Preliminary and incomplete; please do not cite

OPERATING RULES AND SEDIMENT MANAGEMENT FOR HYBRID RUN-OF-RIVER/PEAKING CAPACITY HYDROELECTRIC PLANTS

R. David Simpson[1]

rdsimpson3@live.com

This version prepared for 10 January presentation at the Delhi School of Economics

ABSTRACT

The loss of hydroelectric reservoir capacity to sediment accumulation is a serious problem in much of the world. Hydroelectric plants are typically classified as either "reservoir" or "run-of-river" (RoR) facilities. A reservoir facility stores water during the rainy season so as to be able to generate more power in the dry season. The allocation of stored water for dry-season power production has been extensively explored using the tools of dynamic optimization.

There is also a sort of hybrid system: RoR plants with small reservoirs for daily peaking. Such plants store only enough water to run their turbines for a part of the day, and allocate daily flows during the dry season to meet peak daily electricity demand. During the rest of the year, they operate as conventional RoR plants. In this work I consider the optimal allocation of daily peaking capacity, deriving a sort of peak-load pricing rule relating the value of power generated during peak and off-peak periods. I also nest the solution to the *daily*-timestep optimization problem into an *annual*-timestep optimization to characterize expenditures on sediment management. This inquiry is motivated by work I have been doing on sediment management at the Kali Gandaki A Hydroelectric Plant in Nepal, and I will illustrate findings with some results calibrated for that facility.

____

This work has its origins in a project I have been working on with the Natural Capital Project for the World Bank. My understanding of the issues I am addressing, while still far from perfect, has been improved by interactions with a number of people. I would like to thank all, without implicating any, of the following: Mahesh Acharya, Raj Bista, Benjamin Bryant, Himesh Dhungel, Arjun Kaandel, Gregory Morris, Urvashi Narain, Hara Raj Neupane, Annu Rajbhandari Rafael Schmidt, Rabin Shresthra, Govinda Timilsina, Adrian Vogl, and William Young.

---

[1] https://www.perc.org/people/r-david-simpson/

I. Introduction

Hydroelectric power is important in many parts of the world. Two variants are common. Reservoir plants impound a large volume of water behind dams to compensate for seasonal variability in river flow.[2] Run-of-river (RoR) plants divert water from a river channel to generate electricity.

Allocation of stored water from a reservoir facility has been studied in some detail (Førsund 2015; Kawashima 2007). It is treated as a problem in the allocation of an exhaustible resource: the water impounded in the reservoir during a dry season during which it is not replenished. The general principle of exhaustible resource allocation then applies: water should be allocated so its marginal value is equal at all times. The marginal value of water used in hydroelectric generation is comprised of the product of its marginal contribution to power generation and the marginal value of the power generated. The latter typically varies with the time at which power is provided; there is typically a peak demand period, and a continual variation in consumer willingness to pay for power at other times. Consequently the relatively simple rule that the marginal value of water discharged should be equal at all times when it is scarce may be translated into much more complicated operating rules.

The analysis can be further complicated when a dam operator faces both short- and long-term optimization problems. We have just described the short-term problem: how to allocate water from a fixed supply over the duration of an extended period when replenishment is limited. There is also a problem, though, in that many hydroelectric reservoirs are losing capacity over time. Reservoirs fill with sediment so, absent measures to control accumulation, capacity may become more and more scarce – and hence, valuable – over time. It is often possible to check sediment accumulation by dredging, siphoning, controlling erosion in the catchment area, or with periodic management operations such as reservoir flushing. This poses another optimization problem: how much cost should be incurred at any point in time so as to control sediment accumulation that will affect future operations? The general principle here is that the reduction measures should be pursued until the net present value of the marginal cubic meter of reservoir capacity in all *future* periods just balances the *current* marginal cost of sediment control (see Kawashima 2007)

These problems of large hydroelectric plants with reservoirs that impound water for seasonal use have been studied in the literature, and prescriptions offered for their solution.[3] Such issues would be obviated for RoR plants, which are, by definition, constrained to operate with whatever flow is available at the time. There is, however, an interesting hybrid: an RoR plant

---

[2] Impoundment might also be used to create a greater vertical drop ("hydraulic head") in a more compact space. Power generation varies proportionally with head, other things being equal.

[3] Similar issues also arise in the analysis of dams that impound water during the winter for irrigation in the summer.

with “peaking capacity”. In such a facility a dam impounds a relatively small volume of water; it is only enough to supply power for a fraction of a day. The operator of such a plant faces an even shorter-term optimization problem than does operator of a reservoir plant with seasonal storage: what daily schedule of filling and discharging the limited capacity of the reservoir will optimize power users’ benefts?

The first step in solving such a problem is to determine the conditions under which the capacity constraint binds. Peaking capacity plants may be built on rivers where flow is sufficient to operate at full power most of the year, but not during a dry season when flow is reduced. Obviously, when river flow is sufficient to serve the full capacity of the plant to produce power the capacity of the reservoir to store water is irrelevant. When does peaking capacity bind, and hence, become valuable, then? When there is a period of the day during which reservoir is full and can store no more water, and another when it is empty, and power production is constrained by river flow.

A short-term solution when capacity binds involves satisfaction of different conditions depending on whether the reservoir is filling or emptying. The marginal value of a cubic meter discharged should be constant over an interval during which the reservoir is filling (or emptying), but it will be lower (respectively, higher) in the former than the latter.

The long-term solution also differs from that of the seasonal reservoir. The value of an extra cubic meter of reservoir capacity maintained by sediment management is not the marginal value of the electricity that can be produced with the extra cubic meter of water stored, but rather, the *difference* between peak and off-peak values of the electricity that could be generate from the marginal cubic meter. The fundamental insight is that the amount of electricity that can be produced by a peaking plant depends ultimately on the flow of water in the river. Reservoir capacity does not affect this flow. Rather, it affects *when* during the course of the day the fixed flow of water can be discharged. The long-term optimization problem is solved, then, when the current marginal cost of sediment reduction is equated to the net present value of the *difference* between peak and off-peak power values.

