

Long-Term Hydro Scheduling in a Pool Market

P. J. Neame, A. B. Philpott, G. Pritchard

University of Auckland

New Zealand.

email: p.neame@auckland.ac.nz

Abstract

A generator offering power into a pool market like New Zealand's needs to construct an offer stack for each half-hour period, indicating how much power they will generate at each price. For a hydroelectric generator, decisions in each period are linked by the quantity of water in the dam. We consider a simple model where the generator is unable to influence the market price by withholding capacity, and hence we model the uncertainty with a price distribution function. We have devised a dynamic programming algorithm that can then find the optimal stack for each half-hour period. However, this is only practical for short-term models. Hence, we consider a model with weekly periods, initially by offering the same stack in for each half-hour period within that week. Because the price is uncertain, the revenue gained is uncertain, and so is the release (the total quantity of water used). We model the release with a normal distribution, and consider the combined problem of selecting a mean and a variance release for each week, and constructing stacks with this mean and variance. This leads to some elegant theoretical results and gives some preliminary insights into constructing a practical algorithm.

1 Introduction

In many electricity markets, some or all of a generator's production will be sold on the spot market. Usually, in each period (half an hour in New Zealand), the generator is required to construct an *offer stack* made up of several tranches specifying a price range at which a given quantity (specified by the generator) of electricity will be provided. From the combination of all generators' offer stacks, and the demand for electricity, a clearing price is found at each node of the network, and generators are paid the price at their node. There is substantial uncertainty as to what this price will be, which also affects how much electricity a generator will be producing in each period.

For a hydroelectric generator, the major cost is an opportunity cost of using the water in their dam now, as opposed to saving it for later use. This links the price

they would like to receive now, to what the prices may be in the future. Even after offering a stack to the market, there is uncertainty in the amount of production the market will assign to the generator and hence how much water it will need to use.

Most of the previous work on hydro scheduling in the literature has focussed on the centrally-planned case where the uncertainty is in the demand; see for example Pereira and Pinto [4] and Jacobs et al. [2]. One notable exception is the work of Scott and Read [5], although their focus was more on constructing equilibria than optimal offers.

If we make the assumption that the generator is not large enough to influence the market price, then the price uncertainty may be modelled by a price distribution function. In [3], Neame et al. considered how to produce a near-optimal offer stack for a single period in this case. In this paper, we discuss how to set up a dynamic programming model with a stage for each half-hour period, and hence produce near-optimal stacks for a short-term model. However, some of the price activity we would like to consider is very obviously seasonal. For example, a New Zealand hydro generator planning their releases in March of 2001 should have considered the prospect of a dry, cold winter ahead.

Hence we need to consider some way to construct a model in which each stage represents a longer time period, for example one week. The information provided by this model could then be supplied to a short-term model. In this paper, we begin by considering the approach of offering the same stack in every period of the week. We can approximate the amount of water used over the week by a normal distribution. Thus, for every stack there is a mean and variance of dispatch. This can also be related to a distribution in return on the water used. If the value-of-water function is nonlinear, then the value-to-go depends on both the mean and variance of dispatch.

This leads to an optimisation model, where the aim is to select a mean and variance, along with a stack which produces that mean and variance. This model has a number of nice properties, allowing us to make a number of observations about how to find optimal stacks, and about the form of the value-to-go function.

2 Solving the Short-Term Model

In [3], we develop a dynamic programming algorithm to construct a near-optimal stack for a given time-period with a given cost function. If the price is represented by a probability density function $f(p)$, the one period problem is

$$\max_{q(p)} \int [pq(p) - c(q(p))] f(p) dp, \quad (1)$$

subject to $q(p)$ being non-decreasing and piecewise constant. (Throughout this paper, we omit the limits on the integrals; they are minimum and maximum prices.) In the hydro case with no inflows, assuming that switching costs and running costs are small, the cost function can be written as $c(q) = V(x_0) - V(x_0 - W(q))$, where x_0 is the current water level, $V(x)$ is the value-of-water function, and $W(q)$ gives the units of water required to produce q units of electricity.

The algorithm partitions the quantity axis into a finite number of potential production levels and produces a stack that gives an expected return that is close to the optimal expected return. Details on choosing the partition, and bounds on sub-optimality can be found in [3].

