Chapter 31

On the Marginal Value of Water for Hydroelectricity

Andy Philpott[21]

31.1 ▪ Introduction

This chapter discusses optimization models for computing prices in perfectly competitive wholesale electricity markets that are dominated by hydroelectric generation. We revisit the relationship between partial equilibrium in perfect competition and system optimization of a social planning problem, and we use the latter to show how perfectly competitive electricity prices correspond to marginal water values.

Most industrialized regions of the world have over the last 20 years established wholesale electricity markets that take the form of an auction that matches supply and demand. The exact form of these auction mechanisms varies by jurisdiction, but they typically require offers of energy from suppliers at costs at which they are willing to supply, and clear a market by dispatching these offers in order of increasing cost. Day-ahead markets, such as those implemented in most North American regions, seek to arrange supply well in advance of its demand so that thermal units can be prepared in time. Since the demand cannot be predicted with absolute certainty, these day-ahead markets must be augmented with balancing markets to deal with the variation in load and generator availability in real time.

In this chapter, we study markets with hydroelectric generators that use reservoirs of stored water. Generators with hydroelectric reservoirs face an inventory problem. They would like to optimize the release of water from reservoirs to maximize profits using a stochastic process of prices, but this process is not known and must be deduced by each agent using current and future market conditions and hydrological models of future reservoir inflows. For an agent controlling releases from a hydroelectric reservoir, the marginal cost of supply in the current period involves some modeling of opportunity cost that includes possible high prices in future states of the world with low inflows.

[21]The research described in this paper has benefited from discussions with Ziming Guan, Tony Downward, Geoffrey Pritchard, and Golbon Zakeri.

In this chapter, we restrict our attention to a setting in which all agents are price-takers who do not act strategically. Under this assumption, it is well known that electricity prices for a single trading period and a single location can be computed as shadow prices from a deterministic convex economic dispatch (ED) model that maximizes total social welfare. In a stochastic setting, where there is a stochastic process of inflows that is known to all agents, and they maximize expected profit using the probability law determining these inflows, a competitive (partial) equilibrium will correspond to the water-release policy that maximizes the expected welfare of a social planner. A stochastic optimization (SO) for maximizing expected social welfare can therefore be used to estimate the marginal cost of electricity supply in future states of the world. These values are estimates of the wholesale electricity prices in these states. This approach is demonstrated in Section 31.2 using a simple model.

In a system with a single storage reservoir, electricity prices can be shown to be equivalent to the expected marginal value of water in the reservoir, in other words the extra social benefit that would be obtained from the optimal deployment of an extra cubic meter of stored water. Observe that this value comes from the optimal marginal use of the water and so the planner must have an optimization problem in mind. Policy iteration algorithms based on an early paper by Stage and Larsson [1713] are often used in practice to compute these marginal values. In Section 31.3, we apply these methods to a simple example.

Most hydroelectricity systems comprise several reservoir storages, and so the social planner's optimization problem becomes a stochastic dynamic programming (DP) problem with a matching number of state variables. When this number exceeds the acceptable state dimension for classical DP (about three or four), the optimization method used most in practice is stochastic dual dynamic programming (SDDP) due to Pereira and Pinto [1458]. In Section 31.4, we outline this method, and we apply it to an example problem from the New Zealand electricity system in Section 31.5. The estimated prices from a case study in 2008 are then compared with their historical counterparts.

31.2 ▪ The Social Optimum

We consider a social planning problem in an electricity system of hydroelectric generating plants fed by reservoirs ($i \in \mathscr{H}$) that are subject to uncertain inflows over some decision horizon of T periods. The system also has generating plants ($j \in \mathscr{T}$) that run on thermal fuel at some known generation cost, and consumer segments $c \in \mathscr{C}$. For simplicity we shall assume that all generating plants sell power into an unconstrained transmission grid so that consumer demand can be aggregated at a single location to which power is sent.

In general, the reservoirs are connected in a set of river chains represented by a network of m nodes (reservoirs and junctions) and l arcs (canals or river reaches). The topology of the network can be represented by the $m \times l$ incidence matrix A, where

$$a_{ik} = \begin{cases} 1 & \text{if node } i \text{ is the tail of arc } k, \\ -1 & \text{if node } i \text{ is the head of arc } k, \\ 0, & \text{otherwise.} \end{cases}$$

By adding dummy nodes if necessary, we can ensure that every pair of nodes is joined by at most one arc. If there are several river chains, then the network need not be connected.

We first describe a deterministic model over the planning horizon $t = 1, 2, \ldots, T$. We let $u(t)$ be a vector of flow rates (cubic meters per period) in the arcs in the network in period t. Some arcs correspond to generating stations. Some might correspond to spill around a station. For each arc $k = 1, 2, \ldots, l$, we specify a generation function $G_k(u_k)$ that is zero when the arc does not contain a station, and otherwise is a strictly concave function of flow rate u_k for arcs corresponding to stations, giving the energy output in a period when the flow rate is u_k. Let $x(t)$ denote a vector of reservoir storages in each node at the beginning of period t and $\omega(t)$ be a vector of uncontrolled reservoir inflows (in cubic meters) that have occurred in period t. These define a water balance constraint

$$x(t+1) = x(t) - Au(t) + \omega(t),\ t = 1, 2, \ldots, T.$$

The hydro generators are accompanied by thermal generators $j \in \mathcal{T}$ that produce a vector v of energy that costs $\sum_{j \in \mathcal{T}} C_j(v_j)$ to produce, where C_j is a strictly convex function. Demand of consumer segment $c \in \mathcal{C}$ is represented by d_c, where the strictly concave function $D_c(d_c)$ denotes the welfare accrued by consumer segment c. Note that the derivative D_c' represents the inverse demand curve for consumer c, which is strictly decreasing by assumption. Demand must be satisfied in each time period, so, in a deterministic model,

$$\sum_{k=1}^{l} G_k(u_k(t)) + \sum_{j \in \mathcal{T}} v_j(t) \geq \sum_{c \in \mathcal{C}} d_c(t),\ t = 1, 2, \ldots, T.$$

When inflows to the reservoirs are random variables, the sequence of time intervals $1, 2, \ldots, T$ is replaced by a stochastic process. We represent this by a scenario tree, as shown in Figure 31.1. Here each scenario tree node n spans a period t (which in our models is a week) and corresponds to a realization $\omega(n)$ of reservoir inflows in that period. To account for different inflow sequences, all variables previously indexed by t are now indexed by n. Each node has probability $\phi(n)$, where the probabilities of nodes in a given period sum to one. We denote by $n-$ the unique parent node of node n and by $n+$ its set of children.

The leaf nodes $\mathcal{L}$ of the tree correspond to the end of the decision horizon (period T). For each leaf node $n \in \mathcal{L}$, we specify a strictly convex future-value function $V_n(x(n))$ that measures the total future welfare in state n of having reservoir levels $x(n)$.

We are now in a position to formulate a stochastic social planning (SSP) problem. This is

$$\text{SSP: min} \quad \sum_{n \in \mathcal{N}} \phi(n) \left(\sum_{j \in \mathcal{T}} C_j(v_j(n)) - \sum_{c \in \mathcal{C}} D_c(d_c(n)) - \sum_{n \in \mathcal{L}} \phi(n) V(x(n)) \right)$$

$$\text{subject to (s.t.)} \quad \sum_{k=1}^{l} G_k(u_k(n)) + \sum_{j \in \mathcal{T}} v_j(n) \geq \sum_{c \in \mathcal{C}} d_c(n), \qquad n \in \mathcal{N},$$

$$x(n) = x(n-) - Au(n) + \omega(n), \qquad n \in \mathcal{N},$$

$$d(n) \geq 0,\ v(n) \geq 0,\ x(n) \in \mathcal{X},\ u(n) \in \mathcal{U},\ n \in \mathcal{N}.$$

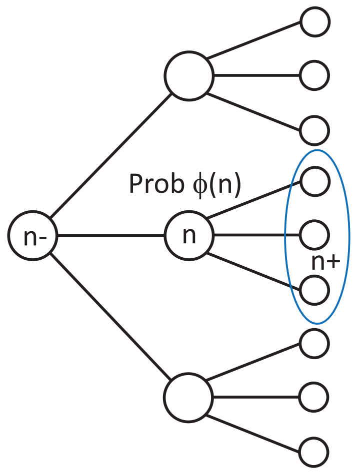

Figure 31.1. *Each node n spans a period and corresponds to a realization $\Omega(n)$ of reservoir inflows in that period.*