This paper is motivated by three concerns. First, these issues arise in the management of the Kali Gandaki “A” Hydroelectric Plant in Nepal. There is a practical interest in better understanding how its peaking capacity might be best managed in short term and best preserved in the long term, and this inquiry proceeds from that interest.[4] Second, this topic involves an interesting application of dynamic optimization methods, and fills in an as-yet-unexplored niche in the analysis of hydroelectricity plant operations.

[4] It should, however, be underscored that the analysis and any opinions expressed here are the author’s, and do not necessarily reflect those of the World Bank, which sponsored the work on which this paper is based, the Nepal Electricity Authority, nor any other body.

Third and most interestingly, the application of this analysis is to a part of the world that has historically been underserved, and for which meeting electric power needs has been an important challenge in development policy (Timilsina, et al., 2018). Steps to increase the efficiency with which existing reservoir capacity is allocated and to manage the conservation of that capacity might contribute to real improvements in people's lives. Moreover, those improvements might arise both directly, from more timely power provision, but also indirectly, through other avenues. When electricity is not available via central supply, other sources can include thermal generation, with implications for both greenhouse gas emissions and local air pollutants.

The rest of the paper is presented in four sections. In the next section we derive conditions for when the reservoir capacity constraint binds. Following that, we solve the short- and long-term optimization problems. The third section following applies the results to the Kali Gandaki "A" Hydroelectric Plant, and a final section discusses results and considers how arcane prescriptions from dynamic optimization modeling might be reduced to more operational guidance.

II. When does capacity matter?

This paper is about valuing the preservation of reservoir capacity in a plant with limited daily peaking capacity during a dry season of the year. Reservoir capacity, like any other economic good, is only valuable to the extent that it is scarce. If there is more capacity than is needed to operate the plant optimally, the marginal value of capacity is zero. Stated in this way, this claim seems unexceptionable, and it is. However, it is easy to confuse two different things. Water that is held in the reservoir during periods of low demand for discharge during a later period of high demand *is* valuable. But its scarcity arises because the flow of water in the river is itself scarce during the dry season. The fact that water than can be used later is valuable does not necessarily mean that having extra capacity to store that water would itself be valuable.

This might be demonstrated formally as follows. Let $v(x_\theta, \theta)$ be the consumer surplus derived from electricity generated when water is discharged at rate x_θ at time θ. The index θ should be thought of as clock time, rather than a date. The types of facilities we are interested in are run on 24-hour cycles to meet daily peak demand, so assume that the dam operator's objective is to

$$\max_{\{x_\theta\}} \int_0^{24} v(x_\theta, \theta) d\theta. \quad (1)$$

This objective is to be maximized over repeating 24-hour daily cycles. If S_θ is the amount of water stored in the reservoir at time θ and we can assume the flow of water in the river is approximately constant at rate f at any time of day,[5] then

$$\dot{S} = f - x_\theta. \qquad (2)$$

A dot over a variable indicates its total derivative with respect to time. Expression (2) just says the change in reservoir volume at any point in time is the difference between inflow and discharge at that time.

There is, of course, also a limit as to how much can be stored in the reservoir. It cannot be drained below some zero point, nor filled in excess of its capacity, which I will denote as K,

$$0 \leq S \leq K \qquad (3)$$

The limits in (3) then imply that

$$x_\theta = f \text{ if } S_\theta = 0 \text{ or } S_\theta = K; \qquad (4)$$

if the reservoir were completely full, or if it were completely empty over any interval of time, the discharge rate would have to be the same as the inflow rate over that interval.

Finally, as the plant is operated on a 24-hour cycle, suppose that $S_0 = S_{24}$: at the end of a 24 hour period there must be as much water in the reservoir as there was at the beginning, so the cycle can be repeated again.[6]

Consider now the solution to the problem just described. It is instructive first to consider the case in which the constraints in (4) do *not* bind, and hence that the operator's choice of discharge rate is never constrained to equal river flow. Put another way, the reservoir is never completely filled or completely emptied. To solve this problem, introduce a costate variable, λ, to append (2) to the integrand of (1) (in the language of optimal control theory, forming the Hamiltonian). A solution must satisfy

$$\frac{\partial v}{\partial x} = \lambda \qquad (5)$$

and

---

[5] Precipitation during a day might affect flow, but such variations may not be large and, of course, the "dry season" is characterized by a general dearth of rain.

[6] We could also suppose the beginning and end times of the cycle are choice variables under the operator's control. This would introduce additional boundary conditions that would add to the complexity of the analysis, but repeat elements already explicated in the discussion of "switch points" below. Inasmuch as it seems both intuitively and empirically obvious that the plant would and should be operated on a 24-hour cycle to meet daily variation in demand, these endpoints will be considered fixed.

$$\dot{\lambda} = 0 \qquad (6)$$

Heuristically, the economic interpretation of the costate variable, λ, is as the implied price of the state variable, which is, in this case, the volume of water held in the reservoir, S. Equation (6) says that the value of another liter of water in the reservoir is just the marginal value of the power that could be generated by releasing that liter. Equation (7) says that, because the liter of water could be released at any time of day, the marginal value should be the same during every minute of the day. If it were not, the operator should allocate more discharges when they're more valuable and fewer when they're less valuable.

Denote by $V(K)$ the value of the objective function, (1), when the optimal sequence of discharges, $\{x_\theta\}$, is chosen. We are interested in the marginal value of capacity, K; by how much does the value of the objective, $V(K)$, increase with an incremental increase in K? Differentiating $V(K)$ we find:

$$\frac{dV}{dK} = \int_0^{24} \frac{\partial v}{\partial x}\frac{dx}{dK} d\theta = \lambda \frac{d}{dK}\int_0^{24} x d\theta = \lambda \frac{d}{dK}(24f) = 0 \quad (7)$$

The second equality comes from the maximization conditions that $\partial v/\partial x = \lambda$ and λ is constant; the third equality comes about because over a 24-hour repeating cycle the total amount of water discharged must equal the total flow available; and the final equality results because river flow is independent of reservoir capacity.

While expression (7) is, in a sense, trivial, it is worth underscoring that the mathematical argument substantiates the fundamental economic proposition. The value of an asset depends on how scarce it is. If the reservoir were never fully filled and drained in a 24-hour cycle – that is, if, as assumed in deriving (7), the constraints in (3) and their implication in (4) never arose – then there would be no economic loss associated with lost capacity.