This approach can then be extended to a multi-period case by approximating the value-of-water function $V_t(x)$ at each time period t , and choosing a $V_T(x)$ at the time horizon. This was coded and tested on a model based on the Karapiro dam, which has three turbines and hence $W(q)$ is nonsmooth at two points (where two turbines become more efficient than one, and where three become more efficient than two). An approximation of the inverse of $W(q)$ is shown in Figure 1. We observed shoulder period prices in New Zealand (prior to this year's dramatic price behaviour) and constructed a price distribution $f(p)$ based on four normal distributions (although raw prices could be fed into the code instead). We found that any boundary effects caused by the choice of $V_T(x)$ disappearing after around 20 periods, and using a Pentium II 400, a model planning 24 periods into the future (with 100 partition points on the quantity axis at each period, and 220 water levels) could be solved in just over 5 minutes. Substantial speed-ups can most likely be achieved by using fewer partitions and water levels, but even then, planning for more than a week into the future would take an excessive amount of computing time.




Figure 1: The efficiency curve of a three turbine hydroelectric generator

With $V_T(x) = 0$ for all x and no inflow, an initial value-of-water function $V_0(x)$ (as shown in Figure 2) is produced. Note that the graph is flat beyond the largest throughput possible in the timeframe, and the function is concave (near the maximum throughput, the generator will run at almost any price; with little water, they can wait in the hope of higher prices later). This is demonstrated by the stacks generated in period 0, shown for various initial water values in Figure 3.

3 A Longer Term Model

The dynamic programming algorithm presented previously will take an unacceptable amount of time to solve for more than a few days into the future. To consider




Figure 2: The value of water function

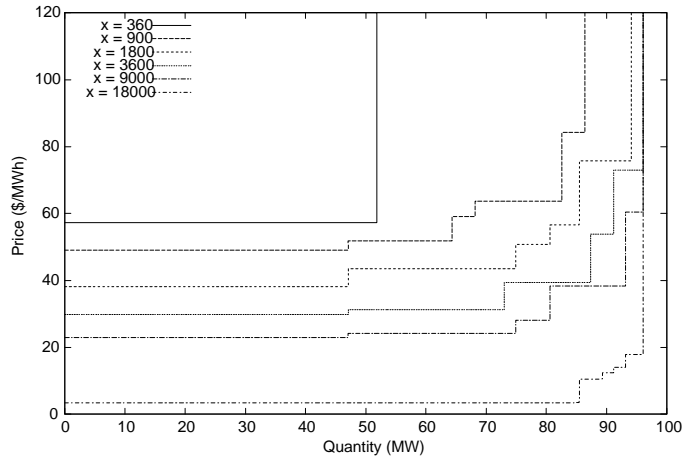


Figure 3: Offer stacks generated for various water levels x

seasonal effects, we would like to consider possible price and inflow behaviour many months into the future. To do this, we coalesce half-hour periods into weekly periods; however, given that the amount of water used and the average price obtained for electricity within that week are both uncertain, this is not a straightforward task. We begin by treating each half-hour period as identical and saying that we will offer the same stack in each half hour.

Throughout this section we assume for simplicity that $W(q) = q$. (The general case will be discussed in the paper's conclusions).

By the central limit theorem, the same offer in each half-hour will generate a total dispatch distribution that is approximately normal with mean μ_W and variance σ_W^2 . If we label the mean and standard deviation of water release in each half-hour period as μ and σ , then

$$\begin{aligned}\mu_W &= N\mu \\ \sigma_W &= \sqrt{N}\sigma.\end{aligned}$$

If we have N identical shorter periods per longer period e.g. if we offer in every half hour for a week, against the same $f(p)$, then $N = 336$. We then try to find a stack for each half-hour period giving the best expected total return for the week. The maximal return possible for given μ and σ is

$$g(\mu, \sigma) = \max_{q(p)} N \int pq(p) f(p) dp$$

$$\text{subject to } \int q(p) f(p) dp = \mu \quad (2)$$

$$\int (q(p))^2 f(p) dp = \sigma^2 + \mu^2.$$

In what follows, we assume that at period t we have a value-of-water function $V_{t+1}(x)$ which is concave in x . This will then give us a value-to-go function

$$h(x, \mu, \sigma) = E_Z \left[V_{t+1} \left(x - N\mu - \sqrt{N}\sigma Z \right) \right].$$

