THE ABCs OF ELECTRICITY PROJECT ANALYSIS

A report submitted under

Contract No. EEM-C-00-06-00008-00

by

Arnold C. Harberger
Chief Economic Advisor to USAID

8283 Bunche Hall, UCLA
405 Hilgard Avenue
Los Angeles, CA 90095-1477
Telephone: 310-825-1011
Fax: 310-825-9528
e-mail: **harberger@econ.ucla.edu**

July 12, 2010

TO: U.S. Agency for International Development
PPC/CDIE/DI -- Attn. ACQUISITIONS

Strategic Objective or Activity: 905-101/Economic Growth

Cognizant Technical Office: EGAT/EG
RRB, Room 2.11.123
1300 Pennsylvania Avenue NW, Washington, DC 20523

U.S. Agency for International Development
1300 Pennsylvania Avenue, NW
Washington, DC 20523
www.usaid.gov

THE ABCs OF ELECTRICITY PROJECT ANALYSIS

Arnold C. Harberger
Chief Economic Advisor to USAID

July 2010

This is the sixth installment of our Introduction to Cost-Benefit Analysis, the first five of which formed the basis of a short course sponsored by USAID in order to orient its cadres in this subject. This course was first given in December 2008 and has been repeated several times since. I began this presentation with the above notice in order to make clear to readers that such a course must necessarily concentrate on essentials -- by its very nature a course that is given in one week cannot replicate the job that would be done in a course that lasted a full year, or even a couple of months.

The bare-bones example of cost-benefit analysis is best represented by an individual enterprise, in which the owners have only their own money at stake. In such a case the benefits as seen by these owners are simply the monetary inflows from a project (including those from the disposition of its assets), and the costs are the monetary outlays that are made, both in the investment phase and in the operating phase of the project. The relevant discount rate, again from the owners' point of view, would be their own relevant opportunity cost of funds -- i.e., the rate of return that they could normally expect to earn on project funds in the natural alternative uses to which these funds would be put in the absence of the project.

As one moves from such a simple owner-operated project to examples that are more relevant for real-world public sector investments, one encounters some additional problems, both in valuing a project's costs and benefits and in establishing the discount

rate that is relevant for calculating a project's net present value. For the most part, the differences between the analysis of an owner-operated venture and that of typical public-sector projects are concentrated in the valuation of benefits. Most project costs are indeed cash outlays, just like those of a business venture, but this is not at all generally true on the benefits side. For example, public parks and highways (other than toll roads) rarely yield any cash inflows. The problem then is to find ways of estimating the true economic value of their benefits. In other cases (e.g., irrigation projects and toll roads), there is usually some charge for the use of the project's output, but that charge is typically a very poor measure of the project's true benefits. Again the challenge is to measure the true benefits of the project.

Electricity projects appear to be in a different category, in that one hardly ever sees attempts to measure the actual benefits that users receive from such projects. Yet paradoxically, we still say we are quantifying the value of such benefits. The explanation of this apparent anomaly lies in what is called the "least alternative cost" principle. This principle states that one should not attribute to a project a value of benefits that is greater than the least alternative cost one would have to incur by providing an equivalent benefit stream in a different way.

This principle is fully general, but often seems quite redundant. Thus irrigation projects provide certain flows of water to a farming area, but with most of them one cannot even dream of a sensible alternative way to provide the same flows. In those cases the alternative cost (say of bringing the water in by truck) is so high as to be irrelevant in the analysis of a project to draw water from a nearby river, so one tries to put an economic value on the river water itself, as it is used for irrigation. However, one

does encounter irrigation projects in which pump irrigation (from underground aquifers) is a quite reasonable alternative to drawing water from the river. In these cases, it would be incorrect to attribute to the river-irrigation project a benefit stream that exceeded the alternative cost of getting the same water by way of pump irrigation.