Here we assume that $\mathscr{X}$ and $\mathscr{U}$ are compact convex sets and SSP has a nonempty feasible region. We shall assume in what follows that all nonlinear constraints satisfy constraint qualifications to guarantee the existence of Lagrange multipliers. In particular, we have Lagrange multipliers $\phi(n)p(n)$ for each $n \in \mathscr{N}$ so that we can solve SSP by minimizing the Lagrangian

$$\begin{aligned}
\text{LSSP:}\quad \min \quad & \sum_{n\in\mathscr{N}} \phi(n)\Bigg(\sum_{c\in\mathscr{C}} p(n)d_c(n) - \sum_{c\in\mathscr{C}} D_c(d_c(n)) \\
& \quad + \sum_{j\in\mathscr{T}} \big(C_j(v_j(n)) - p(n)v_j(n)\big) - V(x(n)) - \sum_{k=1}^{l} p(n)G_k(u_k(n))\Bigg) \\
\text{s.t.} \quad & x(n) = x(n-) - Au(n) + \omega(n), \qquad n \in \mathscr{N}, \\
& d(n) \geq 0,\ v(n) \geq 0,\ x(n) \in \mathscr{X},\ u(n) \in \mathscr{U},\ n \in \mathscr{N}.
\end{aligned}$$

LSSP separates by agent. The consumer solves

$$\begin{aligned}
\text{CP}(c)\text{: max} \quad & \sum_{n\in\mathscr{N}} \phi(n)(D_c(d_c(n)) - p(n)d_c(n)) \\
\text{s.t.} \quad & d_c(n) \geq 0.
\end{aligned}$$

Each thermal generator solves

$$\begin{aligned}
\text{TP}(j)\text{: max} \quad & \sum_{n\in\mathscr{N}} \phi(n)(p(n)v_j(n) - C_j(v_j(n))) \\
\text{s.t.} \quad & v_j(n) \geq 0,
\end{aligned}$$

and the hydro generators together solve

$$\begin{aligned}
\text{HP: max} \quad & \sum_{n\in\mathscr{N}} \phi(n) \sum_{k=1}^{l} p(n)G_k(u_k(n)) + \sum_{n\in\mathscr{L}} \phi(n)V(x(n)) \\
\text{s.t.} \quad & x(n) = x(n-) - Au(n) + \omega(n), && n \in \mathscr{N}, \\
& x(n) \in \mathscr{X},\ u(n) \in \mathscr{U}, && n \in \mathscr{N}.
\end{aligned}$$

If all the river chains are operated by a single agent, then HP represents the problem of maximizing this agent's expected revenue and residual value given prices $p(n)$ in each state of the world. If there is more than one hydro agent and HP is not separable by agent, then we require an additional condition to establish a correspondence between competitive equilibrium and an optimal social plan. Let $\mathscr{H}$ denote the set of hydro agents and x_i the vector of water stocks for reservoirs operated by agent i. We require the existence of functions $V_i(x_i)$ so that

$$V(x(n)) = \sum_{i \in \mathscr{H}} V_i(x_i(n)).$$

Even with this condition, to enable a separation into agent problems, in general we require a price $\phi(n)\pi(n)$ for the flow conservation constraint. This gives a Lagrangian for HP, which is

$$\begin{aligned}\mathscr{L}(x,u,\pi) = &\sum_{n \in \mathscr{N}} \phi(n) \sum_{k=1}^{l} p(n) G_k(u_k(n)) + \sum_{n \in \mathscr{L}} \phi(n) \sum_{i \in \mathscr{H}} V_i(x_i(n)) \\ &+ \sum_{n \in \mathscr{N}} \phi(n)\pi(n)^\top (x(n-) - Au(n) + \omega(n) - x(n)),\end{aligned}$$

that we maximize over $x(n) \in \mathscr{X}$, $u(n) \in \mathscr{U}$. $\mathscr{L}(x,u,\pi)$ can be rearranged to give

$$\begin{aligned}\mathscr{L}(x,u,\pi) = &\sum_{n \in \mathscr{N}} \phi(n) \left(\sum_{k=1}^{l} p(n) G_k(u_k(n)) - \pi(n)^\top Au(n) \right) \\ &+ \sum_{n \in \mathscr{N} \setminus \mathscr{L}} \phi(n) x(n) \left(\frac{\sum_{m \in n+} \phi(m)\pi(m)}{\phi(n)} - \pi(n) \right) \\ &+ \sum_{n \in \mathscr{L}} \phi(n) \sum_{i \in \mathscr{H}} (V_i(x_i(n)) - \pi_i(n) x_i(n)) \\ &+ \sum_{n \in \mathscr{N}} \phi(n)\pi(n)^\top \omega(n).\end{aligned}$$

The values of $\pi(n)$ are water prices, often called *marginal water values*. At the end of the planning horizon, if $x_i(n)$ is not at a bound, then we require for a maximum that

$$V_i'(x_i(n)) = \pi_i(n), \quad n \in \mathscr{L}.$$

For flow $u_k(n)$ through station arc k from h to i, we have

$$-\pi(n)^\top Au(n) = (\pi_i(n) - \pi_h(n))\, u_k(n),$$

so the owner of station k attempts in state of the world n to maximize

$$\sum_{k=1}^{l} p(n) G_k(u_k(n)) + (\pi_i(n) - \pi_h(n))\, u_k(n).$$

In other words, he/she maximize revenue from $u_k(n)$ at price $p(n)$, while paying $\pi_h(n)$ for upstream water and getting paid $\pi_i(n)$ for the water supplied to the downstream reservoir. At the optimal solution u_k^*, we have

$$\pi_h(n) - \pi_i(n) = p(n) G_k'\left(u_k^*(n)\right), \tag{31.1}$$

which relates the spot price of energy $p(n)$ to the marginal water value difference and the marginal conversion factor from water to energy.

A reservoir owner in state of the world n makes $x(n)$ as large as possible when the expected marginal water value at the next time period exceeds that at the current period and makes $x(n)$ as small as possible when the expected marginal water value at the next time period is lower than that at the current period. These prices in equilibrium enable different owners on the river to extract the maximum welfare from the water. If there is no possibility of trading in water at market prices, then the market is incomplete, and the competitive equilibrium may not be the same as the optimal social plan. This observation is well known and is studied by Lino et al. [1198].

Henceforth we will assume that all the reservoirs and generating stations on a given river chain are owned and operated by the same generator, and that all the river chains have separate catchments. This means that HP separates into independent river chains each operated by a single agent. Hydro agent i solves

$$\begin{aligned} \mathrm{HP}(i)\text{: max} \quad & \sum_{n\in\mathcal{N}} \phi(n) \sum_{k=1}^{l(i)} p(n) G_k\left(u_k(n)\right) + \sum_{n\in\mathcal{L}} \phi(n) V_i(x(n)) \\ \text{s.t.} \quad & x(n) = x(n-) - Au(n) + \omega(n), \qquad\qquad n \in \mathcal{N}, \\ & x(n), u(n) \geq 0, \end{aligned}$$

where now $x(n)$ $(\omega(n))$ denotes the vector of reservoir storages (inflows) on river chain i and $u(n) \in \mathbb{R}^{l(i)}$ the vector of flows through the $l(i)$ arcs in the river chain.

This defines a perfectly competitive equilibrium defined by the individual optimality conditions and market clearing condition:

$$\begin{aligned} \text{CE:} \quad & u^i(n), x^i(n) \in \arg\max \mathrm{HP}(i), \\ & v_j(n) \in \arg\max \mathrm{TP}(j), \\ & d_c(n) \in \arg\max \mathrm{CP}(c), \\ & 0 \leq \sum_{i\in\mathcal{H}} \sum_{k=1}^{l(i)} G_k\left(u_k^i(n)\right) + \sum_{j\in\mathcal{T}} v_j(n) - \sum_{c\in\mathcal{C}} d_c(n) \perp p(n) \geq 0. \end{aligned}$$

By comparing optimality conditions, we have the following result.

