sources of project labor, we are looking for the ultimate sources, not the immediate ones.

The idea of the "ultimate sources" is in reality quite simple, and it is based on an absolutely fundamental element in project evaluation, in policy analysis, and in the whole field of applied welfare economics. What we are doing in such analyses is comparing an economic equilibrium that exists in the presence of a given project (or policy) and comparing that equilibrium with the one that would exist in the absence of that project (or policy). When we analyze the effects of a given tax, we usually employ a simple supply-and-demand framework, and insert into that framework first a zero tax, and alternatively a tax of size T. We then compare the equilibria that result from these two exercises to estimate the efficiency cost of the tax, or its incidence on different groups, or its revenue yield.

When we deal with project analysis the story is similar, but more complicated. For here we are looking at the profile of the project's costs and benefits, over its entire life. We thus have to think of comparing two moving pictures -- one depicting how the world would evolve (in its relevant aspects) in the absence of a project, the other showing a similar evolution in its presence. Our project profile is then obtained by each year's difference between these two moving pictures, with respect to the project's direct inputs and outputs. The external effects of the project are similarly obtained by taking, for each period, the difference with respect to government tax receipts and possibly other relevant distortions.[1]

---

[1] Distortions play a very important role in applied welfare economics. Without entering into a full exposition, I hope I can convey the key point, that we are not concerned with changing prices and quantities in other markets when those markets are not distorted. In an undistorted market, the market price measures the economic opportunity cost from the side of the suppliers, and it also measures the benefit of added units from the standpoint of the demanders. Movements up and down of the equilibrium price and quantity thus have incremental costs equal to incremental benefits. When distortions are present, there is an economic gain when quantity goes up and demand price (benefit) exceeds supply price (cost), and an economic loss under the

We have to think of these same two moving pictures as we contemplate the sourcing of labor for a given project. Rather than try to expound this subject in abstract terms, I will try to convey its essence by describing a real-world study in Mexico, on which I served as an adviser. The basic data for this study were labor market surveys in around 40 different labor market areas of the country. In each of these areas the surveys provided data on wages, average hours, average earnings, etc. for some 100 different occupations, with separate data for men and for women.

The first task was to simplify the problem by amalgamating the 100 different occupations into meaningful and manageable groups. Thus we had common, semi-skilled, and skilled labor, ordinary clerical and skilled clerical workers, technicians, professionals, and perhaps one or two more categories. The rule for putting several occupations into one category was that these occupations had to pay similar wages.

So now one had the average wage rates of, say, 10 labor categories, in 40 labor markets covering basically the whole territory of Mexico. For each category and each location, the tax payments linked to that category in that place were estimated (analogous to $70 in Alaska, $15 in Chicago, and $7 in Hawaii).

Next, we had to make assumptions about the geographical sourcing of incremental labor for each of the 40 centers. In all cases it was assumed that half the labor came at the expense of other jobs in that same labor market area. Different assumptions were made about the sourcing of the other half of project labor. Under the first alternative -- called the donut -- the second half of project labor was assumed to come from the contiguous labor market areas, in proportion to their existing numbers of workers in the category (e.g., semi-skilled). Under the second alternative -- called the nationwide -- the second half of project labor was assumed to be drawn

---

same concept if the equilibrium quantity goes down (as between the two moving pictures).

from all 39 other labor market areas, again in proportion to their existing numbers of workers in the category.

So, if w_{ir} is the monthly wage of workers of category i in region r, one estimated the amount of T_{ir} -- equal to the monthly taxes paid, linked to the wage w_{ir}. Then the economic opportunity cost of labor of category i in region s would be

$$w_{is} + [T_{is} - \sum_r a_{ir} T_{ir}].$$

Here a_{ir} for region r would be the fraction of workers in category i (in this case semi-skilled workers) coming from each sourcing region. In the Mexican study, a_{ir} was always 1/2 for the region in which the project was located. The remaining half would be distributed across the other sourcing regions in proportions of their numbers (N_{ir}) of workers in the category in question. That is, a_{ir} for regions other than s would be equal to $(N_{ir} / \sum_{r \neq s} N_{ir})(1/2)$.

This may look more complicated than it really is. What we are doing is taking as a benefit the taxes to be paid by (or on account of) workers in the project, and as a cost that taxes that would have been paid by (or on account of) them in the employment they likely would have had in the absence of the project.