Then, for a given water level x , the total return R is given by

$$R = \max_{\mu, \sigma} g(\mu, \sigma) + h(x, \mu, \sigma),$$

giving a combined problem

$$R = \max_{\mu, \sigma} \left\{ \max_{q(p)} N \int pq(p) f(p) dp + h(x, \mu, \sigma) \right\}$$

$$\text{subject to } \int q(p) f(p) dp = \mu \quad (3)$$

$$\int (q(p))^2 f(p) dp = \sigma^2 + \mu^2.$$

The chief difficulty in solving this problem arises from having the square of the mean in the second equality constraint. This would include products of different tranche quantities, meaning that the stack optimisation could not be solved by the DP given in [3]. Fortunately, this problem has some nice properties.

Lemma 3.1 *If $V_{t+1}(x)$ is concave, then the optimal value of (3) can be obtained by solving*

$$R = \max_{\mu, \sigma} \left\{ \max_{q(p)} N \int pq(p) f(p) dp + h(x, \mu, \sigma) \right\}$$

$$\text{subject to } \int q(p) f(p) dp = \mu \quad (4)$$

$$\int (q(p))^2 f(p) dp \leq \sigma^2 + \mu^2.$$

Proof. Consider arbitrary x, μ and Z and $\sigma_1 \leq \sigma_2$. Then since $V_{t+1}(x)$ is concave, we have

$$V_{t+1} \left(x - N\mu - \sqrt{N}\sigma_1 Z \right) + V_{t+1} \left(x - N\mu + \sqrt{N}\sigma_1 Z \right)$$

$$\geq V_{t+1} \left(x - N\mu - \sqrt{N}\sigma_2 Z \right) + V_{t+1} \left(x - N\mu + \sqrt{N}\sigma_2 Z \right).$$

As the normal distribution is symmetric it follows that

$$E_Z \left[V_{t+1} \left(x - N\mu - \sqrt{N}\sigma_1 Z \right) \right] \geq E_Z \left[V_{t+1} \left(x - N\mu - \sqrt{N}\sigma_2 Z \right) \right],$$

and hence $h(x, \mu, \sigma_1) \geq h(x, \mu, \sigma_2)$. Thus there exists an optimal solution to (4) with the second constraint satisfied as an equation at optimality, and hence the optimal solution value is identical to that of (3). \blacksquare

Lemma 3.1 is important, as (4) has a convex feasible region. Now we show that $g(\cdot)$ and $h(\cdot)$ are concave functions, so long as the set of feasible stacks is convex and the value-of-water function is concave.

Lemma 3.2 *If $V_{t+1}(x)$ is concave, then $h(x, \mu, \sigma)$ is concave in x, μ and σ .*

Proof. By definition

$$h(x, \mu, \sigma) = E_Z \left[V_{t+1} \left(x - N\mu - \sqrt{N}\sigma Z \right) \right],$$

where Z is taken from the standard normal distribution. For fixed Z , $V_{t+1} \left(x - N\mu - \sqrt{N}\sigma Z \right)$ is a linear function of x , μ , and σ and so $V_{t+1} \left(x - N\mu - \sqrt{N}\sigma Z \right)$ is a concave function of x , μ , and σ . The expectation of concave functions is concave, so $h(x, \mu, \sigma)$ is concave in x, μ and σ . ■

Lemma 3.3 *If the set of feasible stacks is convex, then $g(\mu, \sigma)$ is concave in μ .*

Proof. Define a norm $\|\cdot\|$ on the space of feasible stacks by

$$\|q(\cdot)\| = \sqrt{\int (q(p))^2 f(p) dp}.$$

Then (2) can be written as

$$\begin{aligned} g(\mu, \sigma) = & \max_{q(p)} \int pq(p) f(p) dp \\ & \text{subject to } \int q(p) f(p) dp = \mu \\ & \|q(p) - \mu\| \leq \sigma. \end{aligned}$$

