

INTRODUCTION TO COST-BENEFIT ANALYSIS

PART V

APPLICATIONS TO IRRIGATION PROJECTS

Arnold C. Harberger

University of California, Los Angeles

January 2009

Background

Just as the benefits of highway projects are closely linked to the volume of traffic, so those of irrigation projects are linked -- guess what? -- to the amount of water delivered, and its value. But in both these cases claims abound for a variety of external effects that do not hold up under scrutiny. We very often see the double counting of benefits on many different types of projects, but I think that one so-called project report that I once reviewed in India probably holds the record. In that report, benefits were claimed: a) equal to the value of the water, 2) plus the increase in the value of the land that took place as a consequence of the project, 3) plus the increase in the value of crops produced on that land, and 4) plus the wages bill paid for the extra employment that emerged as a result of the project. As I said, I had seen cases of double counting if benefits quite often, but this was a case of triple, and even quadruple counting!!

We start with the value of the water delivered by the project. Water is a productive factor for agriculture, side by side with land, labor, capital (fences, buildings, cattle, farm machinery, etc.), fertilizers, gasoline and so on. Each of these factors of production has what economists call a marginal product, and, broadly speaking, a market system leads to a situation in which the value of this marginal product is brought into equality with the price that has to be paid for that factor. To illustrate, suppose the agricultural wage is \$5 a day, and suppose a farmer estimated that adding one worker to his labor force would generate an increase of \$6 or \$7 per day in the farm's output. Obviously, it would be worthwhile to hire that worker because the benefits of that action (\$6 or \$7) would exceed the cost. Now the farmer's hiring one more worker is not going to have a perceptible (to him) impact on the market wage, so natural economic incentives will work to keep adding to the farm's labor force until the marginal product of an extra worker gets down to around \$5 a day. The same sort of process works with respect to each and every factor of production that can be freely bought (or hired) in the marketplace at the market price. This is the case for labor, fertilizer, gasoline, farm machinery, etc.) But not for the water from an irrigation project. This water is distributed in quotas to the farmers who have "water rights" linked to that project. These rights are usually set on the basis of equal amounts of water per hectare (or acre) of the area served by the project. But no one knows in advance how much water that will be. That depends on the forces of nature (rainfall, snowfall, etc.) and on how much of the river's water is taken by others (upstream) or reserved for the use of others

(downstream). The irrigation quota for a given farm might thus be 100 cubic meters one year, 20 the next, and only 5 in the one after that. The market principle would work if farmers in the project were able to trade irrigation water among themselves, a process that would lead to a high price when water was scarce and a low price when it was abundant. Economists have long argued for freedom to trade water -- within irrigation projects and even outside them -- but progress in this direction has been very slow. The norm is still that farmers each get "their share" of the available water, a share that will be big or small, month by month and year by year, as nature and the priorities of other users dictate. Usually, the farmers have to pay something for their water rights, but more often than not it is a fixed charge per hectare, rather than a price per cubic meter of water actually delivered. And when it is a price per cubic meter, it is usually a very low price, which is far below the productive value of the water. So, whereas for other factors of production there is a strong tendency to use more and more of a factor, up to the point where its marginal product matches its market wage or market price, this is rarely true for the water delivered by irrigation projects.

Thus, we cannot measure the marginal productivity of irrigation water by what the farmers pay for it. They usually pay much less than its marginal product, which leaves us with the problem of ascertaining the economic value of irrigation water (if we are doing an ex post evaluation of our existing project) or of predicting that economic value for a new project being analyzed. What we must remember here is that the value we are seeking is closely linked to the market price that would prevail, varying from month to month and year to year, if the farmers served by the project were freely able, among themselves, to buy and sell water deliveries period by period.

We are going to argue that by far the best way to estimate the likely value of irrigation water is to keep our eyes always on the water -- not looking at indirect ways of getting estimates. But first I want to explain the principal indirect method that is actually widely used, and that has its own underpinnings in good economics. This is often called the farm budget method, but I prefer to call it the "residual value method", because this label much better conveys how the method really works.