Where one has sourcing among major cities in a country, one typically does not expect to find significant difference in the taxes paid (per worker) as between source and destination. However when a project is in the main metropolitan area of a country, and is likely to be sourced largely by migrants from the rural hinterland, then there can be large differences in taxes paid, which will tend to lead to an economic opportunity cost of labor somewhat lower than the market wage. This is likely to be the case for low-skilled labor in many developing countries (notably

China and India). There the wage paid in the urban destinations of migration tends to be very substantially higher than the wage in the rural sources of migration. This is taken as a genuine economic cost, as it reflects the fact that labor probably would not move to the cities in the absence of a significant premium (as in the case of truck drivers in Alaska, but with a more modest premium (e.g., city wage perhaps double the rural wage).

The bottom line of this section is simply that most of the time one takes the wage actually paid (or to be paid) by the project as the first approximation measure of the economic opportunity cost of labor. This is then adjusted by a difference in taxes linked to the wage (payroll taxes, personal income taxes, and sales and excise taxes) as between employment on the project in question and the alternative employment "at the source" (conceptualized as the difference in employment at each source as between the two moving pictures -- "with" and "without" the project in question).

Dual Labor Markets -- Protected vs Unprotected Sectors

In many developing countries we observe the phenomenon of "dual" labor markets. Sometimes these are referred to as "formal" and "informal" sectors; sometimes as "modern" and "traditional" sectors. I like to use the terms "protected" and "free" sectors, because these terms most readily reflect the theoretical underpinnings of a dual labor market. The easiest way to visualize this syndrome is to think of everybody in the labor force wanting to have a job in the protected sector, but only a fraction of them are lucky enough to get one. The rest fall into two categories -- those who take a job (at a significantly lower wage) in the free sector, and those that I call the quasi-voluntary unemployed -- who are eager to get a job in the protected sector but who are not willing to accept the wages and working conditions of the free sector.

To put a little meat on these conceptual bones, let me recount some relevant experiences from India. I spent the academic year 1961-62 working on a project of the MIT Center for International Studies, which served in a collaborative and advisory capacity with the Planning Commissions of India. The Center's headquarters were in the diplomatic section of New Delhi, where most embassies were located. And the project's financing came mainly from the Ford and Rockefeller Foundations. It should come as no surprise, then, that the wages and salaries paid by the project were significantly above the free market level -- being in this respect similar to those paid by the U.S. Embassy, the U.K. High Commission, and nearly all the other missions representing modern industrial countries.

It should come as no surprise, then, that just about everybody in the Delhi labor force would like to have a job in this sector. Many embassies had waiting lists of hundreds of applicants. In our project's case, we didn't keep a waiting list, but a steady stream of applicants came to our office door. Sensing the relevance of our situation, I asked Mr. Lakshman, our office manager, to peek into my office when he finished interviewing each applicant. If I was not too busy, he would then pass that applicant on to me, for a second session. My practice in such cases was to pull out the "help wanted" page of the daily paper, and go through the ads with the applicant looking over my shoulder. Quite regularly, we would find ads for salesmen, particularly for travel to rural areas. People with middle-school education were typically wanted, for example, as salesmen to sell pump-irrigation equipment to farmers. The pay, as I remember was quite good, but not a single one of my interviewees expressed any interest in a pump salesman job. A number of them expressed actual disdain for such a job, thinking it to be far below their own social status. When pressed as to what kind of jobs they would like, they almost always mentioned the U.S. Embassy and the British High Commission, plus other diplomatic

representations. Usually, too, they mentioned the Indian government as an employer they would like. In such cases, I would turn again to the want ads, and usually found some that had been placed by the Indian government. Typically, these might be for educated Indians to serve as public sector representatives and technical advisers in rural villages. Once again these jobs were spurned, usually with comments like "What would I ever be able to do in a village?" The fact that these types of ads came up repeatedly convinced me that they did not represent protected sector jobs. What I was seeing was a whole series of people who were eager candidates for genuine protected sector jobs, but who had no interest in the available, specific, free-sector jobs.