While the least-cost principle can thus sometimes come into play for irrigation projects, one can say that it virtually always is the determining factor in estimating the benefits of electricity projects. The reason is that a next-best alternative of reasonable cost nearly always exists. Indeed, much of the time this next-best alternative is simply the standard way of doing things. The projects being analyzed in such cases then represent attempts to find new or different ways of doing things that are better than the standard alternative. The benefits of the "new or different" way are measured in such cases by the standard costs that would be saved, if the "new or different" project is in fact undertaken.

To see the deeper meaning of this approach, consider what would happen if we somehow left the standard alternative out of the picture. In the example that I have used for years in my classes, I consider an electricity project that is really bad -- much more expensive than the standard alternative. Now let us for the moment think of carrying out a cost-benefit analysis of this project without bringing the standard alternative into the picture in any way. Inevitably, this leads us to think of "this" project as somehow standing alone. At the moment when it is installed, it will presumably be the newest project in an existing system. But over time the older plants in this system will wear out, so the generating capacity of the system will decline year after year, as, one after the other, the older plants are abandoned. The overall generating capacity of the system

would thus steadily decline over time. Even without a growing demand for energy, this would mean a market or economic price of electricity that would be steadily rising. In the more likely case of a continuously growing demand, this upward trend of price would be even more exaggerated. It would take a truly terrible project to fail a cost-benefit test, when its output was being valued at prices that were increasing exponentially throughout its economic life. One can almost say that cost-benefit analysis carried out under these assumptions would lose virtually all of its power to discriminate between good and bad projects. All would look good in the face of an ever-rising price of energy.

Now let us return to the real world. It obviously makes no sense to assume that our project -- call it plant E -- is going to be the last project to be built in our city's electricity network. We simply have to think of the system being operated in a sensible way. In the beginning plant E is added to an existing system consisting of plants A, B, C, and D. Over time plant A, the oldest, is likely to be the first to be retired. It will perhaps be "replaced" by plant F. But by that time the area's energy demand would likely have grown by enough to justify the addition of yet more capacity, say plant G. So, maybe 5 or 10 years down the line, the system would likely consist of plants B, C, D, E, F, G. Further on, plants B and C will presumably also reach retirement age, and maybe plants H, I, J, and K will be added. The final step (in the analysis of "our" project -- plant E) will come when plant E itself reaches the point of being retired from the system. At that point the cost-benefit profile of project E would come to an end, perhaps with a blip of extra benefit representing the salvage value of the plant, perhaps with a blip of cost (e.g., for a nuclear plant) for the safe disposal of its remains.

The image I have tried to conjure up here is that of a motion picture representing the costs and benefits attributable to plant E, not standing alone, but imbedded in a system which is being managed intelligently, with other plants being retired when their staying in the system would entail more cost than benefits, and with new plants being added in a pattern that reflects the continuing use of cost-benefit principles. All of this lies behind the development of our basic tool of analysis, the "moving picture" of how the system would operate in the presence of our project, i.e., "with" the project E.

But this is not the end. In order to get the cost-benefit profile of project E, we have to make a forecast of how the system would operate in the absence of this project. In this scenario we do not do project E, but follow some alternative strategy in managing the electricity network. What strategy? There are only two good answers here -- a) the best alternative strategy, if we are able to identify such a strategy in specific terms, or b) a "standard" alternative, defined by our best estimates of the typical costs of energy (varying by time of day, season of year, etc.) that we consider would emerge from a proper continuing application of sound cost-benefit analysis.

Answer a) is likely to be feasible, if at all, only in sophisticated modern electricity systems, whose operations are governed by up-to-date computer systems designed to take into account all relevant factors in order to come up with a minimum-cost strategy for the system as a whole. More likely is answer b) which is based on a more general knowledge of the costs of equipment, fuels and other material inputs, labor and other services, etc. This is the line that we will explore in this paper, using examples that move progressively from the simplest to the more complex.