Note that expression (7) does *not* say that discharges never vary over the daily cycle, nor does it say that there is no value in having the ability to store water at some points so as to be able to discharge more at another. If demand varies over the course of the day, discharges would certainly vary so as to better serve demand. Expression (7) says, rather, that if discharges do not vary by *enough* to invoke capacity constraints, then capacity (again, as distinguished from the ability to vary discharges) will have no marginal value.

The distinction between the value of additional river *flow* and that of reservoir *capacity* might be underscored by paralleling equation (7), but this time differentiating with respect to the river flow rate, f rather than reservoir capacity, K:

$$\frac{dV}{df} = \int_0^{24} \frac{\partial v}{\partial x}\frac{dx}{df} d\theta = \lambda \frac{d}{df}\int_0^{24} x d\theta = \lambda \frac{d}{df}(24f) = 24\lambda \quad (8)$$

By the usual interpretation of the costate variable as the shadow price of the state variable, λ is the amount by which the objective would be increased with a marginal increase in water stored. Over the course of a day during which the reservoir starts empty, is filled, and then is emptied again, total discharges must equal total inflows. So the value of water in storage is just one hour's worth of the value if flow could be increased in every hour (as the units of flow are measured as volume per hour). Again, the point to be made here is that *water* is valuable on the margin when flow is low even if the *capacity* to store it may not be scarce.

This begs the question of the conditions under which the capacity constraint binds. Suppose that there is a maximum flow a plant can safely handle. At the Kali Gandaki "A" Plant used as an example in this paper, the maximum flow of water the each generating turbine is designed to handle is 47 cubic meters per second, so the overall limit for all three of its turbines operating simultaneously would be 141 meters $m^3 \cdot s^{-1}$. Designate the maximum rate of discharge at which power can be generated as $\overline{x}$. When the rate of flow in the river, f, exceeds $\overline{x}$ there is obviously no need for storage capacity. When flow is less than $\overline{x}$ the need for storage capacity depends on the difference between actual flow, f, and the maximum that can be used, $\overline{x}$.

This can be illustrated by considering an extreme example. Recall from expression (7) above that the marginal value of capacity would be zero if the reservoir were never fully filled and emptied within the same 24-hour operating cycle. It may be instructive to ask, then, under what conditions of actual and maximum flow it would be *physically possible* to fully fill and drain the reservoir in the same 24-hour cycle. It cannot be economically optimal to do what is physically impossible.[7]

Recall that K is the capacity of the reservoir. Let us introduce one additional quantity: define x_0 as the minimum allowable flow. At the Kali Gandaki Plant a minimum flow of 4 $m^3 \cdot s^{-1}$ should be maintained to sustain aquatic life in the river. If we take as given that a plant operates on a 24-hour peaking cycle, a necessary (it is not generally sufficient) condition for reservoir capacity to bind would be that

$$K/(f - x_0) + K/(\overline{x} - f) = \frac{K(\overline{x} - x_0)}{(f - x_0)(\overline{x} - f)} \leq 24. \qquad (9)$$

The quickest possible option for filling the reservoir is to let water accumulate at the rate of $f - x_0$, releasing only the minimum amount required, x_0, while the reservoir is filling. This would take $K/(f - x_0)$ hours to accomplish. Then, when it is full, the quickest way to empty the

[7] We might note in passing that if the marginal willingness to pay for power is determined by factors exogenous to the operation of the plant, such as the cost of supply from alternative sources, a so-called "bang-bang" operating rule could be optimal. The reservoir should first be filled as rapidly as possible, then emptied as rapidly as possible, as we describe here.

reservoir would be to discharge water at rate $\overline{x}$. As more water is flowing in all the time, though, the net rate of discharge would be $\overline{x} - f$. It would, then, take $K/(\overline{x} - f)$ hours to empty at this rate. Note that discharging at a gross rate faster than $\overline{x}$ is ruled out on the argument that there would be no point in discharging water that could later be used for generation.

The *shortest possible* duration for an empty-to-full-to-empty cycle results when the denominator of the middle expression in (9), $(f - x_0)(\overline{x} - f)$ is maximized. Differentiating this expression and setting the result to zero to find a maximum,

$$(\overline{x} - f) - (f - x_0) = 0 \Longrightarrow f = (\overline{x} + x_0)/2 \quad (10)$$

Substituting from (10) into (9), if reservoir capacity comprises a binding constraint, we would need to have

$$K \leq 6(\overline{x} - x_0) \quad (11)$$

To see how this might play out in the case of the Kali Gandaki Plant, its maximum designed flow is 141 $m^3 \cdot s^{-1}$, or 507,600 $m^3 \cdot h^{-1}$. Minimum environmental flow is 4 $m^3 \cdot s^{-1}$, or 14,400 $m^3 \cdot s^{-1}$. So

$$K \leq 6 \cdot (507{,}600 - 14{,}400) = 2{,}959{,}200 \text{ m}^3. \quad (12)$$

Reservoir capacity would have to be less than about 3 million cubic meters to be binding *under any circumstances*. The rate of flow that minimizes the refill cycle length is about 72.5 $m^3 \cdot s^{-1}$. The minimum average flow in the Kali Gandaki River at the dam site is about 55 $m^3 \cdot s^{-1}$. At this flow, capacity would have to be less than about 2.77 million cubic meters to be binding.

It may be worth underscoring again the distinction between the value of *having more water* to generate electricity as opposed to *having more capacity to store water*. Intuitively, when water flow in the river is at its lowest, having more water would be most valuable. But by the same token, when water flow in the river is at its lowest, having more capacity to store water might not help, if the flow were arriving too slowly to fill the added capacity. It is also worth underscoring one important assumption underlying this analysis: that the plant is operated on a 24 hour peaking cycle. More capacity might be used if water were to be stored for use, say, only every second or third day. For present purposes, though, we assume that the plant is operated on a daily cycle: being able to provide *more* peak power on even-numbered days would not adequately compensate for providing *less* on odd-numbered ones.

