Risk-Averse Reservoir Management in a De-regulated Electricity Market

M. Craddock, A. D. Shaw and B. Graydon
Strategic Development Group
Electricity Corporation of New Zealand*
New Zealand
mark.craddock@ecnz.co.nz, andrew.shaw@ecnz.co.nz
bgraydon@phb.co.nz

Abstract

In a deregulated electricity market, there are three major influences on the strategy of a generation company, market interactions, contract positions and attitude towards risk. An optimisation tool to address the medium-term scheduling problem should incorporate all three factors when formulating a strategy. Recent research has shown much progress in market interactions and contract position, and independently progress has also been made in looking at risk aversion. This paper draws together this work and presents a medium-term model which combines all three factors. We provide an overview of the formulation of this model and discuss the behaviour of the model. We also outline some future directions for this model.

1 Introduction

With the advent of deregulated electricity markets around the world there is a need for a paradigm shift in the design of optimisation models for companies operating in these environments. The emphasis has shifted from systemwide cost minimisation to revenue maximisation for each participating company. In addition, financial tools for risk management, such as two-way hedge contracts, coupled with a growing awareness at board level of the need to manage risk exposure have added to the complexity of the problem to be solved.

Hedge contracts are financial agreements to compensate for a difference between the contract price and the underlying spot price. For both the seller and customer they provide price certainty for a fixed quantity of demand. However, the generator faces an uncertain cost stream. In compensation for assuming this increased risk exposure the generator charges a contract fee. Utility and utility functions are

*The content of this paper represents the views of the authors and not that of ECNZ.

commonly used tools in risk management (see [1]). When a company's revenue stream is subject to a stochastic process (such as hydrology) typical bottom-line figures such as annual net revenue are defined by a probability distribution. Utility is simply a dimension-less value encapsulating an entire distribution, and the utility function transforms the distribution of net revenue into utility.

In the electricity industry when a market participant holds significant hydro storage the strategy employed in its management has substantial effects upon the distribution of the annual result. While high net revenues are obviously desirable, they become less desirable if under a different stochastic outcome the same strategy would produce an extremely low net revenue. A risk-averse strategy would reduce the variation of the annual net revenue distribution, as measured by the utility.

Previous work [2] has looked at the problem of risk-aversion in a centrally coordinated cost minimisation problem. Other work [3] has developed a market equilibrium model including hedge contract effects but without any account of risk attitude. This paper draws together these two concepts and outlines an algorithm to solve the complete problem. The following three sections outlines the method that we have developed and then present some preliminary results from the model.

2 Market Equilibrium Conditions

The optimisation problem for a risk-averse firm with an aggregated thermal and hydro station operating in a market with contracts can be stated as

$$\max_{g_t^b, g_t^h} \left\{ E\left[\alpha\left(w_T, s_T\right) \mid a_t, \forall t = 1, ..., T\right]\right\},$$

subject to

$$\begin{array}{lll} s_t = s_{t-1} + a_t - g_t^h, & & \forall t = 1, ..., T, \\ w_t = w_{t-1} + p_t(g_t^b + g_t^h + \hat{g}_t)(g_t^b + g_t^h - k_t) - c_t(g_t^b), & & \forall t = 1, ..., T, \\ \underline{s}_t \leq s_t \leq \overline{s}_t, & & \forall t = 1, ..., T, \\ \underline{g}_t^i \leq g_t^i \leq \overline{g}_t^i, & \forall i \in (h, b), & \forall t = 1, ..., T. \end{array}$$

The market spot price in period t is given by $p_t(.)$, a function of the total market generation $(g_t^b + g_t^h + \hat{g}_t)$. The thermal station produces g_t^b, the hydro station produces g_t^h, and the contracted quantity is k_t. The remaining quantity $\hat{g}_t$ is produced by the other players in the market. The amount of water stored in the reservoir at the end of period t is given by s_t, and the initial storage is a fixed quantity s_0. Thermal generation incurs a production cost $c_t(.)$ and the reservoir is subject to an uncertain inflow a_t during period t. Both the reservoir storage and generation quantities are subject to upper and lower bounds. In [2], w_t was introduced as the accumulated net revenue to period t. We define w_t thus