The Simplest Case -- A Homogeneous Thermal Alternative

This section is meant to introduce readers to some very basic aspects of electricity economics. It should be thought of as dipping your toes in the water, not as a full-body immersion. In this exercise we will have one standard alternative -- a homogeneous thermal generator. By homogeneous I mean that the actual machines used in plants A, B, C, D, and E would all be physically the same (though of different ages), assuming all of them to be thermal plants (using fuel to generate energy). We will derive costs per kilowatt hour (kwh), based on the use of this standard generating equipment.

Assume we have data telling us that, given the current fuel price, the operating cost of this standard piece of equipment amounts to 4¢ per kwh. This mainly covers fuel, but it also takes into account the labor and other inputs involved in the actual operation of the equipment. It definitely <u>does not</u> include any return to invested capital. This will enter our picture at a later stage.

For now, let us simply concentrate on the idea that 4¢ per kwh is the appropriate cost of energy, measured at the plant, when that energy is being produced during off-peak hours. Why do we not add a charge for the use of the generating equipment itself? Simply because the simple addition of some extra kwh of output during off-peak hours does not require any more capital equipment than we already have.

When does the system require additional homogeneous thermal plants? Quite naturally, when demand threatens to push beyond the level of energy that our existing plants can deliver. We measure that capacity in kilowatts (KW), and also use the term "power" to refer to KW. Power in common parlance is something we can have without using it at all, and certainly without using it fully. But economists ask, when a system

has more capacity (power) than it needs to satisfy demand, why charge for the use of that capacity? Nobody loses anything if idle capacity is put to use, while further excess capacity still exists. Thus, economists say, when idle capacity is present, the appropriate charge for energy would cover running costs (variable costs), not capital costs (fixed costs).

This line of thinking naturally leads to the idea of electricity charges that vary through time, being higher during peak-time hours and lower during off-peak hours. Charges linked to capital equipment (generating capacity) should appropriately be concentrated during peak-time hours, because if demand increases significantly during these hours, one will actually have to install additional capacity if that demand is to be met. And if one does not install new capacity in such a case, it will take an increase in the price of peak-time energy in order to constrain demand within the limits of the existing capacity.

An example will probably help to clarify how these concepts are actually used. Suppose that new capacity (of the homogeneous thermal variety) costs $800 per KW, that the relevant discount rate $\underline{r}$ is 10%, and that the relevant depreciation rate (δ) for this equipment is 5% per year. Then, in order to justify the addition of a new KW of capacity, the necessary benefit is (.15)($800), or $120 per year. If capacity is added that cannot generate such an annual benefit, the investment is that capacity is not justified. Keep in mind, then, that $120 per year is the target revenue that should be expected, from new increments of capacity. How can one think of getting this revenue? From the sale of energy during peak-time hours. Thus, if the system's peak is of 3000 hours per year, the

needed peak time surcharge would be 4¢ per kwh. And if the peak were 2000 hours, the relevant peak-time surcharge would be 6¢ (= $120/2000) per kwh.

Let us proceed on this latter assumption. Our "standard costs" for energy are 4¢ per kwh, off-peak, and 10¢ (= 4¢ for operating costs plus 6¢ of peak-time surcharge) during 2000 peak hours. Out next step is to apply this assumption to different types of hydroelectric projects. (Recall that we have only one type of thermal capacity, reflected in the modifier "homogeneous"). We will consider, in turn, run-of-the-stream hydro projects, daily reservoirs, and seasonal hydrostorage projects.

Run-of-the-Stream Hydro Projects

The key characteristic of run-of-the-stream projects is implicit in the title—energy is generated using river water "as it flows". Typically, a run-of-the-stream project will be situated on an incline, where water is flowing down a hill, or over a waterfall. Such projects typically channel the water though large tubes (penstocks) which carry it from the top to the bottom of the incline, and which lead directly into one or more turbines at the bottom of the hill. The running water turns the turbine, generating electric energy.