Quasi-voluntary unemployment is quite common in developing countries, particularly among the more skilled and better educated portions of the population. In India, at least when we were there, the highest rate of recorded unemployment was among recent university graduates, who often chose to wait for a job with good working conditions, ample pay, and favorable career promise. Sometimes they would be on the labor market as long as two years before they found a job they were willing to accept. Of course, in order to be classified as unemployed they would have to respond positively to questions like "What efforts have you made during the past month in order to seek employment?" And of course their answer would be "I've been to the American Embassy, the U.K. High Commission, the Indian Foreign Office, the Delhi branches of Monsanto and General Electric, etc., etc." They would thus end up classified as unemployed. But quite clearly their unemployment was quasi-voluntary. They were involuntarily unemployed, looking up at the protected sector wage, and voluntarily unemployed looking down at the free-sector wage.

What can we say about the economic opportunity cost of labor in such cases? The answer is that the opportunity cost of labor for a protected sector job (at wage w_p) is higher than

the free-market wage, and lower than w_p itself. The most fruitful assumption to make here is that the people who end up filling a set of protected sector jobs are selected by criteria other than their individual supply prices of labor for these jobs.[2] Those who are working in the free market at a wage w_f are deemed to have that as their relevant supply price (i.e., they would gain a personal surplus of the extra earnings (w_p-w_f) if they were selected for a protected sector job). Those in the quasi-voluntary category are assumed to have supply prices spread out evenly over the range between w_f and w_p, so their average supply price is taken to be $(w_f+w_p)/2$. Hence the EOCL for a protected sector job at wage w_p is estimated to be a weighted average of w_f (with a weight proportional to the number of workers of the given category who are employed in the free sector) and $(w_f+w_p)/2$, with a weight proportional to the number of quasi-voluntary unemployed in that category.

Of course, each occupation, each labor category, will have a different free market wage, w_f. And so far as w_p is concerned, there can be many different entities paying more than the free-market wage to a given category of labor. We have mentioned the diplomatic sector and certain types of government jobs. In addition we find protected sector jobs where the "protection" comes from minimum wages, or from strong labor unions. One also finds, as alluded to above, protected sector jobs in the offices and plants of multi-national companies. These companies are particularly vulnerable to labor unrest, because they pay so much more (for what appears to be the same type of work) in their home base than they do in developing countries. Their best response to ward off such unrest is to offer wages and working conditions

---

[2] Formally, the conditions are that from among the available applicants, the successful ones are chosen by criteria <u>uncorrelated</u> with their individual supply prices.

such that they always have a large backlog of willing applicants waiting for each kind of job.

The EOCL for a protected sector job then will be a function, both of the free-market wage, and the w_p for that particular job. In general, the EOCL will be higher for employment at a wage of $10 a day than at one of $7 a day, assuming the free-market wage to be, say $4 a day. Why is this? -- simply because there are more quasi-voluntaries waiting around for the $10 job than there are waiting for the $7 job. Thus not only do we have ($4 + $10)/2 entering into the weighted average, rather than ($4 + $7)/2 in the case of the $7 job, but we also have a heavier weight in this part of the formula for the $10 job vis-a-vis the $7 one.

What should also be clear is that the presence of quasi-voluntary unemployment is different from the phenomenon of cyclical unemployment. In the latter case, there is the presumption that, the higher the observed unemployment rate, the lower is likely to be the economic opportunity cost of labor. In the case of quasi-voluntary unemployment, it will be higher, for a given labor category, the higher is the relevant protected sector wage, w_p. Hence higher quasi-voluntary unemployment is likely to be linked to a <u>higher</u> economic opportunity cost of labor for protected-sector jobs at a higher wage.

The previous discussion, leading to the conclusion that the EOCL for labor hired at the free-market wage (w_f) is w_f itself, and that for labor hired at any given protected-sector wage w_{pj} (higher than w_f) would be a weighted average of w_f itself on the one hand and $(w_f+w_{pj})/2$ on the other, was intentionally kept simple, so as to focus directly on the essential features of a dual labor market. Thus we did not introduce the complications linked to taxes paid on the project job and taxes forgone on the locations from which the labor was sourced. Our having simplified in this way does not mean that we should neglect these tax externalities, in a

real-world dual-market case. Rather, one should in such circumstances graft them on to the bare-bones analysis that we have presented. In this type of case, the full amount of taxes paid by a new protected sector job would be a benefit, as before. Similarly, the assumption is that no taxes would have been paid by the quasi-voluntary unemployed, and that those drawn (in net terms) from the free-market sector would have paid taxes in amounts linked to the free-market wage, w_f.