The basic idea of the residual value method is to build up two typical farm budgets -- one in the absence of the project in question, the other in its presence. These two budgets are often quite different -- the first involving a dryland rotation, with crops that do not need much water, and the second dealing with a very different rotation built on the presence of irrigation water. But, also quite often, an irrigation project will deal with land that is already irrigated, drawing water from a river as it passes. Such projects typically involve building a dam on the river, leading to greater availability of water during the irrigation season, plus a degree of control over precisely when, during the agricultural year, the quotas of irrigation water are delivered to the farms. In these cases -- of dams simply enhancing the capacities of pre-existing river-irrigation projects -- cropping patterns will undergo little or no change as a result of the project.

The residual value method focuses on the typical farm's estimated profit-and-loss statement. Quantities are for the different crops produced per year, and also for labor, fertilizer, machinery and other factors of production. The estimated prices are linked to each of these quantities (except for the land itself). Then the project analysts take the sum total of the value of

all of the farm's outputs, and subtracts from it the estimated costs of all of the farm inputs, other than land and water. The result is the farm's profit (including return to land, independent of whether that land is owned or rented). This whole exercise is estimated for all the future life of the project: a) assuming the project is not built, and b) assuming the project is in fact undertaken. Thus we have, year by year over the project's expected life, a residual income with the project, and without it. The difference between these two flows for each year becomes the estimate of that year's expected value of irrigation water.

| | Cost of Inputs Other | | |
|--|-----------------------|----------------------------|-----------------------|
| | <u>Value of Crops</u> | <u>Than Land and Water</u> | <u>Residual Value</u> |
| With Project | 1500 ± 150 | 1000 ± 100 | 500 ± 250 |
| Without Project | 1000 ± 100 | 700 ± 70 | 300 <u>±170</u> |
| Difference In Residual Value Attributed To Water | | | 200 ± 420 |

One can easily appreciate the complexity of such estimation procedures, and particularly the degree of error or uncertainty that is involved each step.

The above table reveals the Achilles' heel of the residual value method. I have made the table simple by dealing with "plus or minus" ranges, rather than standard deviations and variables, because it is easier to understand for those who may not have a background in statistics.¹

¹In statistics one learns that the variance of the sum of two independently distributed variables is the sum of their respective variances. Most students are then quite surprised to learn that the variance of the difference between two such variables is also the sum of their variances. Reworking the table in these terms, suppose that the numbers indicated by ± for "value of crops" and "costs of inputs" are standard deviations. Then the variance of the "with project" residual value would be $22,500 + 10,000 = 33,500$, and the same variance without the project would be $10,000 + 4,900 = 14,900$.

The estimated project benefit in the table is 200 (= 500 minus 300). The variance attaching to that 200 figure would, using the numbers just calculated, be 48,400, and its standard deviation would be 220. This is not as huge as the ± 420 emerging from the table, but the point to be made is exactly the same. When the value we calculate is based on a "difference of differences", that value is subject to a very large standard error, larger than any of the components from which the differences were derived.

The residual value that we calculate in this way is a combined estimate of the contributions of both land and water to farm product. The table shows an extra 200 (for a given year) to be contributed by the irrigation project. Suppose now that the irrigation authority itself collects this amount from the farmers as an irrigation charge. Then, clearly, there would be no reason why land values should rise. But if no irrigation charge were collected, land values would presumably rise by the full present value of this “difference in residual value attributed to water” for all future years of the life of the project. In the simple case of a perpetuity of 200, using a 10% discount rate, land values would rise by 2000 ($= 200/.10$). If the irrigation authority were to collect 80 per year, then the rise in land values (in the same simplified calculation) would be 1200 ($= 120/.10$).

Thus we see the folly of counting both the value of the water and the rise in land values if that were to mean counting 200 per year as the value of the water plus 2000 (= the same 200, capitalized at 10%) as the rise in the value of the land. It would not be a mistake to calculate the benefit as the actual irrigation charges collected (80 per year in the example) plus an induced rise in the land value of 1200 (based on the 120 of irrigation benefit not covered by cash payments for water).