To evaluate the benefits of such a project one typically starts on the purely hypothetical assumption that the turbine capacity of the project will be fully used, through all the 8760 hours of the year. We then divide these hours into 2000 of peak-time and 6760 of off-peak use. Employing this information, we get, for each KW of turbine capacity

2000 hours @ 10¢/kwh =	$200
6760 hours @ 4¢/kwh =	$270.40
Total	$470.40 per KW

Now $470.40 per KW per year is what the installed capacity would produce, if it were fully used all the time. Thus, of course, is not at all likely to be the case, as streamflow always varies quite significantly, mainly by season of the year, reflecting changes in rainfall and/or snowmelt. Thus, on average, the installed turbines will be used at only a fraction of their full capacity. Simply for our example, we will assume this fraction to be 60%. Hence

$$\text{Estimated benefits} = .6 \times \$470.40 = \$282.24$$

We introduce these benefits into our profile of the run-of-the-stream project, deducting capital costs during the construction phase of the project and maintenance plus (very minor) operating costs during its operating phase. The result is a project profile which we then can evaluate, using, of course, the same discount rate (here 10%) that we employed in the derivation of the 6¢ peak-time surcharge.

* * * * *

Readers may have noticed that in the above exercise, the analysis was carried out "per kilowatt (KW)" of installed capacity, and complemented by an assumption of the fraction (in this case 60%) of that capacity which was expected to be utilized in the course of a typical year. But as one focuses on this feature, one quickly comes to wonder, why 60 and not 40 or 80 percent? It is to this question that we now turn.

First, let us recognize that in any real-world case, the answer will depend on the hydrological characteristics of the stream (and site) in question. There may be rivers whose streamflow is so steady that there is only a 20 or 30 percent difference between the lowest and the highest daily flow during the year. In such a case there is not much range for choice as to how many KW of turbine capacity to install. But such cases would be

hard to find in the real world. Most rivers and streams are subject to very heavy streamflow in the rainy season (or the period of biggest snowmelt). Some even dry up completely in the driest period of the year, and for most the lowest streamflow is only a modest fraction of the highest one.

Thus the designers of a typical run-of-the-stream project are faced with a serious problem of choice. If they build the project so as to make use of nearly all the streamflow of the year, they will have to install enough turbine capacity to process the huge rainy-season flows. But then, during the rest (which is most) of the year, the much lower streamflow will leave most of the installed turbines idle for many months running. On the other hand, if the designers decide to keep the fraction of capacity use high, they will have to install turbine capacity geared principally to the rate of streamflow in the drier part of the year. They will then end up using their turbines most of the time, but they will be allowing a lot of the stream's annual waterflow to go to waste. The dilemma is -- build big, and a lot of that turbine capacity will be idle a lot of the time; build small, and your problem will not be one of idle turbine capacity but rather one of a lot of water passing by unused (for electricity generation), simply because you don't have the turbines to process it.

This dilemma represents an economic problem -- one of weighing benefits against costs. This problem is best tackled at the design stage so as to ensure that whatever choice is made as to how much turbine capacity to install, that capacity is achieved at the lowest economic cost. The key facts needed for solving this problem are: a) a graph of the expected (likely) streamflow, period by period (perhaps day by day) throughout the year. This graph will most likely have a single peak sometime in the wettest season and a

single trough (sometime in the driest period). The benefits of adding turbine capacity get smaller and smaller, as we contemplate adding successive increments of capacity. The first few KW of capacity will promise 100% usage, as long as they involve using less than the lowest expected streamflow. On the other hand, the last few KW of capacity that we might add, would promise usage only in the few days of absolutely highest streamflows. Very likely, then, benefits will amply exceed cost for the first few KW, while cost will almost certainly exceed benefits for additions to capacity that can expect to be utilized only a few days per year. Somewhere between these extremes, then, we should be able to find an optimum level for turbine capacity -- a point up to which benefits exceed costs for each successive increment to design capacity, and beyond which costs exceed benefits for each successive increment. This is the sort of calculation that should be carried out in the process of designing any run-of-the-stream project. Obviously, it involves repeated applications of the procedure outlined in the first part of this section, with the ultimate choice being for that turbine capacity which yields the greatest expected net present value (i.e., the greatest expected excess of the present value of benefits over the present value of costs).