Proposition 31.1. *Suppose all the reservoirs and generating stations on any given river chain are owned and operated by the same generator and*

$$V(x) = \sum_{i\in\mathcal{H}} V_i(x_i).$$

The (unique) welfare-maximizing solution to SSP is the same as the (unique) competitive equilibrium solution to CE.

Each model HP(i) can be solved by computing an optimal water release policy to maximize the expected revenue at energy prices $p(n)$. The optimal policy will generate Lagrange multipliers $\pi(n)$ for the water conservation constraints. These will be related to the energy prices through (31.1). When each hydro agent controls a single reservoir, then it is tempting to suppose that its optimal water release policy can be computed by price decomposition and DP. Unfortunately, this is not always straightforward because the energy prices $p(n)$ are not guaranteed to be stagewise independent even if the inflows are independent [166]. On the other hand, DP can be applied to the system optimization problem as long as the number of reservoirs (corresponding to states) is not too large. Given a system optimal solution to the social planning problem, it is possible to derive the competitive equilibrium prices and agent policies using the above proposition.

In the next section, we study the computation of competitive equilibrium prices for the simplest possible system, with one hydro agent and one reservoir. We then show how models with several reservoirs might be solved using SDDP [1458, 1671]. We conclude with some results of estimating competitive prices in the New Zealand wholesale electricity market.

31.3 ▪ A Single-Reservoir Model

One of the first models for reservoir water valuation was developed by Stage and Larsson [1713]. It seeks to compute marginal water values $w_i(t)$ corresponding to discrete reservoir levels $i = 0, 1, \dots, N$ using a form of policy iteration. When the reservoir is full ($i = N$), the marginal value of water is set to zero (here we assume that surplus water can be spilled with no penalty). When the reservoir is empty ($i = 0$), the marginal value of water is set to some high value that we denote as C. Typically this will be a shortage cost, also known as the *value of lost load*. Note that this convention implies that the marginal value at level i is the directional derivative

$$w_i = W_i - W_{i+1},$$

where W_i is the average expected cost of meeting demand in each future period if the current reservoir level is i. If W_i is linearly interpolated, then we obtain a nonsmooth convex function. This implies

$$W_{i-1} - W_i = w_{i-1} \geq w_i = W_i - W_{i+1},$$

with strict inequality occurring as the rule.

The expected marginal value of water at level $i = j$, say, is determined in [1713] by constructing sample paths of reservoir levels that would be produced by an optimal release policy (say to meet a known demand) under random sequences of inflows. Starting at level $i = j$ at time of year t, s out of S such sample paths will reach $i = 0$ before $i = N$. The expected marginal water value $w_j(t)$ is then estimated to be $\frac{s}{S}C$. This process is repeated for all starting levels and times.

Although the Stage and Larsson method is quite general, we will examine it here as applied to a specific problem instance. We assume that reservoir inflows are stagewise independent and have the same distribution at each time of the year. Thus we can dispense with the time index in the formulation and seek values w_i. To simplify the analysis further we will assume that the reservoir has inflow in each period that equals two with probability p and zero with probability $(1-p)$. Demand for energy requires that we release exactly one unit of water in each period. Thus in each time step the reservoir level is a simple random walk, increasing by one unit with probability p and decreasing by one unit with probability $1-p$. Observe that there is no policy choice since the water usage is prescribed by demand. In this setting, the Stage and Larsson method amounts to iterating toward the solution of a fixed-point problem.

In other words, we obtain the following procedure:

$$
\begin{array}{llll}
\text{Set} & \nu = 0, & w_i^\nu = 0; & \\
\text{repeat} & & & \\
\quad \text{for} & i = 0 \text{ to } N \text{ do} & & \\
 & \text{if } i = 0, & \text{set } w_i^{\nu+1} = & C; \\
 & \text{if } i = N, & \text{set } w_i^{\nu+1} = & 0; \\
 & \text{if } 0 < i < N, & \text{set } w_i^{\nu+1} = & (1-p)w_{i-1}^\nu + p w_{i+1}^\nu; \\
\quad \text{end} & \text{for}; & & \\
\quad \text{set} & \nu := \nu + 1 & & \\
\text{until} & w^{\nu+1} = w^\nu. & &
\end{array}
$$

It is easy to see that this procedure converges to the fixed point defined by

$$w_i = (1-p)w_{i-1} + p w_{i+1}, \quad i = 1, 2, \ldots, N-1, \quad w_0 = C, \quad w_N = 0. \tag{31.2}$$

In fact this simple model can be solved exactly for w because there is a closed-form solution for the probability that any random walk starting at $i = j$ will hit barrier $i = 0$ before it hits barrier $i = N$ (see, e.g., [854]). When $p = 0.5$, this gives $w_j = C\frac{N-j}{N}$. Otherwise, letting $r = \frac{1-p}{p}$, we get

$$w_j = \begin{cases} C\frac{1-r}{1-r^N}\left(r^j + r^{j+1} + \cdots + r^{N-1}\right), & j < N, \\ 0, & j = N. \end{cases}$$

The solution corresponds to the fixed-point calculation

$$
\begin{aligned}
&\eta + W_0 = C + (1-p)(W_0) + p W_2, \\
&\eta + W_i = (1-p)W_{i-1} + p W_{i+1}, \qquad i = 1, 2, \ldots, N, \\
&W_{N+1} = W_N,
\end{aligned}
\tag{31.3}
$$

for an average cost per period Markov decision process, where η is the long-run average cost per period of meeting demand. Given a solution W to (31.3), it is easy to see that setting $w_i = W_i - W_{i+1}$ solves (31.2).

An alternative model uses discounting. Suppose costs are discounted in each period with rate α, giving a discount factor $\gamma = \frac{1}{1+\alpha}$. This gives a slightly different fixed-point problem:

$$w_i = (1-p)\gamma w_{i-1} + p\gamma w_{i+1}, \quad i = 1, 2, \ldots, N-1, \quad w_0 = C, \quad w_N = 0. \tag{31.4}$$

The discounted solution corresponds to the fixed-point calculation

$$
\begin{aligned}
&W_0 = C + \gamma((1-p)W_0 + p W_2), \\
&W_i = \gamma\left((1-p)W_{i-1} + p W_{i+1}\right), \quad i = 1, 2, \ldots, N, \\
&W_{N+1} = W_N.
\end{aligned}
\tag{31.5}
$$

Here W_j is the expected discounted cost of meeting demand starting from reservoir level $i = j$. Given a solution W to (31.5), setting $w_i = W_i - W_{i+1}$ solves (31.4).

In both of these models there is an implicit assumption on how shortage costs are incurred. When the reservoir level hits zero, it is assumed that the cost C is incurred immediately. This assumes that none of the inflow occurring over the next period can be used to cover demand. In other words, this is a *here-and-now* or *decision-hazard* problem, which is modelled using (31.3) or (31.5).

In contrast, one might use a *wait-and-see* or *hazard-decision* assumption. Given a reservoir level $i = j$, the random inflow is observed and used with reservoir water to meet demand if possible. Inflow that cannot be stored is spilled. Thus, when the reservoir is empty, with probability $1-p$ we must meet demand with no inflow, incurring a cost of C, resulting in an empty reservoir, or we meet demand with an inflow of two units, leaving the reservoir with one unit in it. When the reservoir is full and inflow is two, we can use half of it and leave the reservoir full. If the inflow is zero, then the reservoir moves to state $N-1$. The fixed-point problem (31.3) is now

$$\begin{aligned} &\eta + S_0 = (1-p)(C + S_0) + pS_2, \\ &\eta + S_i = (1-p)S_{i-1} + pS_{i+1}, \qquad i = 1, 2, \ldots, N, \\ &S_{N+1} = S_N, \end{aligned} \tag{31.6}$$

and (31.5) becomes

$$\begin{aligned} &S_0 = (1-p)(C + \gamma S_0) + p\gamma S_2, \\ &S_i = (1-p)\gamma S_{i-1} + p\gamma S_{i+1}, \qquad i = 1, 2, \ldots, N, \\ &S_{N+1} = S_N, \end{aligned} \tag{31.7}$$

where in (31.6) S_j is the expected (wait-and-see) discounted cost of meeting demand starting from reservoir level $i = j$, and in (31.7) S_j is the corresponding value in the average cost per period model.