Migration-Fed Unemployment

In a number of developing countries one encounters the special case of migration-fed unemployment. Typically, this phenomenon occurs in a growing urban labor market, where migration from rural areas is an important source of that labor-force growth.

To begin, one should recognize that there is nothing at all anomalous about rural-urban migration -- a process that appears to be an integral part of economic development, no matter where. A properly-working labor market would absorb migrants as they came, allowing of course for a reasonable quota of frictional or search unemployment upon arrival. As in any supply-and-demand situation, one expects the market price (the wage in this case) to continually adjust so as to maintain approximate equilibrium between supply and demand.

The special feature of situations of migration-fed unemployment is the fact that in these cases the wage is <u>not</u> performing its natural function of equilibrating supply with demand. The easiest way to visualize this is to assume that something (typically a legal minimum wage) is artificially holding wages (quite generally in the urban destination) above their equilibrium level. People in the rural area -- getting news of such an attractive opportunity, are motivated to move in large numbers to the city. This process would go on and on and on, if each migrant settled quickly into a regular urban employment. But that would bring a Niagara of migrants --

something that the city couldn't absorb and, more important, something that we do not observe. What we see is a process of migration that in a rough way keeps pace with the increasing demand at the urban destination. That is, we see something similar to what we would observe if the labor market were working as it should with the wage rate performing its equilibrating role, but with one big difference -- the unemployment rate stays high, year after year. That is the anomaly.

The explanation is not hard to find, nor hard to understand. What happens in these cases is that the unemployment rate becomes the equilibrating variable, taking over that role form the wage rate. With a very attractive, high urban wage rate a heavy flow of migrants is to be expected. But demand is not increasing as fast as this supply so unemployment grows. With each increase of the unemployment rate, the incentives perceived by potential migrants grow less attractive. At a zero rate of urban unemployment they all want to move to the city. At, say, a 30% urban unemployment rate, they all want to stay home on the farm. Somewhere between these two extremes there has to be a rate of urban unemployment that makes moving to the city just barely attractive to potential migrants. This is the rate that will tend to prevail in a situation of migration-fed unemployment. When the unemployment rate is lower than this, net migration will be stimulated; when the unemployment rate is higher than this, net migration will slow, or even turn negative, until the equilibrium rate is restored.

Table 1 gives a stylized picture of how unemployment serves as an equilibrating force. The table is based on the assumption that each worker in the urban labor force has the same probability of getting a job. This is an unrealistic assumption, but it greatly simplifies the exposition. Readers should be assured that the example can be adapted to incorporate more

realistic assumptions (e.g., that migrants, when they first arrive face a very high unemployment rate, which then gradually declines to a significantly lower rate as with the passage of time they

TABLE 1

Exercises in Migration-Fed Unemployment

(supply price of migration = $4/day)

Panel 1: Urban Wage = $5/day; Equilibrium Unemployment Rate = 20%

		Employed	Unemployed	Labor Force	Expected Earnings
1)	Initial Equilibrium	800	200	1000	$4/day
	Now Create 100 New Jobs				
2)	Short-Run Response	900	100	1000	$4.50/day
3)	Long-Run Response (new equilibrium)	900	225	1125	$4/day
4)	Change From 1) to 3)	+100	+25	+125	-0-

Panel 2: Urban Wage = $6/day; Equilibrium Unemployment Rate = 33 1/3%

1)	Initial Equilibrium	800	400	1200	$4/day
	Now Create 100 New Jobs				
2)	Short-Run Response	900	300	1200	$4.50/day
3)	Long-Run Response (new equilibrium)	900	450	1350	$4/day
4)	Change from1) to 3)	+100	+50	+150	-0-

assimilate more fully into the urban setting). Such an adaptation process would change neither the basic structure of the analysis nor its lessons and conclusions.

Table 1 is built on the assumption that the equilibrium rate of unemployment is 20% and that the "supply price of migrants" is $4 a day. That is to say, migrants will be satisfied with their move, so long as their average earnings amount to $4 a day or more. An expectation of less than $4/day previous migrants will want to go back to their rural home; with a higher earnings expectation, a flood of new migrants would come to the city. In the initial equilibrium, we have a labor force (of a given category of labor) of 1000, of which 800 are employed and 200 unemployed. The employed workers earn the established wage of $5 a day but their expectation is to be employed one fifth of the time, so their expected earnings are $4 per day, just equal to the "supply price of migration". We thus have an equilibrium, with no incentive for either out- or in-migration.