Daily Reservoir Hydro Projects

The daily reservoir can be thought of as a sort of add-on to a run-of-the-stream project, either at the design stage, or later. Here we will assume that we are dealing with a run-of-the-stream project that is already operating. Moreover, we will assume that this existing project is well-designed for its purpose, following the principles presented above.

In our example of a run-of-the-stream project, the installation was expected to produce peaktime energy during 2000 hours, and off-peak energy during the remaining 6760 hours of the year. This means that over the course of the year, around 22% (= 2000/8760) of the waterflow would be used to produce electricity that was worth 10¢ per kwh, while the remaining 78% (= 6760/8760) would be used to produce energy worth only 4¢ per kwh.

A daily reservoir project has two principle objectives: a) to convert to peaktime production much of the water that would normally go to produce offpeak energy in a run-of-the-stream operation, and b) to utilize some of the water that would otherwise go to waste (again in a run-of-the-stream case). To accomplish objective a), a small dam is built upstream of the ROS project. It accumulates water during the offpeak hours, and then releases that water during peaktime. In this way it produces 10¢/kwh energy with the same water that would otherwise end up generating energy worth only 4¢/kwh. The substantial net gain of 6¢/kwh is the principal benefit of the daily reservoir project.

This benefit comes with a significant cost, however. First and foremost, there is the cost of the dam itself, and possibly of a regulating dam downstream of the project, designed to deliver a steady streamflow to downstream users. If the amount of turbine capacity in the project is not increased, the benefit of extra peaktime energy would be limited to the amount by which the turbine capacity of the ROS project exceeded the streamflow-determined ROS output of each day. Thus a project with 5 MW (= 5000 KW) of turbine capacity might in one part of the year be processing a streamflow that generated 2000 kwh during each hour of peak. This could be brought up to 5000 kwh per peaktime hour, but not further, if one left turbine capacity unchanged. All this benefit

would be of type a). Obviously, in periods of the year where the streamflow itself was enough to generate 5000 kwh/hr, the installed capacity would be fully utilized by the ROS project, and the addition of the daily reservoir would have no effect of shifting water (and therefore energy output) from offpeak to peak hours.

Thus, assume

T_o = initial turbine capacity, in KW

S_i = expected streamflow per hour of day i expressed as the number of KW which that streamflow can generate.

Hp_i = number of peaktime hours on day i

Hn_i = number of offpeak hours in day i (= 24 - Hp_i).

Then we have $Hp_i(T_o-S_i)$ = the maximum number of kwh that can be converted from offpeak to peak, if T remains unchanged. This is a maximum, because in periods of very low streamflow, one may not reach an output of Hp_iT in peaktime, owing to low streamflow throughout the day.

The amount of water available to be shifted from offpeak to peak is simply $H_{ni}S_i$, the total waterflow in offpeak hours. All this would represent offpeak energy actually produced by the pre-existing ROS project, if $T_o > S_i$. In case $S_i > T_o$, the difference $(S_i-T_o)H_{ni}$ would represent water that passed by without producing any electricity in the original ROS project. So if the increase in generating capacity ΔT is such that $Hp_i(T_o+\Delta T) > 24S_i$, that means that, with a capacity increase of ΔT, the full streamflow of the day (= $24S_i$) can be processed during peaktime hours. This would include a shift of all the water that had previously gone to produce offpeak energy, plus

all the water that had “gone to waste” because of limited turbine capacity. The benefit for day i of the DR project would then be the excess of this value over the value of the energy that the original ROS project would have produced.