31.3.1 ▪ Thermal Generation

The Stage and Larsson methodology can be extended to a setting where the policy being simulated generates energy using thermal plant (at some cost) as well as using releases from the hydro reservoir. This is the basis for a number of reservoir optimization models in practical use (see, e.g., [326]). It is important to observe that the values w_i derived from such a simulation will only be an expected marginal value for an optimization problem, if the policy being simulated is indeed optimal for that problem. In other words, the expected marginal value of water is the expected extra benefit that accrues from using this water in the most beneficial way. The optimal policy ultimately requires the solution of a stochastic dynamic optimization problem.

To include thermal generation in our example, we alter demand so that it is two units per period and suppose there is a thermal generator that costs $c < C$ to produce one unit of energy. Thus, to meet demand, we can either release two units of water or release one unit of water and run the thermal plant. The thermal generator is assumed to be flexible in that it can be switched on and off with no penalty. It is typical to specify an optimal policy in terms of a marginal water value as follows. If $c < w_i$, then we run the thermal plant. If $c > w_i$, then we release water. If $c = w_i$, then the policy can be any combination of thermal generation and water release.

In this model, the Stage and Larsson methodology (with discounting) might yield the following marginal water value algorithm:

$$\begin{array}{llll}
\text{Set} & \nu = 0, \quad w_i^\nu = 0; & & \\
\text{repeat} & & & \\
\quad \text{for} & i = 0 \text{ to } N \text{ do} & & \\
 & \text{if } i = 0, & \text{set } w_i^{\nu+1} = & C; \\
 & \text{if } i = N, & \text{set } w_i^{\nu+1} = & 0; \\
 & \text{if } 0 < i < N, & \text{if } w_i^\nu < c, & w_i^{\nu+1} = (1-p)\gamma w_{i-2}^\nu + p\gamma w_i^\nu; \\
 & & \text{if } w_i^\nu > c, & w_i^{\nu+1} = (1-p)\gamma w_{i-1}^\nu + p\gamma w_{i+1}^\nu; \\
\quad \text{end} & \text{for}; & & \\
\quad \text{set} & \nu := \nu + 1; & & \\
\text{until} & w^{\nu+1} = w^\nu. & &
\end{array}$$

If we relax the integrality constraint so that the thermal plant can generate any amount between zero and one, then the marginal water value will equal c when this plant is partially dispatched. There is a range of fractional water levels between x and $x+1$, say, for which this will be the optimal action. For almost every instance of this relaxed problem, the marginal value of water will equal c at exactly one integer level i and be different from c for other levels. The correct stationarity conditions (for any choice of γ) that correspond to the minimum cost solution are

$$\begin{aligned}
&w_0 = C, \\
&w_i = \gamma\big((1-p)w_{i-1} + p w_{i+1}\big), && i = 1,\dots,k-1, \\
&w_k = c \in [w_{k+1}, w_{k-1}], && i = k, \\
&w_i = \gamma\big((1-p)w_{i-2} + p w_i\big), && i = k+1,\dots,N-1, \\
&w_N = 0.
\end{aligned} \tag{31.8}$$

The correct value of k can be found by computing an optimal policy. Let W_i denote the expected future discounted cost of meeting demand. The here-and-now fixed-point problem (with discounting) is

$$\begin{aligned}
&W_0 = C + c + \gamma((1-p)W_0 + pW_2), \\
&W_1 = c + \gamma((1-p)W_0 + pW_2), \\
&W_i = \min\{c + \gamma\big((1-p)W_{i-1} + pW_{i+1}\big), \gamma\big((1-p)W_{i-2} + pW_i\big)\}, \quad i = 2,\dots,N, \\
&W_{N+1} = W_N.
\end{aligned} \tag{31.9}$$

The optimal solution gives a threshold type of policy, i.e., we release two units of water if

$$\gamma\big((1-p)W_{i-2} + pW_i\big) < c + \gamma\big((1-p)W_{i-1} + pW_{i+1}\big)$$

or

$$c > \gamma\big((1-p)w_{i-2} + pw_i\big).$$

The right-hand side is the expected discounted marginal water value if we meet demand from hydro only, and water is released if this is less than the cost c of thermal generation.

An optimal policy can be found using linear optimization (LO) and determines a threshold index k below which the thermal plant is always run at capacity and above which it is not. Given this value of k, the optimal policy satisfies

$$\begin{aligned}
&W_0 = C + c + \gamma((1-p)W_0 + pW_2), \\
&W_1 = c + \gamma((1-p)W_0 + pW_2), \\
&W_i = c + \gamma\big((1-p)W_{i-1} + pW_{i+1}\big), && i = 2,\dots,k, \\
&W_i = \gamma\big((1-p)W_{i-2} + pW_i\big), && i = k+1,\dots,N \\
&W_{N+1} = W_N.
\end{aligned}$$

By setting $w_i = W_i - W_{i+1}$, we can derive the marginal values w that will hold at optimality. These satisfy

$$\begin{aligned} w_0 &= C, \\ w_i &= \gamma\big((1-p)w_{i-1} + p w_{i+1}\big), \quad i = 1,\ldots,k-1, \\ w_k &= c \in [w_{k+1}, w_{k-1}], \\ w_i &= \gamma\big((1-p)w_{i-2} + p w_i\big), \quad i = k+1,\ldots,N-1, \\ w_N &= 0, \end{aligned} \tag{31.10}$$

which are the same equations as (31.8).

The case without discounting is similar. If we form an average cost per period Markov decision problem, then we seek the solution to

$$\begin{array}{ll} \max & g \\ & W_0 = -\eta + C + c + (1-p)W_0 + pW_2, \\ & W_1 = -\eta + c + (1-p)W_0 + pW_2, \\ & W_i \le -\eta + c + (1-p)W_{i-1} + pW_{i+1}, \quad i = 2,\ldots,N, \\ & W_i \le -\eta + (1-p)W_{i-2} + pW_i, \quad i = 2,\ldots,N, \\ & W_{N+1} = W_N. \end{array} \tag{31.11}$$

As before, the optimal policy can be found using LO and determines a threshold index k below which the thermal plant is always run at capacity and above which it is not. Given this value of k, we can derive the marginal values w that will hold at optimality. This gives the system

$$\begin{aligned} w_0 &= C, \\ w_i &= (1-p)w_{i-1} + p w_{i+1}, \quad i = 1,\ldots,k-1, \\ w_k &= c, \\ w_i &= (1-p)w_{i-2} + p w_i, \quad i = k+1,\ldots,N-1, \\ w_N &= 0, \end{aligned} \tag{31.12}$$

which is the same as (31.8) when $\gamma = 1$. In fact in the average cost per period Markov decision problem for this system, k turns out to be always equal to $N-1$, and

$$w_i = \begin{cases} C, & i = 0, \\ \frac{ci + C(N-1-i)}{N-1}, & i = 1,\ldots,N-2, \\ c, & i = N-1, \\ 0, & i = N, \end{cases} \tag{31.13}$$

when $p = 0.5$ and

$$w_i = \begin{cases} C, & i = 0, \\ \begin{array}{l} C\frac{1-r}{1-r^{N-1}}\left(r^i + r^{i+1} + \cdots + r^{N-2}\right) \\ +c\frac{1-r}{1-r^{N-1}}\left(1 + r + \cdots + r^{i-1}\right) \end{array}, & i = 1,\ldots,N-2, \\ c, & i = N-1, \\ 0, & i = N, \end{cases} \tag{31.14}$$

otherwise.

Proposition 31.2. *If $C > c$, then there is a unique monotonic solution to* (31.12) *defined by choosing $k = N-1$.*

Proof. If $k < N-1$, then

$$w_{k+1} = \big((1-p)w_{k-1} + p w_{k+1}\big),$$

yielding

$$w_{k+1} = w_{k-1},$$

and hence by monotonicity of w_i, we obtain

$$w_{k-1} = w_{k+1} = w_k = c.$$

Solving

$$w_i = (1-p)w_{i-1} + p w_{i+1}$$

for $i = 1, \ldots, k-1$ gives

$$w_i = c, \quad i = 0, \ldots, k+1,$$

contradicting $w_0 = C > c$. Thus $k = N-1$, and the unique solution is given by (31.13) if $p = 0.5$, or (31.14) if $p \neq 0.5$. □

31.4 ▪ Multiple-Reservoir Models

In most real applications, marginal water values are obtained from the optimal operation of many reservoirs, sometimes linked in a cascaded river system, and almost always linked by some electricity transmission system. When the number of reservoirs is larger than two or three, representing their optimal operation using Markov decision problems becomes more difficult, due to the curse of dimensionality. Approximations of various forms are necessary.