Next, 100 new jobs are created. In the immediate short run, the labor force stays the same, yielding 900 employed (at $5 per day) and 100 unemployed. Now expected earnings are $4.50 per day, well above the $4 supply price of migration. This attracts new migrants, who will keep coming until a new (long-run) equilibrium is established. That will occur when the labor force of this skill category has swelled to1125, of whom 900 will be employed and 225 will be unemployed on an average day. The expected earnings would again be $4 per calendar day, once again equal to the supply price of migration.

We now have to ask, what economic opportunity cost should be assigned to the 100 new jobs created in this example? The answer is $5, the actual urban wage rate applying in this case. Recall that this wage is higher than would prevail in a straight supply-and-demand equilibrium. Being artificially higher, it attracts migrants to the point where they end up just barely satisfied

with having made the move. The real "supply price" of the 125 migrants is (125 x $4/day), or $500 per day. But this is exactly what is earned by the 100 of such migrants who will be employed on an average day. Looked at in this way, the $500 actually paid just barely covers the resource cost of the 125 migrants who were attracted as a result of the 100 new jobs.[3]

The lower panel of Table 1 reflects a case where the urban wage has been raised to $6 a day, and where as a consequence the equilibrium rate of unemployment has moved up to 1/3. The initial equilibrium thus has a labor force of 1200, of which 800 are employed and 400 unemployed on any given day. Actual earnings of the 800 are (800 x $6), or $4800 per day. Expected earnings of the 1200 are $4800/1200, or $4 a day, precisely equal to the supply price of migration.

When we add 100 jobs, the labor force initially stays at 1200, but employment goes up to 900. Expected earnings rise to $4.50, again triggering a new wave of migration. A new equilibrium is reached when the labor force has increased to 1350, employment is 900 and unemployment 450 in this new equilibrium. Thus the creation of 100 jobs has led to increases of 100 in employment and of 50 in unemployment. This is the bane, the curse of migration-fed unemployment. If one's aim is to reduce the number of urban unemployed, creating urban jobs is actually counterproductive!!

The EOCL in this case turns out to be, as before, the same as the urban wage. Creating 100 jobs has caused 150 migrants to move to the city. The economic cost (supply price) of those

---

[3] An alternative way of reaching the same conclusion is to note that the financial cost (to the project) of the 100 jobs is $500 per day. On any given day 100 migrants will tend to be employed at $5, each earning a positive surplus (above the $4 supply price of migration) of $1 per day, for a total surplus of $100 per day. But in addition we have 25 migrants earning zero, representing a deficit of $4 a day, in comparison with their supply price. These 25 cases generate a deficit of $100 per day. When we adjust the $500 financial cost to take these external effects into account, we get $500 + -100 - 4100 = $500. Thus the economic opportunity cost turns out to be equal to the financial cost.

moves was 150 x $4/day, or $600 per day. That is exactly equal to $6 per day (the actual wage) for each of the 100 new jobs.

Though our simple example may seem remote from the real world, the analysis of migration-fed unemployment has genuine roots in reality. Harris and Todaro were puzzled by the chronic unemployment they saw in Africa. I was puzzled by a constant high rate (15%) of urban unemployment in Panama during the 1960s. Harbison noted how, when a sudden large spate of job creation occurred in Kenya, new migrants came in numbers sufficient not only to fill all the new jobs, bur also to increase the pool of unemployed in the same proportion, leaving the unemployment rate the same.

To repeat the main lesson emerging from this section -- when the conditions of migration-fed unemployment exist, new projects created in the urban destinations should get no credit for absorbing unemployment. New urban jobs in such circumstances actually add to urban unemployment, when one moves from one labor market equilibrium to another.

Cyclical Unemployment

This topic was intentionally left for last because, in spite of its being a natural subject of policy interest and at the same time posing fascinating problems for the analyst, it is probably a matter of rather minor importance in a world of institutionalized, formal economic project evaluation. The reason for this is simple -- when such project evaluation is being done as a serious, real-world exercise, it typically precedes by some significant amount of time the actual implementation of a project. Most of the time, groundbreaking in a major project will not occur until three or four years after the project has been evaluated and approved. All the detailed engineering, the packaging of project finance, the assembly of leadership teams will take place during the interim.