It is not clear whether and when the capacity constraint binds at the Kali Gandaki Plant. While bathymetric measurements[8] give some idea of the volume of the reservoir, and concerns have

[8] Bathymetric measurement refers to mapping the bottom of the reservoir to determine its depth at various points.

arisen that its capacity is being reduced by sediment accumulation, current estimates of the reservoir's capacity are approximations. Its original design capacity of about 3.7 million m^3 has clearly been reduced, but it is not clear by how much. Correspondence with experts indicate that the plant may have been designed with excess capacity so as to allow some period of operation before its depletion became a concern.

III. Optimal operating rules in the short and long runs

In this section we return to the optimization set out in expressions (1) – (4), assuming this time that the capacity constraint *does* bind. This allows a characterization of the optimal *daily* operating rule during times when river flow is low and the plant is used for daily peaking. There is also a longer-term problem to be considered. Over a period of time as short as a single season reservoir capacity may be regarded as fixed. Over a time span of multiple years, however, an operator may be able to take actions that will affect capacity. We also consider this long-term optimization problem. As the solution to the short-term problem is required to solve the long-term problem, the former is considered first.

IIIA. Daily optimization

Return to the constrained optimization problem defined by (1) – (4) but suppose now that each constraint in (3) binds at some point in the day. That is, at some point the reservoir is empty, and at some other point, it is full. Over any interval during which the reservoir is either full or empty the rate of discharge is necessarily constrained to be equal to the rate of flow in the river, f.

Suppose for simplicity, but not unrealistically, that daily demand is a single-peaked function. Assume, then, that the operator follows a pattern such as the following:

- Starting at time 0, when the reservoir is empty, until some later time θ_1, when it is full, she discharges less water than flows in: $x_\theta < f$ for $0 \leq \theta < \theta_1$.
- Between θ_1 and some later time, θ_2, the reservoir is maintained at full capacity. Whenever the reservoir is full, discharge must equal flow in the river: $x_\theta = f$ for $\theta_1 \leq \theta < \theta_2$.
- Between θ_2 and θ_3 the operator discharges more water than flows in, meeting peak demand, and emptying the reservoir: $x_\theta > f$ for $\theta_2 \leq \theta < \theta_3$.

- Between θ_3 and 24 hours, the operator is again constrained to discharge as much water as flows in, as she cannot maintain less than zero volume in the reservoir: $x_\theta = f$ for $\theta_3 \leq \theta < 24$.

- At time 24 the cycle begins again, with discharges held to less than flow to begin to refill the reservoir.[9]

These assumptions lead to a multipart elaboration of the objective, (1):

$$V(K) = \max_{\{x_\theta\}} \left[\int_0^{\theta_1} v(x_\theta, \theta) d\theta + \int_{\theta_1}^{\theta_2} v(f, \theta) d\theta + \int_{\theta_2}^{\theta_3} v(x_\theta, \theta) d\theta + \int_{\theta_3}^{24} v(f, \theta) d\theta \right] \quad (13)$$

Now note that the intervals $[0, \theta_1)$ and $[\theta_2, \theta_3)$ are "free" in the sense that the operator can choose x's on these intervals without being constrained by the capacity of the reservoir. This means that expressions (4) and (5) above hold on these intervals, *but, importantly, for different values of the costate variable λ on the different intervals.* So

$$\frac{\partial v}{\partial x} = \lambda_1 \quad (14)$$

And

$$\dot{\lambda}_1 = 0 \quad (15)$$

on the interval $[0, \theta_1)$, while

$$\frac{\partial v}{\partial x} = \lambda_3 \quad (16)$$

and

$$\dot{\lambda}_3 = 0 \quad (17)$$

on the interval $[\theta_2, \theta_3)$.[10]

---

[9] We are abstracting from one concern that could be important. The amount of power that can be generated by a cubic meter of water varies linearly with *hydraulic head*: the vertical distance between the surface of the reservoir and the turbines below. There is, then, the possibility that what I have described as the fourth part of cycle might be dispensed with, in order to more rapidly build up head so that more power can be generate more quickly. As variation in head tends to be relatively small – on other order of 5 meters in roughly 110 – we abstract from this consideration for now, but will bring it up again later.

[10] One might imagine situations in which the operator would want to discharge massive amounts of water in order to produce a pulse of power to meet a brief, high, demand peak. In practice, such possibilities may be constrained by an upper limit on generation imposed by the capacity of the equipment.

Note that equations (14) – (17) describe a sort of "complementary slackness" condition; when the operator is free to choose the flow rate, x_θ, the corresponding marginal value of capacity, λ_θ remains constant. Heuristically, the operator should allocate flow so that there are no "arbitrage opportunities" to increase overall value on the time interval over which choice is unconstrained. Conversely, constraints on flow would arise when the marginal value of capacity is rising (or falling) and it is impossible to generate any more (or less) power, given the limits of reservoir capacity.

As in the case in which the reservoir capacity constraint did not bind, we find the marginal value of capacity by differentiating (13) with respect to K. Doing so,

$$\frac{dV}{dK} = \int_0^{\theta_1} \frac{\partial v}{\partial x}\frac{dx}{dK}d\theta + \int_{\theta_2}^{\theta_3} \frac{\partial v}{\partial x}\frac{dx}{dK}d\theta + [v(x_1,\theta_1) - v(f,\theta_1)]\frac{d\theta_1}{dK} + [v(f,\theta_2) - v(x_2,\theta_2)]\frac{d\theta_2}{dK} + [v(x_3,\theta_3) - v(f,\theta_3)]\frac{d\theta_3}{dK} \quad (18)$$

Note that we are condensing notation, avoiding subscripting subscripts by writing x_i for x_{θ_i}.