The approximation in widespread use is based on the SDDP algorithm of Pereira and Pinto [1458]. Here the marginal values of water emerge from an approximate stage problem that is modeled as an LO problem. This method can be shown to converge to an optimal policy almost surely under mild assumptions on the sampling process (see [804, 1471]). The subgradients of the optimal value function define the marginal water value.

The class of problems we consider have T stages, denoted by $t = 1, 2, \ldots, T$, in each of which a random right-hand side vector $b_t(\omega_t) \in \mathbb{R}^m$ has a finite number of realizations defined by $\omega_t \in \Omega_t$. We assume that the outcomes ω_t are stagewise independent, and that Ω_1 is a singleton, so the first-stage problem is

$$\begin{aligned} z = \min \quad & c_1^\top x_1 + \mathbb{E}[Q_2(x_1, \omega_2)] \\ \text{s.t.} \quad & A_1 x_1 = b_1, \\ & x_1 \geq 0, \end{aligned} \tag{31.15}$$

where $x_1 \in \mathbb{R}^n$ is the first-stage decision, and $c_1 \in \mathbb{R}^n$ is a cost vector, A_1 is an $m \times n$ matrix, and $b_1 \in \mathbb{R}^m$.

We denote by $Q_2(x_1, \omega_2)$ the second-stage costs associated with decision x_1 and realization $\omega_2 \in \Omega_2$. The problem to be solved in the second and later stages t, given decisions x_{t-1} and realization ω_t, can be written as

$$\begin{aligned} Q_t(x_{t-1}, \omega_t) = \min \quad & c_t^\top x_t + \mathbb{E}[Q_{t+1}(x_t, \omega_{t+1})] \\ \text{s.t.} \quad & A_t x_t = b_t(\omega_t) - E_t x_{t-1} \qquad [\pi_t(\omega_t)], \\ & x_t \geq 0, \end{aligned} \tag{31.16}$$

where $x_t \in \mathbb{R}^n$ is the decision in stage t, c_t its cost, and A_t and E_t $m \times n$ matrices. Here $\pi_t(\omega_t)$ denotes the dual variables of the constraints. In the last stage we assume either that $\mathbb{E}[Q_{T+1}(x_T, \omega_{T+1})] = 0$ or that there is a convex polyhedral function that defines the expected future cost after stage T. For all instances of (31.16), we assume relatively complete recourse, whereby (31.16) at stage t has a feasible solution for all values of x_{t-1} that are feasible for the instance of (31.16) at stage $t-1$. Relatively complete recourse can be ensured by introducing artificial variables with penalty terms in the objective.

The SDDP algorithm performs a sequence of major iterations, each consisting of a *forward pass* and a *backward pass* through all the stages, to build an approximately optimal policy. In each forward pass, a set of N scenarios is sampled from the scenario tree and decisions are made for each stage of those N scenarios, starting in the first stage and moving forward up to the last stage. In each stage, the observed values $\bar{x}_t(s)$ of the decision variables x_t, and the costs of each stage in all scenarios s, are saved.

The SDDP algorithm builds a policy that is defined at stage t by a polyhedral outer approximation (OA) of $\mathbb{E}[Q_{t+1}(x_t, \omega_{t+1})]$. This approximation is constructed using cutting planes called Benders cuts, or just *cuts*. In other words, in each tth-stage problem, $\mathbb{E}[Q_{t+1}(x_t, \omega_{t+1})]$ is replaced by the variable θ_{t+1}, which is constrained by a set of linear inequalities

$$\theta_{t+1} - \bar{g}_{t+1,k,s}^\top x_t \geq \bar{h}_{t+1,k,s}, \qquad k = 1,2,\ldots,K, \quad s = 1,2,\ldots,N, \tag{31.17}$$

where K is the number of backward passes that have been completed and $\bar{g}$ and $\bar{h}$ are defined by (31.20) and (31.21) below.

With this approximation, the first-stage problem is

$$\begin{aligned} z = \min \quad & c_1^\top x_1 + \theta_2 \\ \text{s.t.} \quad & A_1 x_1 = b_1, \\ & \theta_2 - \bar{g}_{2,k,s}^\top x_1 \geq \bar{h}_{2,k,s}, \qquad k = 1,2,\ldots,K, \\ & \qquad\qquad\qquad\qquad\quad s = 1,2,\ldots,N, \\ & x_1 \geq 0, \end{aligned} \tag{31.18}$$

and the tth stage problem becomes

$$\begin{aligned} \tilde{Q}_t(x_{t-1}, \omega_t) = \min \quad & c_t^\top x_t + \theta_{t+1} \\ \text{s.t.} \quad & A_t x_t = b_t(\omega_t) - E_t x_{t-1}, \qquad [\pi_t(\omega_t)] \\ & \theta_{t+1} - \bar{g}_{t+1,k,s}^\top x_t \geq \bar{h}_{t+1,k,s}, \quad k = 1,2,\ldots,K, \\ & \qquad\qquad\qquad\qquad\qquad s = 1,2,\ldots,N, \\ & x_t \geq 0, \end{aligned} \tag{31.19}$$

where we interpret the set of cuts as being empty when $K = 0$.

At the end of the forward pass, a convergence criterion is tested, and if it is satisfied then the algorithm is stopped; otherwise it starts the backward pass, which is defined below. In the standard version of SDDP (see [1458]), the convergence test is satisfied when z, the lower bound on the expected cost at the first stage, is statistically close to an estimate of the expected total operation cost obtained by averaging the cost of the policy defined by the cuts when applied to the N sampled scenarios. In this simulation, the total operation cost for each scenario is the sum of the present cost ($c_t^\top x_t$) over all stages t and any end-of-horizon future cost.

If the convergence criterion is not satisfied, then SDDP amends the current policy using a backward pass that adds N cuts to each stage problem, starting at the penultimate stage and working backward to the first. To compute the coefficients for the cuts, we solve the next-stage problems for all possible realizations (Ω_{t+1}) in each stage t and scenario s. The cut for (31.19), the tth (approximate) stage problem in scenario s, is computed after its solution $\bar{x}_t^k(s)$ has been obtained in the forward pass immediately preceding the backward pass k. Solving the $(t+1)$th (approximate) stage problem for every $\omega_{t+1} \in \Omega_{t+1}$ gives $\bar{\pi}_{t+1,k,s} = \mathbb{E}[\pi_{t+1}(\omega_{t+1})]$, which defines the cut gradient

$$\bar{g}_{t+1,k,s} = -\bar{\pi}_{t+1,k,s}^{\top} E_{t+1} \tag{31.20}$$

and its intercept

$$\bar{h}_{t+1,k,s} = \mathbb{E}[Q_{t+1}(\bar{x}_t^k(s), \omega_{t+1})] + \bar{\pi}_{t+1,k,s}^{\top} E_{t+1} \bar{x}_t^k(s). \tag{31.21}$$

The SDDP algorithm is initialized by setting $\theta_t = -\infty$, $t = 2, \ldots, T$, $K = 0$, $k = 1$. Thereafter, the algorithm performs the following three steps repeatedly until the convergence criterion is satisfied:

1. Forward Pass

   For $t = 1$, solve (31.18) and save $\bar{x}_1^k(s) = x_1$, $s = 1, \ldots, N$, and $\bar{z}^k = z$;

   For $t = 2, \ldots, T$ and $s = 1, \ldots, N$,

   Solve (31.19) setting $x_{t-1} = \bar{x}_{t-1}^k(s)$, and save $\bar{x}_t^k(s)$ and $\tilde{Q}_t(\bar{x}_{t-1}^k(s), \omega_t)$.