On days when $(T_o+\Delta T)Hp_i$ is smaller than $24S_i$, the augmented project will not utilize the full streamflow of the day at peaktime, but may well use it during peak plus offpeak hours. In this case the gross benefit for day i of the DR project would then be measured by $Hp(T_o+\Delta T)$ times the peaktime price of energy, plus $[24S_i-Hp(T_o+\Delta T)]$ times the offpeak price of energy. The term in square brackets represents the water that is available $(= 24S_i)$ minus that amount which is used to generate peaktime energy $[=Hp(T_o+\Delta T)]$. From this, as before, we would have to subtract the value of the energy that the original ROS project would have produced.

Just as, for the ROS project, a separate optimization had to be made to determine the optimal level of T, so here a similar process should be used to determine the best level for ΔT, the increment of turbine capacity.

* * * * *

Special note should be taken of the fact that no considerations of electricity demand entered into the above analysis, either of ROS or of daily reservoir projects. The reason for this is the way in which such projects fit into the operations of most electricity systems. As we will see in greater detail later, the principle governing the management of electricity systems is that when demand is low, one uses only those sources of energy that have the lowest running cost per kwh. Then, as energy demand increases, additional capacity is turned on, starting first with the second cheapest per kwh, then turning on the

third cheapest, and so on up the scale. Only at times of very high (peak) demand does the system resort to its generators with the highest running cost.

The mere contemplation of how an ROS system operates tells us that it has practically zero running costs. All that is needed is for somebody to make sure that the water that is flowing in a given hour is actually channeled through the turbines to generate energy. Because ROS energy is so cheap, it is always, in principle, the first source to be used in cases of low demand. And it is used every hour of the day, all through the year, being interrupted only for maintenance and repair. Since ROS capacity constitutes only a small fraction of the total capacity of a typical electricity system, it is in fact used all the time and is practically never used only partially for lack of demand for its energy output.

In the case of daily reservoirs, once again their running cost is next to zero, but since the great bulk of their output is at peaktime, they are then working along with all or nearly all of the other sources of energy in the system and of these they (along with ROS installations) have the lowest running cost. Any variation in peak demand will thus be absorbed by other, higher-cost contributors to the supply of peaktime energy. This explains why demand considerations did not enter into the preceding analysis of daily reservoir projects.

Seasonal Hydro Dams

While daily reservoirs have the effect of letting managers decide when, within a given day or so, water will be needed to generate energy, seasonal hydro dams aim at allowing such water us to be shifted from one part of the year to another. The typical case for which a seasonal dam will be contemplated is one where there is one season in

which the energy from a given stream would have a high value, and another in which that value would be much lower. Some cases thus could arise because streamflow is very heavy in one part of the year and very light in another, while energy demand is pretty steady over the year. Other cases could arise in which the streamflow is pretty steady but demand is highly concentrated, perhaps in winter for lighting and heating, perhaps in summer for air conditioning.

As in the preceding sections, we will analyze seasonal hydro dams as part of a system in which the "standard" way of generating electricity is via homogeneous thermal capacity. Again, this standard capacity will be assumed to have a running cost of 4¢ per kwh, and a capital cost of $800 per KW, with a depreciation rate of 5% per annum. The relevant discount rate for cost-benefit analysis will, as before, be 10%.

The first question to be answered is whether the energy used in a seasonal storage project should be considered as "baseload" or "peaking" capacity. We have already seen how run-of-the-stream capacity is quite naturally "baseload", while the whole reason for building daily reservoir capacity is to augment the system's supply of peaktime energy. In our simplified case of homogeneous thermal capacity, we would have what is called a "stacking pattern" in which ROS capacity sits at the base, homogeneous thermal capacity occupies the middle, and daily reservoir capacity occupies the top. This means that when demand is very low, only ROS capacity will be used. When demand exceeds ROS capacity, that capacity will first be fully used, and it will be supplemented as needed by the output of our homogeneous thermal plants. Only in the hours of greatest demand during the day, will the water accumulated in daily reservoirs be used to "top off" the energy supply coming from run-of-the-stream and homogeneous thermal sources.