We fix σ at an arbitrary value $\hat{\sigma}$. For arbitrary μ_1 and μ_2 , let the stacks giving the maxima in (2) for $g(\mu_1, \hat{\sigma})$ and $g(\mu_2, \hat{\sigma})$ be $q_1(p)$ and $q_2(p)$ respectively. Then, for some $\lambda \in (0, 1)$, consider

$$\begin{aligned} g(\lambda\mu_1 + (1-\lambda)\mu_2, \hat{\sigma}) = & \max_{q(p)} \int pq(p) f(p) dp \\ & \text{subject to } \int q(p) f(p) dp = \lambda\mu_1 + (1-\lambda)\mu_2 \\ & \|q(p) - (\lambda\mu_1 + (1-\lambda)\mu_2)\| \leq \hat{\sigma}. \end{aligned}$$

We show that the feasible stack $\bar{q}(p) = \lambda q_1(p) + (1-\lambda)q_2(p)$ satisfies each of these constraints. Firstly

$$\int \bar{q}(p) f(p) dp = \lambda\mu_1 + (1-\lambda)\mu_2.$$

Secondly,

$$\begin{aligned} \|\bar{q}(p) - (\lambda\mu_1 + (1-\lambda)\mu_2)\| &= \|\lambda q_1(p) + (1-\lambda)q_2(p) - \lambda\mu_1 + (1-\lambda)\mu_2\| \\ &= \|\lambda(q_1(p) - \mu_1) + (1-\lambda)(q_2(p) - \mu_2)\| \\ &\leq \lambda\|q_1(p) - \mu_1\| + (1-\lambda)\|q_2(p) - \mu_2\| \\ &\leq \hat{\sigma}. \end{aligned}$$

Thus $\bar{q}(p)$ is feasible to the maximisation, so

$$\begin{aligned} g(\lambda\mu_1 + (1-\lambda)\mu_2, \hat{\sigma}) &\geq \int p\bar{q}(p) f(p) dp \\ &= \lambda g(\mu_1) + (1-\lambda)g(\mu_2), \end{aligned}$$

and thus $g(\mu, \sigma)$ is concave in μ . ■

Lemma 3.4 *If the set of feasible stacks is convex, then $g(\mu, \sigma)$ is concave in σ .*

Proof. We fix μ at an arbitrary value $\hat{\mu}$. For arbitrary σ_1 and σ_2 , let the stacks giving the maxima in (2) for $g(\hat{\mu}, \sigma_1)$ and $g(\hat{\mu}, \sigma_2)$ be $q_1(p)$ and $q_2(p)$ respectively. Then, for some $\lambda \in (0, 1)$, consider

$$g(\hat{\mu}, \lambda\sigma_1 + (1 - \lambda)\sigma_2) = \begin{aligned} & \max_{q(p)} \int pq(p) f(p) dp \\ & \text{subject to } \int q(p) f(p) dp = \hat{\mu} \\ & \|q(p) - \hat{\mu}\| \leq \lambda\sigma_1 + (1 - \lambda)\sigma_2. \end{aligned}$$

We show that $\bar{q}(p) = \lambda q_1 + (1 - \lambda)q_2$ satisfies each of these constraints. Firstly

$$\begin{aligned} \int \bar{q}(p) f(p) dp &= \int (\lambda q_1(p) + (1 - \lambda)q_2(p)) f(p) dp \\ &= \hat{\mu}, \end{aligned}$$

and secondly

$$\begin{aligned} \|\bar{q}(p) - \hat{\mu}\| &= \|\lambda q_1(p) + (1 - \lambda)q_2(p) - \hat{\mu}\| \\ &= \|\lambda(q_1(p) - \hat{\mu}) + (1 - \lambda)(q_2(p) - \hat{\mu})\| \\ &\leq \lambda \|q_1(p) - \hat{\mu}\| + (1 - \lambda) \|q_2(p) - \hat{\mu}\| \\ &\leq \lambda\sigma_1 + (1 - \lambda)\sigma_2. \end{aligned}$$

Hence, as before $\bar{q}(p)$ is feasible to the maximisation, so

$$\begin{aligned} g(\hat{\mu}, \lambda\sigma_1 + (1 - \lambda)\sigma_2) &\geq \int p\bar{q}(p) f(p) dp \\ &= \lambda g(\sigma_1) + (1 - \lambda)g(\sigma_2), \end{aligned}$$

and thus $g(\mu, \sigma)$ is concave in σ . ■

For a long term planning model, we seek to maximise at each stage $g(\mu, \sigma) + h(x, \mu, \sigma)$ for each reservoir level x . For practical implementation, it is highly desirable that this objective function is concave at each stage. The following theorem shows that concavity is inherited from V_{t+1} to V_t , so in theory at least, a dynamic programming algorithm iterating back from a time horizon will be tractable.