The Indian project analysis that I mentioned at the outset also counted the increase in value of crops. This has no business appearing as an extra benefit (of 500 in the table), because there are clearly extra costs involved (of 300 in the table). When we got the residual value, we actually counted the increment to crop value as a plus, but then we deducted the associated costs, as we should do. To count the residual value and then add to it the increment in crop value is clearly double counting and is completely unjustified.

When it comes to the increment to employment resulting from the project, the basic lesson is the standard one in labor economics. The wage represents “in principle” an economic cost, not a benefit. Sometimes the wage might overestimate the true economic cost of the labor in question, in which case one would consider that, after first counting the full wage payment as a cost, we would introduce an external benefit (e.g., for taxes collected on the basis of those wages, and/or for a “producer surplus” representing the excess of the actual net-of-tax wage payments over the true supply price of the workers involved).

Working with the data of the table and summarizing the worst-case interpretation of my Indian example, we have

| | |
|-----------------------------|--|
| Value of Irrigation Water | = 200 per year |
| Increment to Value of Land | = 2000 (200 per year capitalized at 10%) |
| Increment to Value of Crops | = 500 per year |
| Increment to Labor Use | = 180 per year (say, 60% of the 300 increment to costs). |

What we should have is only the first of these, which could be correctly represented as an 80 per year actual payment of irrigation water charges, plus 120 per year of “economic rent”, generated because the full value of the water was not being collected. This 120 could be directed counted as an economic rent (the preferred way), or capitalized as an increase of 1200 ($= 120/.10$) in the value of land. On labor we should never count the full wages bill as a non-cost. The correct

procedure would either take simply the voluntary net-of-tax supply price of the added labor as a cost (to which no externalities could be appended), or one could initially take the full outlay on labor as a cost (as in a financial analysis), and then consider external benefits equal to the extra taxes paid in connection with those wages, plus the estimated producer surplus received by those workers. In a well-functioning labor market, the project would not generate significant producer surplus because the workers would be expected to have other employment in the absence of the product.

Direct Estimates of the Value of Irrigation Water

Many years ago I was asked by the Argentine authorities to lead a very small team to do an evaluation of an irrigation project (the Ullum Dam) on the San Juan River in western Argentina. The team consisted of two former students, Lucio Reza and Juan Antonio Zapata, and myself. A major engineering firm, the Harza Engineering Company of Chicago had been active in the evaluation process for several months; they had been doing farm budget studies and were asked to make their data available to us. But before we reached that stage, we wanted to familiarize ourselves with the economics of agriculture in the region, and the role of irrigation water within agriculture.

Luckily, Reza had previously spent a couple of years in the city of San Juan, as the local representative of Argentina's Department (Secretariat) of Agriculture. He thus knew many key people in the area, so we installed ourselves on the terrace of our hotel for a couple of days, receiving a steady stream of local experts.

Early in that process, we received a briefing on the existing system of river irrigation, which had been in existence for more than half a century. During that briefing, we were told that some 120,000 hectares were covered by the project, and that the available river water for each month was distributed by two technicians (called tomeros), one on each side of the river. They worked their way up or down the river, opening the sluice gates of one farm after another, then closing them after each farm's water quota had been delivered.

The following day, another interviewee was talking about irrigation matters, and in the course of his conversation casually mentioned "the 60,000 hectares or so that are irrigated". That tiny phrase opened the door to our entire evaluation of the project. We asked him, is it not true that the project covers 120,000 hectares, not 60,000? Yes, he said, 120,000 hectares have irrigation rights, but the water that is delivered to each owner is actually used on about half of his eligible hectares. Why is that? Because putting 4 inches of water on one hectare produces more than one would get from putting 2 inches of water on each of 2 hectares. That is to say, the economically optimal strategy was to leave half the eligible area without irrigation water. What this meant, in this particular case, was that land was super-abundant; and its marginal productivity was essentially zero. The scarce factor was water, and if one had a residual value due to "water plus land", that value arose because of the scarcity of water alone.