2. Standard Convergence Test (at $100(1-\alpha)\%$ confidence level).

   Calculate the upper bound: $z_u = \frac{1}{N} \sum_{s=1}^{N} \sum_{t=1}^{T} c_t^{\top} \bar{x}_t^k(s)$,

   $$\sigma_u = \sqrt{\frac{1}{N} \sum_{s=1}^{N} \left(\sum_{t=1}^{T} c_t^{\top} \bar{x}_t^k(s) \right)^2 - z_u^2}.$$

   Calculate the lower bound: $z_l = \bar{z}^k$;

   Stop if

   $$z_l > z_u - \frac{Z_{\frac{\alpha}{2}}}{\sqrt{N}} \sigma_u,$$

   where Z_α is the $(1-\alpha)$ quantile of the standard normal distribution; otherwise go to the backward pass.

3. Backward Pass

   For $t = T, \ldots, 2$, and $s = 1, \ldots, N$,

   For $\omega_t \in \Omega_t$, solve (31.19) using $\bar{x}_{t-1}^k(s)$ and save $\bar{\pi}_{t,k,s} = \mathbb{E}[\pi_t(\omega_t)]$ and $\tilde{Q}_t(\bar{x}_{t-1}^k(s), \omega_t)$;

   Calculate the kth cut for s in stage $t-1$ using (31.20) and (31.21).

   Set $K = K + 1$, $k = k + 1$.

SDDP terminates when the estimated value z_u of the current policy is closer than $\frac{Z_{\frac{\alpha}{2}}}{\sqrt{N}} \sigma_u$ to a lower bound z_l on the optimal value. The expected value of an optimal policy for the model (with finitely many inflow outcomes) can then be assumed to lie in the interval $[z_l, z_u + \frac{z_{\frac{\alpha}{2}}}{\sqrt{N}}\sigma_u]$ with probability at least $(1-\alpha)$.

Energy prices are deduced from the solution to SDDP by simulating a single stage of the (approximately) optimal policy (using cutting planes to represent the Bellman function) and recording the shadow prices on the demand constraint. Marginal water values in each reservoir at level $\bar{x}$ are determined from the subdifferential of $\mathbb{E}[\tilde{Q}_t(x_{t-1}, \omega_t)]$ evaluated at $\bar{x}$. Often this will be a singleton defined by the coefficients of the highest cut at $\bar{x}$.

It is important to observe that the above algorithm operates under a wait-and-see assumption. In other words, we solve (31.19) with full knowledge of the current period's inflow, which can be used to generate electricity to satisfy demand. This means that shortage costs might be avoided even when all reservoir levels are at their minimum (see Section 31.2 above). Since shortages arising in simulations of an optimal policy yield marginal water value estimates, the wait-and-see assumption is likely to underestimate prices.

Based on the single-reservoir analysis, an optimal generation schedule will use a merit order of water release and thermal generation based on marginal water value and marginal cost. With linear dynamics and with linear production functions, this produces a bang-bang type of control policy. It is well known, however, that many such systems admit singular control (i.e., one not uniquely determined by the adjoint solution). For example, the separated continuous LO problems discussed in [81] have this property. Examples in a hydro-scheduling context are given in [1468].

31.5 ▪ Application

We now show how SDDP can be applied to estimate wholesale prices for the New Zealand electricity system under the assumption of perfectly competitive risk-neutral agents. We begin by solving a model in which the inflows are assumed to be stagewise independent. We use the DOASA software described in [1473].

We seek to estimate New Zealand wholesale electricity prices for the week beginning March 10, 2008. The physical characteristics of the system are described in [1473] and [1472] (which together provide a more detailed study of the New Zealand wholesale electricity market over the period 2001–2008). The electricity network shown in Figure 31.2 is approximated by a four-reservoir system by aggregating the storage reservoirs in the Waitaki river system into a single lake (Lake Waitaki). The resulting Waitaki system is shown in Figure 31.3.

Lake Ohau (OHU), which has limited storage capacity, is now treated as a run-of-river resource, and reservoirs Benmore (BEN), Pukaki (PKI), and Tekapo (TEK) are treated as a single lake. Since these occur in a cascade in the river system, the aggregation of water in each must account for their specific energies. Thus Benmore and Pukaki storage and inflows are discounted by their specific energies before being added to the aggregate Lake Waitaki, which has the same specific energy as Tekapo. To allow for the storage ability of Benmore and Pukaki, Lake Waitaki water can be restored to these lower lakes through the arcs connecting Lake Waitaki to its children. These arcs have multipliers that inflate the water volume transferred to correspond to the difference in specific energies of Lake Waitaki and the lower lakes.

The transmission network is approximated by three nodes, one for the South Island and two for the North Island. Limits on power transfer between nodes correspond to reported transmission limits. Energy deficit in any stage is met by load shedding at an

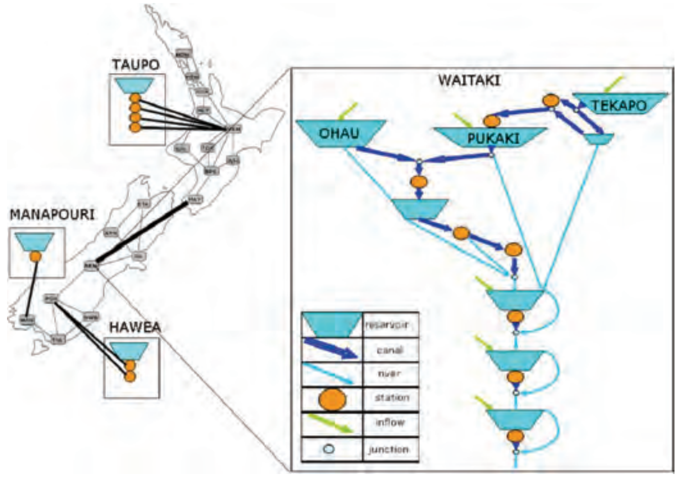

Figure 31.2. *Approximate network representation of New Zealand electricity network showing main hydroelectricity generators. Reprinted with permission from Elsevier* [1473].

increasing shortage cost in three tranches. This is equivalent to having three dummy thermal plants at each node with capacities equal to 5% of load, 5% of load, and 90% of load, for each load sector, and costs as shown in Table 31.1. Load shed above 10% of demand costs a value of lost load of $20,000/MWh, which is roughly equivalent to five hours of load curtailment per year assuming annual amortized peaking-plant costs of $100/kW.

Table 31.1. *Load reduction costs (NZD/MWh) and proportions of industrial load, commercial load, and residential load in each island.*

	Up to 5%	5% to 10%	Above 10%	North Is.	South Is.
Industrial	$1,000	$2,000	$20,000	0.34	0.58
Commercial	$2,000	$4,000	$20,000	0.27	0.15
Residential	$2,000	$4,000	$20,000	0.39	0.27

We choose historical inflows from each week in the years 1997–2006 to give an empirical inflow distribution with 10 outcomes per stage. We assume that these are stagewise independent.

The first model attempts to compute an optimal policy for 52 weeks starting at midnight on March 10, 2008. The computed electricity prices for the week beginning March 11, 2008, are the same at each node. They are $38.45/MWh, $38.76/MWh, and $39.84/MWh, corresponding to off-peak, shoulder, and peak periods. We can simulate the corresponding policy using inflows sampled from the assumed stagewise-independent distribution. This gives storage levels in Lake Waitaki from 100 random sequences, as shown in Figure 31.4. If we simulate the same policy with the 30 most recent historical inflow sequences, then we obtain storage levels in Lake Waitaki as shown in Figure 31.5. The storage levels shown hit both extremes more often than the simulation that assumes stagewise independence.