Now we come to the question at hand -- what place in this stacking pattern should be occupied by seasonal hydro capacity? Since we already know that ROS capacity belongs in the base, and that daily reservoir capacity should be used for peaking, we can concentrate on the question of which, as between homogeneous thermal and seasonal hydro capacity, should be turned on first. In particular, should seasonal hydro be thought of as part of the base, or as an appropriate way to serve peaktime demand?

The best way to focus on this question is to assume that we have one or more seasonal hydro dams already built. They will accumulate water in the wet season, and deliver energy in the dry season.[1] To make our analysis quite clear and straightforward, we must bear in mind that the storage capacity of our seasonal dams will not change, depending on our decision of how to use them. The amount of water they can store was determined when they were built. However, the amount of energy they can generate in any given hour typically is subject to change, because such dams are designed to leave room for adding turbines (up to some limit). In a simplified system consisting of seasonal hydro plus homogeneous thermal capacity, we would have a given level of peak demand, say 1000 MW (one megawatt - MW - equals 1000 kilowatts). If our seasonal hydro capacity is used as baseload, its stored water will be spread over 24 hours a day, for, say, 9 months of the year. To use the water in this fashion perhaps only 200 MW of turbine capacity will be needed (because this capacity will be running virtually continuously.) If on the other hand, the seasonal dams are used for peaking, they may be occupied for many fewer days, and for varying numbers of hours on these days. The

---

[1]Alternatively, we can think of them as accumulating water in the season of low electricity demand (say winter) and delivering it in the season of high electricity demand (say summer).

amount of turbine capacity needed in this case will be much larger, say 600 MW. So we have case A, where seasonal hydro is used for baseload, in which we would have turbine capacity of 200 MW by hydro, supplemented by 800 MW of homogeneous thermal. Alternatively, for case B, we would have to install 400 MW of homogeneous thermal capacity, plus, as indicated, 600 MW of turbine capacity in our hydro dams. As we move from case A to case B, then, we are subtracting 400 MW of homogeneous thermal capacity, and adding 400 MW of turbine capacity in our seasonal hydro dams. Here the cost-benefit analysis is a no-brainer. First, it is much cheaper to add 1 MW of turbine capacity to an existing dam with a place already prepared for additional turbines, than to add 1 MW of homogeneous thermal capacity, which entails building the whole plant plus its associated turbines. And second, using homogeneous thermal capacity for peaking will involve start-up and shut-down costs, which are practically zero in hydro dams, where turbines can be turned on or off simply by pressing a button or flicking a switch.

How seasonal hydro capacity will be used during the wet season depends mainly on physical conditions such as streamflow and the storage capacity of the hydro dams. Start with the idea of using these dams for peaking all year round -- including the wet season. This would mean that excess streamflow -- over and above that needed to satisfy peak demand, would be stored for later use. But suppose this would lead to the dams being filled before the wet season ends. In such a case it is much better to use the excess water in offpeak hours, rather than allow it to go to waste. This could mean that some thermal plants would be shut down for part or all of the wet season, their place being taken by turbine capacity in the hydro dams.

There is one additional point to be made with respect to the seasonal hydro versus thermal tradeoff. In our numerical example, we had 600 MW of seasonal hydro turbine capacity plus 400 MW of thermal capacity as our preferred solution. Under these circumstances it is clear that for most of the hours of the year when hydro capacity is being used, the system's thermal plants will be operating at full capacity. The resulting number of hours of full-capacity thermal operations can be called the "thermal peak". This can easily be a lot larger than the "system peak" or "demand peak" which we earlier assumed to be 2000 hours per year.