Theorem 3.5 *If the set of feasible stacks is convex and $V_{t+1}(x)$ is concave, then $V_t(x)$ is concave.*

Proof. From Lemmas 3.2, 3.3 and 3.4, $g(\mu, \sigma) + h(x, \mu, \sigma)$ is concave. Recall that

$$V_t(x) = \max_{\mu, \sigma} \{g(\mu, \sigma) + h(x, \mu, \sigma)\}.$$

Let μ_1^* and σ_1^* be optimal in this maximisation at $x = x_1$ and μ_2^* and σ_2^* be optimal at $x = x_2$. For some $\lambda \in (0, 1)$, let $\bar{x} = \lambda x_1 + (1 - \lambda)x_2$, $\bar{\mu} = \lambda\mu_1^* + (1 - \lambda)\mu_2^*$ and

$\bar{\sigma} = \lambda\sigma_1^* + (1 - \lambda)\sigma_2^*$. This gives a feasible μ, σ pair, so

$$\begin{aligned}
V_t(\lambda x_1 + (1 - \lambda)x_2) &= \max_{\mu, \sigma} \{g(\mu, \sigma) + h(\bar{x}, \mu, \sigma)\} \\
&\geq g(\bar{\mu}, \bar{\sigma}) + h(\bar{x}, \bar{\mu}, \bar{\sigma}) \\
&= g(\lambda\mu_1^* + (1 - \lambda)\mu_2^*, \lambda\sigma_1^* + (1 - \lambda)\sigma_2^*) \\
&\quad + h(\lambda x_1 + (1 - \lambda)x_2, \lambda\mu_1^* + (1 - \lambda)\mu_2^*, \lambda\sigma_1^* + (1 - \lambda)\sigma_2^*) \\
&\geq \lambda g(\mu_1^*, \sigma_1^*) + (1 - \lambda)g(\mu_2^*, \sigma_2^*) + \lambda h(x_1, \mu_1^*, \sigma_1^*) \\
&\quad + (1 - \lambda)h(x_2, \mu_2^*, \sigma_2^*), \text{ from previous lemmas} \\
&= \lambda V_t(x_1) + (1 - \lambda)V_t(x_2).
\end{aligned}$$

Hence $V_t(\cdot)$ is concave. ■

The Lagrangian for (2) is

$$\begin{aligned}
L(\mu, \sigma, q(p), \theta, \phi) &= N \int pq(p) f(p) dp + \theta \left(\mu - \int q(p) f(p) dp \right) \\
&\quad + \phi \left(\sigma^2 + \mu^2 - \int q^2(p) f(p) dp \right).
\end{aligned}$$

To maximise L we seek a stack $q(p)$ which maximises

$$\int [Npq(p) - \theta q(p) - \phi (q(p))^2] f(p) dp. \tag{5}$$

This is equivalent to solving (1), with $c(q) = \frac{\theta q(p) + \phi (q(p))^2}{N}$.

Suppose that the solution to this problem with $\theta = \theta^*$ and $\phi = \phi^*$ yields a stack giving a distribution of water release with mean μ^* and standard deviation σ^* . Then, by the Lagrangian Sufficiency Theorem (see Whittle [6]), this stack yields the optimal solution in (2), and hence gives the value of $g(\mu^*, \sigma^*)$.

In particular, if we have no restrictions on the stack we can offer, then the optimal solution for $q(p)$ in (5) is a straight line, since the marginal cost is given by $c'(q) = \frac{\theta}{N} + \frac{2\phi}{N}q$ and the optimal solution to offer at marginal cost. Setting $p = c'(q)$, we find

$$q(p) = \frac{Np - \theta}{2\phi}.$$

Suppose that the price distribution has mean \bar{p} and standard deviation s . We seek θ^*, ϕ^* giving a stack satisfying the constraints of (2), so we consider

$$\begin{aligned}
\mu &= \int q(p) f(p) dp \\
&= \frac{N}{2\phi^*} \bar{p} - \frac{\theta^*}{2\phi^*}
\end{aligned}$$

and

$$\begin{aligned}
\mu^2 + \sigma^2 &= \int (q(p))^2 f(p) dp \\
&= \frac{1}{4(\phi^*)^2} [N^2 (\bar{p}^2 + s^2) - 2N\theta^* \bar{p} + (\theta^*)^2]
\end{aligned}$$