A heuristic that has been suggested to overcome the optimistic biases from assuming stagewise independence is called *inflow spreading* [1548]. Assuming independence over k weeks results in a lower variance in total inflow over this period than the data

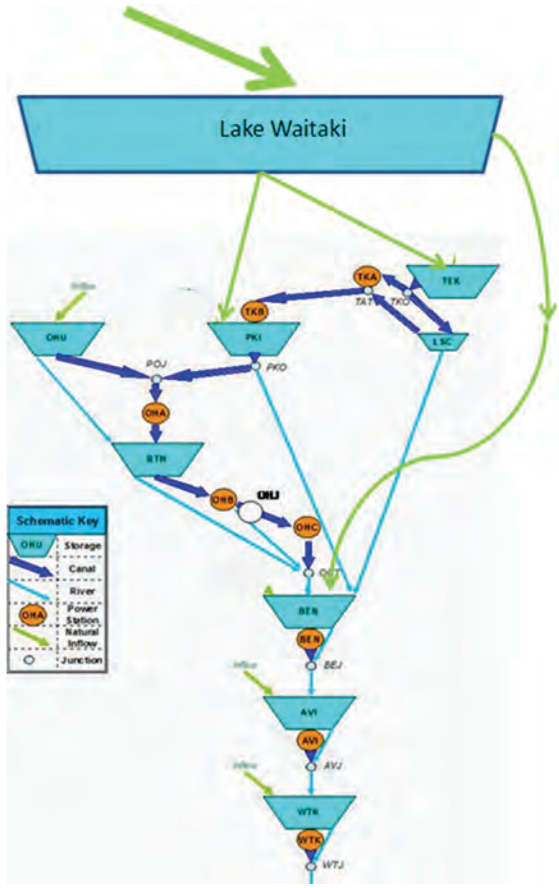

Figure 31.3. *Aggregated model of Waitaki system with storage lakes Pukaki, Tekapo, and Benmore combined to make Lake Waitaki. Their inflows (scaled by specific energies) combine to produce an inflow to Lake Waitaki.*

support. Inflow spreading modifies the inflows using a heuristic that increases the variance of assumed independent inflows so that the variance of k consecutive independent inflows matches the observed variance of this sum.

Consider historical inflow sequence $h(t,y)$ of inflows in week t of year $y = 1,2,\ldots,N$, and suppose $k = 4$. Let

$$w(t,y) = h(t,y) + h(t+1,y) + h(t+2,y) + h(t+3,y),$$

$$a(t) = \frac{\sum_y h(t,y)}{N}, \quad b(t) = \frac{\sum_y w(t,y)}{N},$$

and

$$d(t,y) = \max\left\{0, a(t) + \frac{(w(t,y) - b(t))}{2}\right\}.$$

The model then uses adjusted inflows

$$\omega(t,y) = \frac{N d(t,y) a(t)}{\sum_y d(t,y)}, \quad t = 1,2,\ldots,52, \quad y = 1,2,\ldots,N.$$

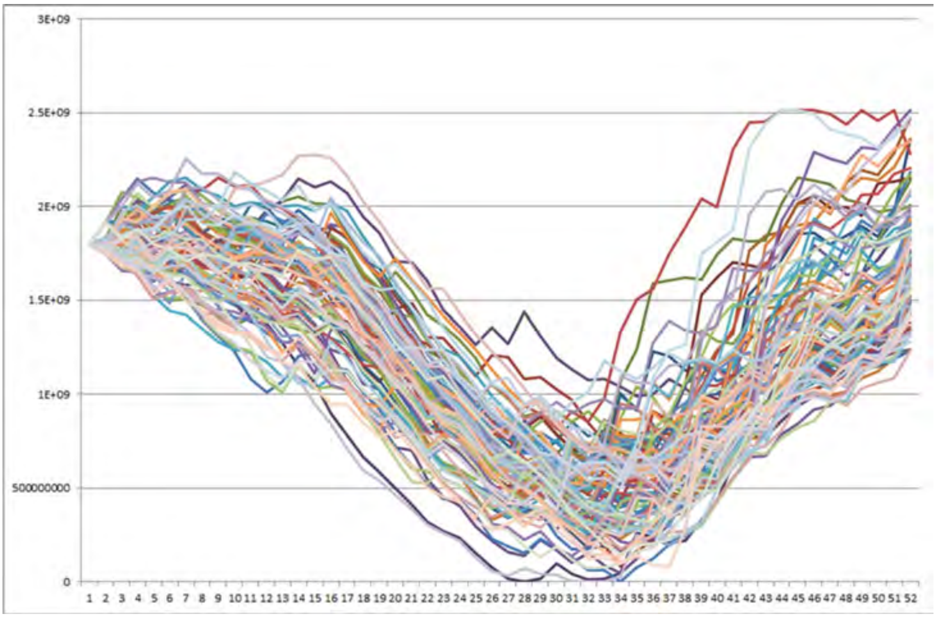

Figure 31.4. *Simulated volumes in Lake Waitaki reservoir for* 100 *randomly generated inflow sequences.*

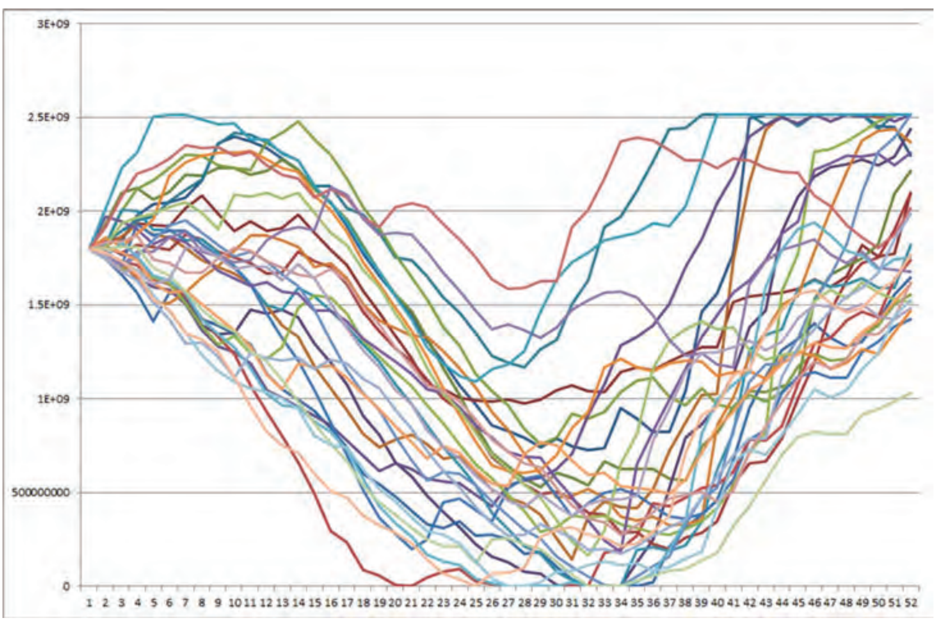

Figure 31.5. *Simulated volumes in Lake Waitaki reservoir for historical inflow sequences* 1979–2008.

An optimal policy for 52 weeks starting at midnight on March 10, 2008, can be computed using adjusted inflows with $k = 4$. We can simulate the corresponding policy using inflows sampled from the assumed stagewise-independent distribution. This gives storage levels in Lake Waitaki from 100 random sequences (with spread inflows)

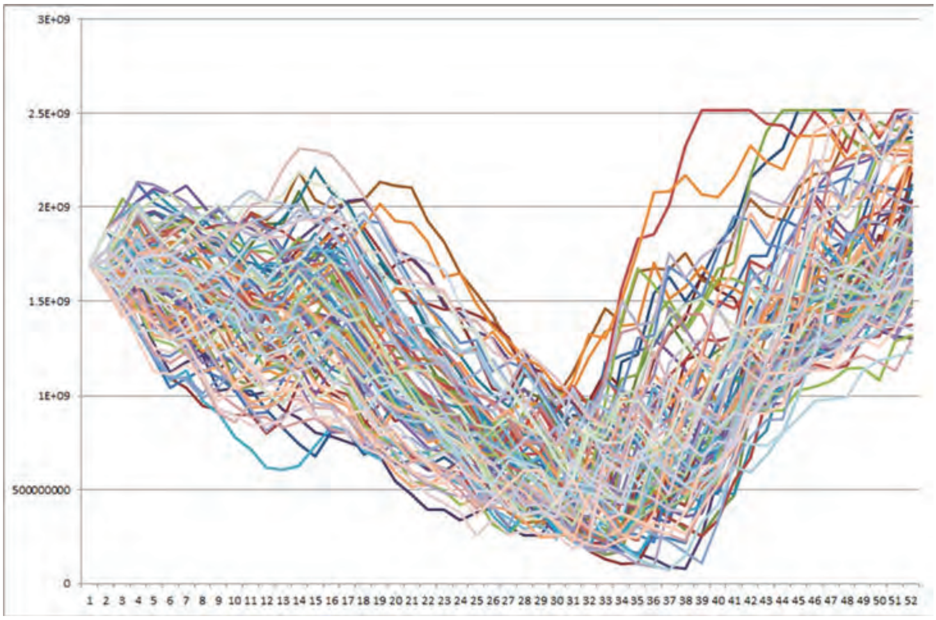

Figure 31.6. *Inflow spreading policy simulated volumes in Lake Waitaki reservoir for* 100 *randomly generated inflow sequences.*