This insight, which every student of elementary economics has learned (or should have learned) became the key to our entire study. Right away we asked whether water rights could be bought and sold -- the answer was no. Then we asked whether an owner of two pieces of land

could use the water rights of one piece to have the water delivered to the other. The answer here was yes, provided the two pieces of land were on the same side of the river. Why was that? Because it would be too complicated to have water taken away by the left-bank tomero and then delivered by the right bank tomero. If any “transfers of water” were to take place, then, both parts of the transfer had to be handled by the same tomeros.

Our next key question was, could one find among the land sales of recent years, any properties where the principal purpose of the buyer of plot B was simply to have “its” water transferred to that buyer’s existing farm A? The answer was yes, there had been quite a number of such transactions. We then had these candidate transactions scanned, so as to omit any on which there were important non-land assets, such as houses or barns, etc. This culling was necessary, for our plan was to use the land prices at which our plots were sold, as estimates of the current market value of their water rights.

Next came the question, taking for granted that the buyers of these properties were really buying water rights, what kind of product were they actually paying for? Certainly they were not getting a certain number of cubic meters of irrigation water, every year, for sure. What they were really buying is what could be characterized as a series of lottery tickets, one for each month of each future year. When the previous winter’s snowfall in the Andes was big, they got a lot of water, when that snowfall was small, they received little water.

We went back to the irrigation records and developed a histogram showing a frequency distribution of water deliveries in each of the past 50 years. That distribution had quite a range, with the maximum water availability being something like 10 or 15 times the minimum. It was obvious that the farmers would not value very highly amounts of water that would come only once every 10 or 15 years. On the other hand, water that they could pretty much count on was extremely valuable. If this amount were to increase by half, they could plant half again as many hectares to their main cultivations -- namely, vineyards and olive trees.²

Given the water delivery experience of the past 50 years, we had to somehow take account of the fact that “sure water” was much more valuable to the farmers than “occasional water”. We had little time for nuances, so we adopted a quite robust scheme of “weighted cubic meters” of irrigation water. We started with an index $I_1 = .4D_1 + .3Q_1 + .2M_e + .1Q_3$, where I_1 is index #1, and D_1 is the first decile of the histogram, Q_1 the first quartile, M_e the median, and Q_3 the third quartile. This index obviously gives much heavier weight to sure as against occasional water, but it is also arbitrary, and we had no time to write a Ph.D. dissertation examining what would be the appropriate weights. So we resorted to a trick that is standard fare for cost-benefit analysis -- a “sensitivity test”. In this case the sensitivity test entailed employing

² Both of these plantings last for many years, but both will die out if deprived of water. The way it works is that in years of minimal water availability, the farmers spread that water very sparingly, trying simply to keep their plants alive. Little or no harvest (of grapes or olives) comes out of such years. It makes no sense to plant olives and vines and if they are going to die off in a few years, so total plantings are largely governed by the expected bottom end of the probability distribution of water availability.

an alternative index, $I_2 = .33D_1 + .27D_2 + .23M_e + .17Q_3$, which gave less weight to “sure” water and more weight to “occasional” water. We carried out parallel calculations using both I_1 and I_2 , all the way to the end of our study. Happily, our conclusion -- that the dam was indeed a worthwhile investment -- was the same, regardless of which of the two indexes was used.

We next set $120,000 P_H = P_{I1} I_1^0$, where P_H is the recent price (in real pesos) of land in our key transactions, and I_1^0 is the first index calculation that we derived from the histogram representing the irrigation experience of the past 50 years. This gave us a price P_{I1} for the unit of quality-adjusted water measured by I_1 . A parallel calculation was then carried out to give on P_{I2} , the price of quality-adjusted water as measured by I_2 .