Consider the two integrals on the right-hand side of the equal sign in (18) first. The choice of x is only free when the reservoir is neither full nor empty. From (14) – (17), optimization over intervals in which the choice of discharge is not constrained implies that the marginal value of discharges is constant over such intervals. So

$$\int_0^{\theta_1} \frac{\partial v}{\partial x}\frac{dx}{dK}d\theta = \lambda_1 \int_0^{\theta_1} \frac{dx}{dK}d\theta = \lambda_1 \frac{d}{dK}\int_0^{\theta_1} x_\theta d\theta \quad (19)$$

Time θ_1 is defined implicitly as the duration required to fill the reservoir, starting from time zero. So, by definition,

$$\int_0^{\theta_1} (f - x_\theta)d\theta = K \quad (20)$$

Differentiating (20) with respect to K, and rearranging

$$(f - x_1)\frac{d\theta_1}{dK} - 1 = \int_0^{\theta_1} \frac{dx}{dK}d\theta \quad (21)$$

Using (21) in (19),

$$\int_{0}^{\theta_1} \frac{\partial v}{\partial x}\frac{dx}{dK}d\theta = \lambda_1\left[(f - x_1)\frac{d\theta_1}{dK} - 1\right] \quad (22)$$

A similar set of machinations, invoking the implicit definition of θ_2 and θ_3 by

$$\int_{\theta_2}^{\theta_3} (x_\theta - f)d\theta = K \quad (23)$$

gives

$$\int_{\theta_2}^{\theta_3} \frac{\partial v}{\partial x}\frac{dx}{dK}d\theta = \lambda_3\left[1 + (f - x_3)\frac{d\theta_3}{dK} - (f - x_2)\frac{d\theta_2}{dK}\right] \quad (24)$$

Using (22) and (24) in (18)

$$\begin{aligned} \frac{dV}{dK} &= \lambda_3 - \lambda_1 \\ &+[v(x_1, \theta_1) + \lambda_1 \cdot (f - x_1) - v(f, \theta_1)]\frac{d\theta_1}{dK} \\ &+[v(f, \theta_2) - v(x_2, \theta_2) - \lambda_3 \cdot (f - x_2)]\frac{d\theta_2}{dK} \\ &+[v(x_3, \theta_3) + \lambda_3 \cdot (f - x_3) - v(f, \theta_3)]\frac{d\theta_3}{dK} \end{aligned} \quad (25)$$

Expression (25) can be greatly condensed, as each of the quantities in square brackets must be zero. The reason for this conclusion can be found in texts on dynamic optimization (see, e. g., Kamien and Schwartz 1982), but the heuristic argument is straightforward. Each of the terms in square brackets compares the instantaneous value of discharging a cubic meter of water immediately before and immediately after the capacity constraint starts to bind. Consider, for example, the quantity in the first set of square brackets. Its first term, $v(x_1, \theta_1)$, is the instantaneous contribution to the operator's objective at time θ_1. The second term, $\lambda_1 \cdot (f - x_1)$, is the implicit value of the additional water stored in the reservoir at the last moment before it reaches full capacity. Balanced against the sum of these two terms is the third, $v(f, \theta_1)$, the value realized at time θ_1 when discharges are constrained to the flow rate, f, by having reached full capacity (note that we could also add $\lambda_1 \cdot (f - f)$, to make the difference we are considering symmetric, but of course, net additions must be zero when the reservoir is full). Since the operator has the ability to choose a strategy yielding a different "switch time,"

θ_1, if she has chosen that switch time optimally, a small variation in it should not yield a higher overall value.[11]

Thus we have

$$\frac{dV}{dK} = \lambda_3 - \lambda_1. \quad (26)$$

While the mathematics underlying the derivation of (26) may have been tedious, the intuition underlying the result is straightforward. If reservoir capacity constrains the operator's choices, it is not because it forces her to produce *less* power over the course of a day than she would have liked to. Her ability to generate power is constrained, but the source of the constraint is the flow of the river, not the capacity of the reservoir. What the capacity of the reservoir constrains is the operator's choice of *when* she can produce the power the river's flow rate allows. A marginal cubic meter of storage means that the operator can produce the corresponding amount of additional power when its value is high (λ_3), but by choosing to produce more power when it is most valuable, she necessarily forgoes the option of using that cubic meter of water to produce the same amount of power when its value is lower (λ_1).

We might note finally that other constraints might come into play in the actual determination of optimal discharges. It may be, for example, that the plant would optimally be operated at full power to meet peak demand (see also footnote 3 above). For present purposes, however, we will simply note that expression (26) establishes our main point: the marginal value of capacity is the *difference* between consumers' willingness to pay for power between peak and off-peak periods.

IIIB. Annual optimization

Expression (26) shows that the marginal value of capacity during the dry season is given by the difference between the marginal value of power during peak and off-peak periods. In deriving (26) we have assumed that reservoir capacity is fixed over periods as short as days or seasons. In deriving conditions to characterize the optimal allocation of reservoir capacity over the course of a day, $V(K)$ has been defined as the value of the planner's (daily) objective function when she is constrained by reservoir capacity K.

Over longer periods – years – reservoir capacity varies as sediment accumulates and measures are taken to control that accumulation. Recognizing this, we now ask how those sediment control measures might be conducted so as to maximize the net present value of the plant. Using the results derived for daily management of the reservoir, consider next optimal

[11] Note also that the terms in square brackets in (25) also provide formulae that could be used in conjunction with (14) – (17), (18), and (20) to calculate the optimal allocation of discharges.

management over longer time horizons. In the analysis above the volume of water, S, held in the reservoir was defined as the state variable, and we asked how discharges, x, should be scheduled so as to maximize value over the course of a day. Over a longer time horizon suppose that the capacity of the reservoir, K, is a state variable whose magnitude can be affected by controlling the deposition of sediment, which will now be denoted by σ. We will now subscript these variables with an index t to denote the passage of time in years: K_t and σ_t.

Define

$$W(K_t) = D_t \cdot V(K_t) - c(\sigma_t) + (1-\delta)W(K_t - \sigma_t), \quad (27)$$

where $W(K_t)$ is the net present value the operator's objective starting at time t, $c(\sigma_t)$ is the cost of restricting sediment deposition at time *t* to be no greater than σ_t, and δ is the discount rate. The other as-yet undefined variable, D_t, is the number of days in a year that river flow is low enough that capacity, K_t, constrains operations. More generally, one might suppose that flow in the river which, recall, we have assumed to be constant *within* a day varies *between* days over the course of a season. In the interest of brevity, however, suppose for now that D_t is fixed. So, expression (27) is just a standard intertemporal valuation formulation: the net present value of operations with capacity K_t at time t is the value of operation over the D_t days of the dry season, less the cost of sediment control, plus the next present value of future operations commencing one year in the future. The value of all plant operations would also include the value of power generated when reservoir capacity does not constrain operations, but as such values would, by construction, not depend on reservoir capacity, we ignore them here.