So if the economy happens to be in a cyclical downturn at the time a project is being evaluated, it will very likely be out of those woods by the time the project actually gets under way. One then looks forward to all the years of the project's expected economic life -- very often 30, 40, or 50 years. What can one say about the projected cyclical stage of the economy in those years? The best-assumption by far, is that we will be dealing with normally-functioning labor and product markets as the economy lives through the expected lifespan of the project.

Moreover, in project evaluation we normally are pitting our project against an alternative, either explicitly or implicitly. Capital for the project is deemed to come mainly at the expense of other investment or of consumption. Labor in the project comes mainly from other employments, etc. There is a real question of how to fit such a framework into a setting of recession or depression. Certainly we would not want to count the project's absorption of some number of unemployed as a big project benefit, without recognizing that most other expenditures of the same money would also have similar effects.

Seeking to resolve the puzzles posed by cyclical unemployment in a reasonable way, I like to think in terms of the subset of projects that actually have unemployment absorption as one of their important objectives. The setting would be something like the present (December 2008) recession, or that of the mid-1970s or the early 1980s in the United States, or the 1980s debt crisis period in Latin America. Each project is then looked upon as drawing its labor force: a) from other employments, b) from new labor-force entrants and c) by drawing down the pool of the unemployed.

It is obviously item c) that makes the cyclical case differ from the norm. What we seek here are two fractions -- first, the fraction of our hirings that will be represented by the net

reduction that they induce in the pool of the unemployed; and second, the amount by which the project's outlay on such workers exceeds their true economic opportunity costs.

There is no way of getting anything like a precise estimate of either of these two fractions, but a good starting point is to recognize that the fractions, as it were, move up and down together. Thus, in a mild recession, perhaps only 10% of those hired would come (in net terms) from the pool of the unemployed, and they in turn might have an economic opportunity cost only 10% below the wage they receive from the project. In a deep depression, perhaps as many as half the workers hired might represent net reductions in unemployment and the EOCL might be 40% below their project wage. In between these two extremes we might have a serious recession (not yet a depression) in which, say, a quarter of those hired would represent reductions in unemployment, and their EOCL might be 20% below the project wage.

Since years of recession and depression are not the norm, and since the calibration of new projects in terms of their unemployment absorption is only for a specific subset of evaluations (focusing on projects ready for essentially immediate adoption), one can probably do quite well testing a couple of alternative fractions -- say a 10% absorption fraction plus a 10% wage discount on the one hand, and a 20% absorption fraction plus a 25% wage discount on the other.

Some readers may be surprised at the fact that the suggested wage discounts are so modest. The answer is that we have to be modest in this regard if we want to respect the evidence. Ultimately, voluntary supply price is the best measure of opportunity cost as seen by the worker himself or herself. Many economists start out with the idea that unemployed workers make major adjustments in their supply price, bringing it down step by step as their unemployment spell gets longer and longer. Surprisingly, however, the data do not support this view. By and large the stated supply price of the unemployed remains much the same,

independent of the length of the spell. Yet these unemployed workers celebrate when they finally get a job. Here we have a contradiction. If one's stated supply price is $2000 per month, economics tells us to infer that at that wage the worker is indifferent between working and not working. But the celebration on finally getting a job negates such indifference. Resolving this contribution is a standing challenge to economists -- we certainly cannot do it here. But the contradiction itself signals us to proceed with caution. Certainly assigning a zero or very low opportunity cost, as many are prone to do, is ill-advised in the extreme, in the face of the evidence.[4]

Special note should be taken of a phenomenon that can indeed lead to a low EOCL for labor that is absorbed from the ranks of the unemployed. I refer to cases where the government has a functioning program of unemployment compensation, so that a reduction in the ranks of the unemployed generates an important fiscal saving for the government (and, more fundamentally, the country's taxpayers). In this case the employment of an unemployed worker in a $3,000 a month job might mean an economic opportunity cost of only $1,200 a month -- the difference stemming from a saving of $1,500 a month in unemployment compensation payments to that worker, plus, say, $300 a month in taxes that will be paid on the basis of the $3,000 salary. This element, however, is rarely present in low-income developing countries. Serious programs of unemployment compensation are largely concentrated in the so-called "first world".

---

[4] I tell my students every year that I am quite willing to assign a zero EOCL to any worker who is willing to work for zero compensation. And there are many of them -- volunteer workers in hospitals, sport leagues, churches, etc. But we don't find them very often as members of a country's regular labor force!