This prepared us for the next step, of trying to assess the value of the Ullum Dam project. This required us to have a reasonable projection of how water availabilities would change if the dam were built. Our procedure here was to simulate how the dam would have functioned, over the past 50 years, if it had been in existence all that time. To do this we obtained month-by-month streamflow data, plus month-by-month actual irrigation deliveries over this 50-year period. For the simulation, we followed a very simple strategy, something that was a virtual necessity due to our time constraints. We divided the year into two seasons -- irrigation and non-irrigation. The strategy was to accumulate water behind the dam during the non-irrigation season, and deliver it during the irrigation season. The amount accumulated each month (in the simulation) was the full amount by which that month’s streamflow exceeded the amount that had to be left in order to cover the water rights of downstream users. This accumulation was allowed to go on, during the entire non-irrigation season. However, for past years of abundant water, the accumulation had to stop, once the dam’s capacity of 440 cubic hectometers was reached.

For deliveries from the dam, we again needed a simple strategy. The one we chose was to assume that water deliveries from the dam would simply be used to proportionally increase the natural streamflow of each irrigation season, month by month. Here we reached a different limit at the point where this strategy called for deliveries over and above the delivery capacity of the canals. In such cases our simulation saved the excess water (above and beyond canal capacity) for the next irrigation season.

This simulation resulted, quite obviously, in much greater deliveries of irrigation water than what had actually occurred during the past 50 years. Then, using the simulated data, month by month and year by year, we were able to develop two new histograms I_1^* and I_2^* representing what values I_1 and I_2 would have taken, had the dam been in existence for the past 50 years. These procedures yielded one component of the benefits of the dam. This is represented by $P_{I1} (I_1^* - I_1^0)$ for the first index and $P_{I2} (I_2^* - I_2^0)$ for the second index. This measure assigns a value to the “dam water” which is equal to that which the market assigned to “river water”.

Our next step was to recognize that dam water is more valuable than river water. This is because the dam managers have some degree of control over when the stored water will be delivered to the farms. Obviously they will try to time their deliveries so as to come as close as they can to giving farmers water at the times they want it most. Thus, they will certainly not emulate our simulation, by simply giving farmers a proportional expansion of each year's natural streamflow. They will make significantly better use of the water. How much better -- we really do not know. So we again made alternative assumptions, both of them extremely conservative. Under the first of these, dam water was assumed to be 5% more valuable than river water, and under the second it was taken to be 10% more valuable. However, this increment of value would apply not just to $P_{I1}(I_1^* - I_1^0)$ or to $P_{I2}(I_2^* - I_2^0)$, but rather to $P_{I1} I_1^*$ or to $P_{I2} I_2^*$. Why?

Because once the dam is there, it can manipulate the timing of all deliveries during the irrigation season, the only limitation coming on those rare occasions where the dam is already full, in which case the natural streamflow must be delivered, else that water would be wasted, as far as the project area is concerned.

To sum up, our measures of the benefits of the dam, up to now, are

$$B_{1a} = P_{I1}(I_1^* - I_1^0) + .05 P_{I1} I_1^*$$

$$B_{1b} = P_{I1}(I_1^* - I_1^0) + .10 P_{I1} I_1^*$$

$$B_{2a} = P_{I2}(I_2^* - I_2^0) + .05 P_{I2} I_2^*$$

$$B_{2b} = P_{I2}(I_2^* - I_2^0) + .10 P_{I2} I_2^*$$

We are still not quite finished. P_{I1} and P_{I2} were derived on the basis of the observed real P_H , which should be seen as the private discounted value of the private benefit that farmers currently get, per hectare of irrigation rights, from streamflow irrigation. Two corrections are called for if we are aiming at the overall economic benefit of the dam project. The first is to adjust the benefit stream upward to take into account the estimated increase in property and income tax revenues that will result from the project. The second is to recognize that the appropriate discount rate for calculating the economic net present value of a project is not the private discount rate but rather the overall economic opportunity cost of capital. A rough adjustment of this type is to multiply the tax-adjusted benefit figure by the private discount rate, and divide it by our best estimate of the true economic opportunity cost of capital in the country.

In our actual study we made these two adjustments. Luckily, under all the alternatives we examined -- B_{1a} , through B_{2b} -- the present value of project benefits exceeded the present value of costs. The project seemed quite definitely to be worth undertaking.