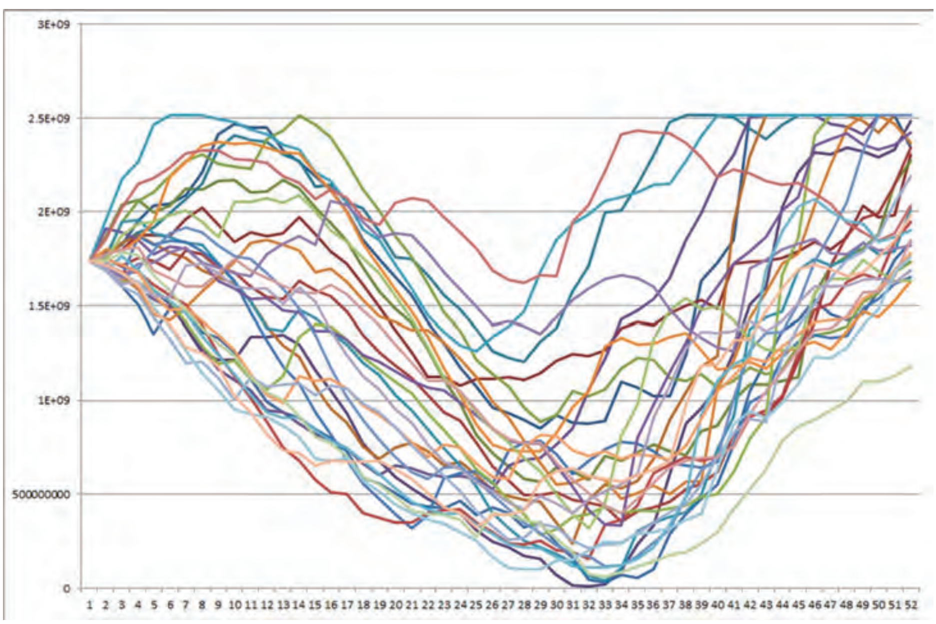

Figure 31.7. *Inflow spreading policy simulated volumes in Lake Waitaki reservoir for historical inflow sequences* 1979–2008.

as shown in Figure 31.6. If we simulate the same policy with the 30 most recent historical inflow sequences, then we obtain storage levels in Lake Waitaki as shown in Figure 31.7. The stock levels hit zero in only one year out of 30. The electricity prices for the week beginning March 11, 2008, are computed and are the same at each node. They are \$40.10/MWh, \$40.81/MWh, and \$41.02/MWh, corresponding to off-peak, shoulder, and peak periods. The inflow spreading has resulted in slightly higher prices.

The prices estimated by our model have been computed using SDDP, which makes use of a sampling procedure. This means that there will be some sampling error in their estimation. We first attempt to estimate this error by simulating the policy that we have obtained with perturbed initial reservoir levels. A simulation of the optimal (inflow spreading) policy over 2,500 random inflow sequences yields a sample average optimal cost of \$354,861,576. Increasing the initial reservoir volume by 1 million cubic meters and simulating with common random numbers yields a sample average optimal cost of \$354,812,679. The difference in costs is \$48,897, with a standard error of \$2,685. Assuming that the SDDP policy is optimal yields an estimate of marginal water value for Lake Waitaki of \$0.0489/m^3, with a 95% confidence interval [0.0435, 0.0543]. To convert a cubic meter to an energy value, we must divide by the specific energy of Lake Waitaki, which is 0.001131 MWh/m^3, giving a 95% confidence interval for the perfectly competitive wholesale Benmore price of [\$38.46, \$48.01].

31.6 ▪ Discussion

One can compare the computed estimates of weekly prices with observed prices in the New Zealand wholesale market in the week beginning March 11, 2008. These range from \$60/MWh to \$190/MWh. The marginal water values computed by our SDDP algorithm are significantly lower than those reflected in observed electricity prices. There are several possible causes of this.

1. Poor representation of stagewise dependence in inflows: We have represented stagewise dependence in inflows using an inflow spreading heuristic based on the approach currently used by some New Zealand electricity generators. We have also ignored the melting of snowpack, which affects the inflow dynamics in the Waitaki catchment. The true reservoir inflow processes in New Zealand exhibit a more complicated long-range dependence than we have in our model. New inflow models that aim to improve the modeling of this dependence are discussed in [1518].
2. Wait-and-see representation of uncertainty: Our SDDP model assumes that inflows are observed in any week before release decisions are made for that week. As discussed above, this assumption implies that the model will underestimate the probability of a shortage event, which will underestimate marginal water values. On the other hand, local meteorologists make reasonably reliable weather forecasts for the next three to five days, so we conjecture that the wait-and-see model should give reasonably accurate estimates, all else being equal.
3. Assumptions about constant fuel costs not reflecting fuel contract structure: In the New Zealand wholesale market, all large electricity generators (including thermal generators) offer increasing supply functions to the system operator. We have assumed that thermal short-run marginal cost is constant for each plant, so these supply functions in our model are assumed to be constant. Increasing supply functions for thermal plant are often interpreted as evidence of market power exercise, but they have alternative explanations. The owners of these plants typically operate under fuel contracts that affect the cost of their day-to-day fuel consumption. For example, coal is stockpiled, and so if replenishment is expensive

and time-consuming, its marginal value increases as a stock is depleted. On the other hand, gas is typically supplied under take-or-pay contracts over a contract period of several months. This effectively makes above-average consumption in any period more costly as it incurs a risk of exceeding the contract volume and having to pay a penalty.

4. Agents making different assumptions about the probabilities of dry inflow outcomes: We have assumed that all agents have the same probability distribution of inflow outcomes, sampled with equal probability from historical records. Recent climate observations (such as El Nino Southern Oscillation forecasts) might affect agents' subjective assessments of these probabilities.
5. Agents acting strategically in offering generation at higher than competitive prices: We have assumed that all agents act as price-takers. This is not required by the rules governing the operation of the New Zealand wholesale electricity market. Some increase in price might be attributed to agents behaving strategically.
6. Agents not being risk neutral: We have assumed that all agents are risk neutral, and so they act to maximize expected profits. However, generators and purchasers are generally risk averse, so there is a possible discrepancy between a social optimum (and corresponding marginal water values) and a competitive partial equilibrium. This discrepancy is explored in detail in [1470], where competitive agents facing uncertain reservoir inflows are endowed with possibly different coherent time-consistent risk measures. In [1470], it is shown that if enough contracts are made available to agents to trade, and they are not too different in their attitudes to risk, then a social risk measure emerges in equilibrium that can be used to compute a risk-averse social plan that coincides with the partial competitive equilibrium. The risk-averse social plan can be computed using a risk-averse version of SDDP (see [1469]). In this setting, prices will increase as agents trade off their current costs against risk-adjusted future costs, which will make them more conservative than a risk-neutral agent when water is scarce. If the market for risk is incomplete, so thermal plant owners and hydro plant owners have no ability to hedge their different risk positions, then a perfectly competitive equilibrium need not correspond to a social planning solution, even if the planning solution is risk averse. Of course, even if we assume completeness in markets for risk, agents in hydro-dominated electricity markets face a plethora of commercial risks apart from uncertain inflows (which are the sole source of uncertainty in our model). It is conceivable that these other uncertainties are regarded as more critical than inflow risk, influencing the price markups we observe in the week beginning March 11, 2008.

Identifying which combination of these factors is driving the difference between modeled and observed prices is difficult. However, for regulators of hydro-dominated electricity markets who monitor inefficiencies arising from imperfect competition or market incompleteness, this is a critically important task. SO models and stochastic equilibrium models are powerful tools to enable this to happen.