The operator decides how much sediment deposition to allow in year *t*, σ_t. The optimal choice will be characterized by

$$-c'(\sigma_t) - (1-\delta)W'(K_t - \sigma_t) = 0; \quad (28)$$

Differentiating both sides of (27) with respect to K_t,

$$W'(K_t) = D_t \cdot V'(K_t) + (1-\delta)W'(K_t - \sigma_t) \quad (29)$$

or

$$\delta W'(K_t - \sigma_t) - [W'(K_t - \sigma_t) - W'(K_t)] = D_t \cdot V'(K_t) \quad (30)$$

Finding a general solution to (30) is difficult. A couple of extreme cases are illustrative, though. First, suppose the right-hand side of (30) were zero. This would be the case if reservoir capacity did not yet constrain daily operations during the dry season.[12] Then

[12] Plants may be constructed with larger reservoirs than are initially required, on the assumption that the accumulation of sediment will reduce capacity over time.

$$\delta = \frac{W'(K_t - \sigma_t) - W'(K_t)}{W'(K_t - \sigma_t)}; \quad (31)$$

The proportionate rate of change in net present value should grow at the discount rate. Though accumulation of sediment might not yet constrain dry-season operations, it hastens the date at which such a constraint will bind, and hence, imposes a cost.

Expression (31) might hold for a recently built facility that had not yet experienced much sediment deposition. At the other extreme, if it were possible to operate a facility in a steady state in which the value is maintained over time and the net rate of sediment deposition were zero, we would have

$$\delta W'(\overline{K}) = D_t \cdot V'(\overline{K}), \quad (32)$$

Or, substituting from (28)

$$-c'(0) = \frac{1-\delta}{\delta} D_t \cdot V'(\overline{K}) \quad (33)$$

The bar over K designates a steady-date value, and by making the argument of the marginal cost function zero we are supposing that net sediment deposition is zero.

Finally, using the short-term results derived earlier,

$$-c'(0) = \frac{1-\delta}{\delta} D_t \cdot (\lambda_3 - \lambda_1) \quad (34)$$

The expression on the left-hand side of (34) is the *current*, *single-period* marginal cost of controlling the marginal cubic meter of sediment deposition when reservoir volume is held constant and no net accumulation occurs. The expression on the right-hand side of (34) is the *net present value over all future periods* of an extra cubic meter of reservoir storage space.[13] As noted in the introduction, the problem of allocating peaking capacity on a daily cycle differs in some important ways from that of allocating fixed reservoir capacity on a seasonal basis. However, the result that *current* marginal costs should be balanced against the net present value of all future marginal benefits is common to the literature on reservoir capacity management more generally (see, e. g., Kawashima 2007).

[13] Slightly more formally, noting the factor $1 - \delta$, the expression on the right-hand side of (29) is the net present value *starting one period in the future* of the marginal cubic meter of reservoir capacity. This is an artefact of the specification of discounting in this model.

IV. Calibrating results

Expression (34) relates the marginal benefit of sediment reduction, on the right-hand side, to the marginal cost, on the left. To see how this balance might be struck in practice, consider again the example of the Kali Gandaki "A" Hydroelectric plant in Nepal.

The three turbines at the Kali Gandaki Plant are intended to operate at a maximum water flow of 141 cubic meters per second ($m^3 \cdot s^{-1}$). Water flow in the Kali Gandaki can vary from as little as 55 $m^3 \cdot s^{-1}$ in the dry season of winter and early spring to more than 1,000 $m^3 \cdot s^{-1}$ during and after the annual Monsoon, when flows are also enhanced by snow and glacial melting at higher elevations. In the dry season, then, water flow in the river is insufficient to support peak power generation.

The right-hand side of equation (34) depends on the discount rate, the number of days that capacity constrains operations, and the difference in values between peak and off-peak power consumption. Let us suppose the discount rate is 8%, so $(1-\delta)/\delta = 11.5$. From the annual water flow statistics, flow is insufficient to allow full power generation for about six months of the year, so suppose that D_t, the number of days for which the capacity constraint binds, is 180.

It has often been difficult to assign a reasonable value to the marginal willingness to pay for power in Nepal, a country in which supply has often been insufficient to meet the demand at the prices the Nepal Electricity Authority charges in peak periods. Recently, however, an increase in power imports and purchases from independent producers has greatly increased supply (NEA 2018). Moreover, construction of a new hydroelectric facility with more than three times the Kali Gandaki "A" Plant's capacity, the Upper Tamakoshi Plant, will expand the country's domestic production capability by about two-thirds.

It seems reasonable to suppose, then, that the gap between consumers' willingness to pay for peak and off-peak power values will narrow, as more sources become available to supply peak power. For the purposes of illustration, suppose that the official tariffs, which historically have not measured willingness to pay (hereinafter MWP) for peak power generation, will measure MWP going forward. Suppose, then, that the value of off-peak power – the price earlier designated as λ_1 – is NR 6 per kWh, and peak power – λ_3 – is NR 12 per kWh.

Finally, we must translate from the value of the additional power that can be produced at the peak to the value of the marginal cubic meter of reservoir storage that will enable its production. Each generating turbine at the Kali Gandaki plant has a rated capacity of 48 MW at a flow rate of 47 $m^3 \cdot s^{-1}$. An additional cubic meter of water storage could, then, run a turbine for 1/47th of a second. Since the turbine could produce 48 MWh of electricity in an hour, which is 3600 seconds, the marginal cubic meter of capacity would enable generation of 48 MWh/(60 · 60 · 47) = 0.284 kWh of additional peak-load power.