We still were not done, however. There remained the possibility that we might somehow have overestimated the true benefits of the project. This would simply be an unavoidable risk if we had no way to check on our calculations. But in this case we did have a way, about which our early interviews on our hotel terrace had informed us. The key to this check was the fact that side by side with streamflow irrigation, there was a fairly wide use of pump irrigation in the area served by the dam.

The key fact here is that (assuming the chemical characteristics of the water do not differ greatly), pump irrigation water has to be better than water from the dam, for basically the same reason that makes dam water more valuable than river water. The reason is that farmers can pump water (up to the capacity of their pipes and pumps) when they really need it most, whereas with dam water they have to wait their turn as the tomero makes his monthly rounds. Thus they might get their water two weeks before they need it most, or two weeks after, and simply have to make the best of those deliveries from the dam. Thus, if we found in that area that people could get good pump irrigation water for less than the value we estimated for water from the dam, this would mean that our benefit estimates were too high.

Pursuing this line of thought, we made a rather careful survey of the situation with respect to pump irrigation in the area. It turned out that while the land surface sloped very gently downward from west to east, the aquifer containing ground water sloped downward more rapidly. This led to a situation where the earliest wells (those farthest to the west) had reached water at a depth of perhaps 30 meters. Those wells were already put in operation before 1920. As time went on and pump technology improved, the area covered by pumps moved eastward step by step, going to 50, then 75, then 100 meters, etc. At the time of our study, the new wells that were being drilled were at a depth of about 200 meters. Our inference was that investments in these wells were just yielding the prevailing private rate of real return on agricultural investments. Happily the resulting costs estimated for pump irrigation water were significantly higher than our own estimates of the value of water from the dam. Our procedures thus easily passed this final test.

* * * * *

I presented this step by step description of an actual real-world project evaluation exercise in an effort to transmit something of the flavor of such work. It is certainly not a routine exercise that anybody with a little instruction could perform. One has to be alert to the particulars of each case. One has to mesh the real-world observation of these particulars with the lessons of economic theory. Then one has to take the result and fit it into the framework of economic cost-benefit analysis, incorporating the key places where economic opportunity costs, and economic values in general will be different from the corresponding prices that we observe in the market.

How far can the lessons from this exercise be extended to a broader range of irrigation projects? I would say quite a distance. The key observation in our analysis of the Ullum Dam project was the finding that the prices paid per hectare in a certain set of agricultural land sales could be seen as in fact paying for the rights to irrigation water that were attached to that land. This enabled us to bypass the laborious farm budget method, with its Achilles heel of a very high variance surrounding its estimates. One could contemplate doing something similar in cases where one finds irrigated land side by side with non-irrigated land of similar quality, on both of which crops are being raised (probably different crops and different rotations in each type). The difference in per hectare prices between irrigated and non-irrigated hectares could in principle play the same role as the price of our key transacted hectares -- i.e., as a measure of the present value of the "series of lottery tickets" that one is in effect getting when one acquires rights to

irrigation water. If one did this, one would have to be quite sure that the two classes of land where prices were being compared really did have similar qualities of soil. Also, one would have to check concerning improvements that might tend to be present on already irrigated hectares and absent on those used for dry-land farming. Particularly important here clearly is land-leveling, which can represent a big capital investment that leaves no immediately perceptible trace. But in many cases irrigated land embodies much more of such investment per hectare than do neighboring unirrigated hectares. The presence of irrigation ditches themselves represent another capital cost not present in unirrigated lands. Also, there might well be fencing and farm buildings linked to one kind of land use and not to the other. Thus unirrigated lands used for pasture are often divided into separate fenced parcels in order to permit orderly grazing (and in particular to avoid overgrazing). In order to use the difference in land prices in the way we used the prices of our key transactions, one would have to estimate the contributions of each of these elements to the respective prices of irrigated and nonirrigated land. Only after correcting for these elements would one go on to take a price difference which we would then interpret as the private-market valuation of the present value of water rights. Obviously, if these elements get to play too big a role in the comparison one gets into the same problem as one encounters in the farm budget method -- that is, the high variance that applied to the difference between two separate elements, each of which is subject to considerable error of estimation.