Solving these equations gives us

$$\phi^* = \frac{Ns}{2\sigma}, \theta^* = N \left(\bar{p} - \frac{\mu s}{\sigma} \right),$$

and hence the equation for the optimal stack is

$$q^*(p) = \mu + \frac{\sigma}{s} (p - \bar{p}).$$

Observe that the quantity offered is perfectly correlated with the price, and we choose the intercept and slope of the stack to give the desired mean and variance. This gives us an expected return of

$$\begin{aligned} L(\mu, \sigma, q^*(p), \theta^*, \phi^*) &= N \int p q^*(p) f(p) dp \\ &= N(\sigma s + \bar{p}\mu), \end{aligned}$$

and hence

$$g(\mu, \sigma) = N(\sigma s + \bar{p}\mu).$$

This particular form for the return arises because, in the case where there are no bounds on q , the quantity dispatched is perfectly correlated with the price.

However, if we use these results with an example $f(p)$, then the optimal mean μ and standard deviation σ gives σs much greater than $\bar{p}\mu$. Effectively, our generator is buying large amounts of water at low prices in order to sell at higher prices. Obviously a lower bound of zero is needed on q . Also, an upper bound reflecting the turbine capacity is needed. The optimal stack will still be a straight line, but the return calculation becomes substantially more complicated.

4 Conclusions and Further Work

We have developed a model with weekly periods with the same stack offered in at each half-hour period within the week. Using a normal distribution to represent the total amount of water used in each week leaves us with three decisions to make for each week — the mean, the variance and a stack with these properties. Fortunately, the functions that arise have nice theoretical properties, concavity in the value-of-water function is inherited, and a dynamic program iterating backwards from a time horizon would be tractable. However, further research is needed to find an efficient algorithm. The dynamic programming algorithm from [3] could easily be used to find a near-optimal stack for given values of θ and ϕ , but it is not obvious how to find good values of θ and ϕ .

Obviously, the assumption that the generator cannot affect the market clearing price is not true for some of the larger generators in New Zealand. Anderson and Philpott [1] address the issue of optimality of stacks in a price-making environment, which we are moving towards including into the framework presented in this paper. Throughout Section 3 we assumed that $W(q)$, the amount of water used by the generator to generate q units of electricity was given by $W(q) = q$. We conjecture

that if $W(q)$ is convex then $g(\cdot)$ will be concave in the mean and standard deviation of electricity generation.

Some further extensions include allowing for contracts (which make no difference to the offering strategy in the price-taking case) and modelling each week as being made up of more than one type of period (e.g. to allow three stacks for peak, shoulder and off-peak periods). The price model could be further improved by considering dependences within a week and between weeks. Finally, we would also need to incorporate inflows, allowing for uncertainty in the inflows and correlations in inflows across weeks.

Acknowledgement

The authors acknowledge the support of FRST under contract UOAX0004.

References

- [1] E. J. Anderson, A. B. Philpott, 2001, "Optimal Offer Construction in Electricity Markets", *Mathematics of Operations Research* (to appear).
- [2] J. Jacobs, G. Freeman, J. Grygier, D. Morton, G. Schultz, K. Staschus, J. Stedinger, "SOCRATES: A System for Scheduling Hydroelectric Generation Under Uncertainty", *Annals of Operations Research* 59(1995) pp. 99–133.
- [3] P. J. Neame, A. B. Philpott, G. Pritchard, "Offer Stack Optimisation in Electricity Pool Markets", submitted to *Operations Research*.
- [4] M. V. F. Pereira, L. M. V. G. Pinto, "Multi-stage Stochastic Optimization Applied to Energy Planning", *Mathematical Programming* 52 (1991) pp. 359–375.
- [5] T. J. Scott, E. G. Read, "Modelling Hydro Reservoir Operation in a Deregulated Electricity Market", *International Transactions in Operations Research* 3 (1996) pp. 243–253.
- [6] P. Whittle, *Optimization Under Constraints*, Wiley (1971).