So, filling in these values for the variables on the right-hand side of (34), we have

$$11.5 \cdot 180 \cdot (12 - 6) \cdot 0.284 = 3{,}525 \quad (35)$$

This figure of about Nepali Rupees 3,500 per cubic meter of sediment deposited is the steady-state value of marginal cubic meter of sediment occupying live storage capacity in the Kali Gandaki "A" Plant's reservoir

The reservoir at the plant was designed with a storage capacity of some 7.7 million m^3, although nearly all of the 4 million m^3 of dead storage was filled almost immediately with sediment (Morris 2014). The plant receives a tremendous volume of silt, sand, and larger particles from glacial runoff, the stream bed, land erosion, and other sources such as landslides. It is estimated that the annual flow of sediment, about 43 million tonnes per year, could completely fill the reservoir in a single rainy season (sediment transport varies nonlinearly with water flow, and is negligible when flows are low). Less than one-tenth of one percent of this flow is trapped in the reservoir, however (IHA 2014). The remainder passes on. The great majority of it is simply transported around the hydroelectric generating equipment during the months when water flow is high and most water is diverted around the intakes. About 15 percent of the sediment enters the intakes, where most is trapped in settling basins, and the remainder – around four percent of the annual load – passes through the generating turbines, where it can do considerable damage, reducing operating efficiency, raising the prospect of unanticipated outages, and necessitating costly repairs.

A variety of practices may be adopted to prevent sediment from settling in the reservoir or dislodging that which has accumulated. Sediment-laden water may be sluiced or flushed. Sluicing is simply diverting it around the hydroelectric intakes, as described above. Flushing is periodically opening all gates and emptying the reservoir, a measure that might, if undertaken for long enough, scour the basin down to the approximate contours of the river's original channel. Some such operations might be costless. When water flow is extremely high it is dangerous to operate the generating equipment due to the high concentration of sediment.[14] Suspending generation could, then, have the beneficial side-effect of also scouring and removing more sediment.

Experts have also suggested that the plant be maintained at a lower level when flows are high and, hence, sediment-laden (Morris 2014). The maximum operating height of the Kali Gandaki reservoir is 524 meters above mean sea level; the minimum, 518. The lower is the level, other

[14] Not only the amount of sediment, but also the composition of materials mobilized, differ with higher flow rates. Greater flows can mobilize heavier sediment, and larger sand, as opposed to finer silt, particles cause more damage.

things being equal, the shorter the residence time of water in the reservoir and, consequently, the less time for sediment to settle.

There is, however, a cost to maintaining a lower operating height. The amount of power a hydroelectric plant produces from a cubic meter of water depends linearly on the height from which it falls – what is referred to as its "hydraulic head". The design head for the Kali Gandaki plant is about 110 meters, meaning that every meter by which the reservoir height is lowered reduces power generation by a little less than one percent. If we suppose that *an element* – it need not be entirely sufficient in and of itself– of the strategy to control sediment accumulation in the reservoir were to maintain a lower operating level, we can estimate the marginal cost of doing so.

The opportunity cost of forgone generation would be higher during peak than off-peak demand periods, so let us suppose that reductions in operating height could be timed to occur in off-peak hours.[15] The rated capacity of the Kali Gandaki plant is 144 MW, so it could generate 144 MWh of electricity in 60 minutes. If the reservoir height were lowered by one meter from an initial level of 110 meters, the resultant loss in power generation would be 144 MWh/110 = 1,310 kWh. At an off-peak price of NR 6 per kWh, this works out to NR 7,860 per hour of off-peak operation.

Suggested operating procedures for the Kali Gandaki Plant call for sluicing operations when flow exceeds twice that required to achieve full power generation – a little less than 300 $m^3 \cdot s^{-1}$ (IHA 2014). Average flow exceeds this threshold for about four months of the year, from June through October: about 120 days. The cost of operating at a lower level for 16 hours per day during 120 days per year would then be NR 7,860 · 16 · 120 = NR 15.09 million.

This is the seasonal cost, in terms of the value of power generation forgone, from lowering the height of the reservoir by one meter during off-peak hours during the rainy season. We have not found information on how effective this would be; that is, would lowering reservoir height by a meter during periods of high flow result in enough reduction in sediment accumulation to justify the cost of lost power generation? In the absence of this information, what can be done instead is to calculate

[15] Sediment transport and deposition is only a significant issue during the Monsoon season and few months after. Far more sediment is suspended when flow is high than when the river is reduced to a relative trickle. Average daily flow often reaches 1,000 $m^3 \cdot s^{-1}$ during the summer months. Only 141 $m^3 \cdot s^{-1}$ need be used to achieve full power generation. The surface area of the reservoir is about 65 hectares – 650,000 m^2. With water arriving at the rate of, say 850 $m^3 \cdot s^{-1}$, the height of the reservoir could be raised at a rate of 850 $m^3 \cdot s^{-1} \cdot$ 3,600 s per hour/650,000 m^2 = 4.7 meters per hour. In short, the height of the reservoir could be varied reasonably quickly between peak and off-peak periods when flow is high, so it should be feasible to lower the level when power is less valuable and raise it again power is more valuable.

$$\frac{NR\ 15{,}090{,}000\ \cdot m^{-1}}{NR\ 3{,}527\ \cdot m^{-3}} = 4{,}280\ m^3 \cdot m^{-1};^{16} \quad (36)$$

that is, if maintaining the reservoir at a lower height were part of a cost-effective strategy for managing reservoir capacity to maximize the value of service in a steady state, a one-meter reduction in reservoir height would need to result in maintenance of incremental volume sufficient to store about 4,280 cubic meters of water. It has been estimated that about 43 millions tonnes of sediment flows down the Kali Gandaki River every year, and of this, less than 0.1%, or 43,000 tonnes, which would displace about 30,000 cubic meters of water, is retained. The figure of 4,280 cubic meters of water, which would translate into about 6,000 tonnes of sediment, is considerably lower than the presumed current rate of deposition. Conversations with experts suggest that capacity loss in the reservoir has, in fact, recently been reversed by more scrupulous adherence to recommendations to reduce the level of water in the reservoir during the rainy season.