Now let us backtrack and ask, what is the logic behind our earlier derivation of a 6¢ per kwh peaktime surcharge, to be applied over a system peak of 2000 hours per year. That logic was that it was the growth of demand at peaktime that called forth the need for more (homogeneous) thermal capacity -- hence the scarcity value of peaktime energy should cover not only the running cost but also the capital cost of thermal capacity. We got the 6¢ per kwh peaktime surcharge by first calculating the annualized capital cost of thermal capacity of \$120/KW = .15 × \$800/KW). We then divided this \$120 by the 2000 hours of system peak. That calculation no longer makes sense in the presence of significant amounts of seasonal hydro capacity. Now the thermal peak is going to be significantly longer -- say 4000 hours per year, because our seasonal dams have enough capacity to deal with more than just the system peak. Hence if demand grows (at peak as well as offpeak hours), what is going to happen is that our <u>given</u> hydro storage capacity will still fill the system peak and more, but the increase of demand will leave a gap which will (under our assumptions) be filled by adding to the number of homogeneous thermal generators in the system. Thus incremental thermal capacity will operate (again under

our assumption of homogeneity) for the 4000 hours of thermal peak, not just the 2000 hours of the system or demand peak. Hence the peak that we should use for the calculation of the peaktime surcharge is 4000 rather than 2000 hours, and the resulting surcharge becomes 3¢ rather than 6¢ per kwh. The total collected to cover capital cost is exactly the same $120.00 per KW per year as before, only now it is spread over 4000 rather than 2000 hours. Why? Because this is the number of hours that newly-added homogeneous thermal plants are expected to operate.

* * * * *

Up to this point we have carried the discussion on the assumption that the price of offpeak energy was equal to its running cost (here 4¢/kwh), and that of peaktime energy was equal to running cost plus a peaktime surcharge of either 6¢ (without seasonal hydro) or 3¢ (with seasonal hydro) per kwh. These assumptions make economic sense, and can be said to carry out, in our simple example, the lessons of modern electricity economics, a branch of economic analysis set in motion by French technocrats working at or with Electricité de France in the early 1950s. It was their great insight that the true economic marginal cost of electricity would naturally vary by hour of the day, day of the week and in many cases season of the year, and it was their recommendation that these variations should be reflected in the prices paid by the users of electric energy. The French started the time-pricing of electricity in the 1950s. They were followed by many (perhaps by now even most) other countries in adopting this innovation. Typically, time-pricing is first applied to large industrial and commercial users, and only gradually and often only partially extended to domestic customers. But by now time-pricing of energy for household use is also pretty widespread -- I have been paying Southern California Edison

separately for peak and offpeak use, at different rates for different seasons of the year, ever since the middle 1990s.

But what do we do in cases where time-pricing is not used, and where the pricing system therefore does not reflect true economic cost? The answer here is very simple. The fact that the prices paid by users do not reflect the true economic cost of energy does not change that true cost. Even if peaktime energy is given away for free to some users, that does not alter the fact that it costs 10¢/kwh (without seasonal hydro) or 7¢/kwh (with seasonal hydro) in our examples.

And since our measure of the benefit of hydro projects (whether they are run-of-the-stream or daily reservoir, or seasonal storage) is based on the amount of thermal generating costs that they end up saving, all the calculations that we have done assuming "prices" equal to 4¢, 10¢, and 7¢ remain valid. But now they should be recognized as measures of system marginal costs of electricity under the relevant assumed conditions. It is by adding up the savings of these costs that a new project accomplishes, that we obtain a measure of the project's direct benefits.

* * * * *

Readers will probably already appreciate that it takes a certain amount of effort for a person to feel at ease doing electricity economics. The path we have traveled in this paper barely opens the doors to the complexity of the subject, but I feel confident that such an introduction is necessary for a person to begin to acquire an intuitive understanding of the subject. This paper will be followed by another, in which the assumptions made will be brought much closer to likely real-world cases.