At the time the Kali Gandaki "A" Plant began operations in 2002 live storage capacity was estimated to be about 3.7 million cubic meters. If, over 15 years, that volume was reduced to the point that absence of storage capacity began to constrain operations – to less than 3 million cubic meters – the annual rate of loss would have to have been on the order of 50,000 $m^3 \cdot y^{-1}$. The fact that experts to credit better operating practices with preserving capacity suggests that maintaining a lower reservoir height during high-flow periods is, in fact, a cost-effective strategy, and perhaps even that more aggressive implementation of that strategy, and/or adoption of other measures for sediment control and removal would be justified.

V. Discussion and conclusion

The above results suggest that policies now in effect for management at the Kali Gandaki "A" Plant are justified. What constitutes an optimal policy may vary considerably with underlying variable and parameter values, however. The marginal cost of sediment reduction prescribed in equation (34) varies in the duration of the period during which river flow is low. The discount rate calculation $(1-\delta)/\delta$ would take on a value of 19 at a discount rate of 5%, but decline to 7 for a discount rate of 12.5%. Another source of considerable uncertainty might be the duration of the season during which the reservoir capacity constraint binds. Water flow is typically low from roughly November through May, but it would be difficult to say exactly when the constraint binds. Moreover, the shadow price of the constraint (the difference between the λ's) would vary with both the capacity of the reservoir and the flow in the river.

[16] The units here, of course, could simply be expressed as "meters squared," but it is useful to underscore that we are asking how much *volume* could be preserved by sacrificing a unit of *height*.

The most critical consideration may be the difference between peak and off-peak marginal willingness to pay for power, $\lambda_3 - \lambda_1$. It is, perhaps, a heroic assumption to suppose that Nepal's power problems have been resolved to the extent there is only a six Rupee difference between marginal willingness to pay between peak and off-peak periods. Only a few years ago some commentators estimated a peak MWP of 40 Rupees or more (Shrestha and Shrestha 2016).

It is worth underscoring, though, that it is the *difference* between peak and off-peak MWP, not the absolute magnitude of the former, that determines the value of maintaining reservoir capacity. It is the flow of the river, not necessarily the capacity of the reservoir, that determines how much power can be generated. Capacity only affects *when* power is generated. If it didn't matter when power were generated – if it were equally valuable at any time – reservoir capacity would have little value.

A related issue is what determines MWP for power. If the Kali Gandaki Plant were Nepal's main source of power generation any increase (or decrease) in the amount of power it provided at a particular time would induce a corresponding decrease (or increase) in consumers' MWP at that time. Recently, however, the Nepal Electricity Authority has relied on power purchased from India or independent domestic providers to meet the nation's demand (NEA 2018). These foreign and independent providers are effectively setting the peak MWP: what consumers are willing to pay is what they will have to pay to the sources that provide it. In the future, and as facilities like Nepal's large new Tamakoshi Plant come online, the peak-load pricing principles of public utility regulation may be invoked. The implication of these observations is that a facility like the Kali Gandaki "A" Plant may in the future be treated even more as a peaking facility than it now is.[17] This may both ease the task of valuing reservoir capacity and obviate some of the finer points of analysis in Section IIIA. When the operator's choices do not affect willingness to pay for power provided from the plant, the operating procedure generally involves what is known in dynamic optimization as a "bang-bang" solution (see also footnote 6): operate at full capacity during the peak period, and reduce discharges to the minimum so as to recharge the reservoir as quickly as possible when demand is lowest.

This last observation begs the question of if, perhaps, the analysis presented in this paper is "too clever by half". It would be a daunting task to prescribe detailed operating rules even if extensive and reliable data were available; they are not, so practical advice needs to be more limited. Moreover, an outside researcher would be displaying great hubris if he were to claim to have learned more from a brief study than professionals who have been intimately engaged in practical operations for years already know.

The conclusions that arise from this inquiry are not so much that anything should be done differently, then, than an affirmation that what is being done seems to make sense. Reservoir

---

[17] In fact, the Nepal Electricity Authority's most recent *Annual Report* underscores that the availability of power from other sources had led to more peak power generation dedication at Kali Gandaki (NEA 2018).

capacity at the Kali Gandaki Plant declined over its first decade and a half of operation. This may well have been because the plant was constructed with excess capacity, in anticipation that some filling would occur. On observing that this has occurred, however, efforts have been made to assure that sediment deposition is controlled by, among other things, sacrificing some power production by reducing the height of the reservoir in high-flow periods.

We can derive some approximate figures for the opportunity cost of maintaining the lower reservoir level, and, subject to the proviso that the estimates both of these costs and the benefits of maintaining capacity are subject to considerable imprecision, the costs seem to suggest that, if anything, still more aggressive control measures might be adopted. Thesemight include occasional reservoir flushing and reduction of erosion in the catchment due to road building and other causes.[18]

Finally, inasmuch as the value of reservoir capacity is determined by the difference between MWP for peak and off-peak power, continued improvement in the supply system might eventually reduce the importance of having reservoir capacity to generate extra peak power. By the same token, however, uncertainties concerning the degree to which expansions in system supply will continue to keep pace with growing demand, it will likely to be prudent to continue to maintain reservoir capacity in case peak demands again spike.

[18] There is little that could realistically be done to dramatically reduce sediment delivery. The headwaters of the Kali Gandaki lie high in the Himalayas, and between glaciers and active seismology, a tremendous volume of material will be transported regardless of land use closer to the reservoir. Local disturbances may, however, contribute larger particles – stones, rather than silt – and so measures like road-building restrictions might be relatively effective on the margin.

References (incomplete)

Førsund, F. R. (2015). *Hydropower Economics*. Springer.

Kamien, M. I. and N. L. Schwartz (1982). *Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management*. Elsevier.

Kawashima, S. (2007). Conserving reservoir water storage: An economic appraisal. *Water Resources Research*, *43*(5). https://doi.org/10.1029/2006WR005090

Nepal Electricity Authority (NEA) (2018). *Annual Report*.

Shrestha, J., and N. T. Shrestha (2016). Expansion Planning of Electricity Generating System Using the VALORAGUA and WASP-IV Models in Nepal. *Hydro Nepal Journal of Water Energy and Environment*.

Timilsina, G., P. Sapkota, and J. Steinbuks (2018). How Much Has Nepal Lostin the Last Decade Due to Load Shedding? An Economic Assessment Using a CGE Model. World Bank Development Research Group Working Paper.