$$w_t = \sum_{\tau=1}^{t} p_\tau(g_\tau^b + g_\tau^h + \hat{g}_\tau)(g_\tau^b + g_\tau^h - k_\tau) - c_\tau(g_\tau^b),$$

which is simply the revenue earned on the spot market by the uncontracted quantity less the operating cost. Note that we are modelling a standard contract for differences hedge and have not included the payment for contracted quantities as this is fixed and can be removed from our optimisation. For reporting purposes the contract payments are added later. Finally, $\alpha(.,.)$ is the *utility function* describing the attitude to risk with respect to net revenue and storage that the firm is prepared to accept in the final period T. Changing the utility function appropriately models different risk averse strategies.

In the above formulation the utility function applies only to the two state variables of the final period. This is useful as each period t is independent of the others and therefore a standard dynamic programming technique can be applied. We choose to employ a dual dynamic programming technique similar to that used by [3]. Following the approach of [3], we form the Lagrangian function thus

$$\begin{aligned}\mathcal{L}(g_t^b, g_t^h, w_t, s_t) = E\left[\alpha\left(w_T, s_T\right) \mid a_t, \forall t = 1, ..., T\right] + \sum_{t=1}^{T} \pi_t \left(s_{t-1} - s_t + a_t - g_t^h\right) \\ + \sum_{t=1}^{T} \lambda_t \left(w_{t-1} - w_t + p_t(g_t^b + g_t^h + \hat{g}_t)(g_t^b + g_t^h - k_t) - c_t(g_t^b)\right) \\ + \text{bound terms.}\end{aligned}$$

We have introduced some lagrange multipliers π_t and λ_t for the two sets of constraints. Neglecting the variable bounds the first derivatives are

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial g_t^b} &= \lambda_t \left(\frac{dp_t}{dg_t^b}(g_t^b + g_t^h - k_t) + p_t(g_t^b + g_t^h + \hat{g}_t) - \frac{dc_t}{dg_t^b}\right), \quad t = 1, ...T, \\ \frac{\partial \mathcal{L}}{\partial g_t^h} &= -\pi_t + \lambda_t \left(\frac{dp_t}{dg_t^h}(g_t^b + g_t^h - k_t) + p_t(g_t^b + g_t^h + \hat{g}_t)\right), \quad t = 1, ...T, \\ \frac{\partial \mathcal{L}}{\partial w_t} &= \lambda_{t+1} - \lambda_t, \quad t = 1, ...T-1, \\ \frac{\partial \mathcal{L}}{\partial s_t} &= \pi_{t+1} - \pi_t, \quad t = 1, ...T-1, \\ \frac{\partial \mathcal{L}}{\partial w_T} &= \frac{\partial \alpha}{\partial w_T} - \lambda_T, \\ \frac{\partial \mathcal{L}}{\partial s_T} &= \frac{\partial \alpha}{\partial s_T} - \pi_T.\end{aligned}$$

Setting the first derivatives to zero and again neglecting the bound terms, we get

$$\begin{aligned}p_t(g_t^b + g_t^h + \hat{g}_t) &= \frac{dc_t}{dg_t^b} - \frac{dp_t}{dg_t^b}(g_t^b + g_t^h - k_t), \quad t = 1, ...T, \\ p_t(g_t^h + g_t^h + \hat{g}_t) &= \frac{\pi_t}{\lambda_t} - \frac{dp_t}{dg_t^h}(g_t^b + g_t^h - k_t), \quad t = 1, ...T, \\ \lambda_{t+1} = \lambda_t, \pi_{t+1} &= \pi_t, \quad t = 1, ...T-1, \\ \lambda_T = \frac{\partial \alpha}{\partial w_T}, \pi_T &= \frac{\partial \alpha}{\partial s_T}.\end{aligned}$$

These are the necessary optimality conditions for a risk-averse firm operating within a competitive market. The first four conditions are the optimality conditions for the firm during a period t, and the final two conditions are the end horizon conditions. We have omitted the bound terms in this formulation as they will be implied by the cost of shortage and spill by the the dynamic programming recursion.

3 Applying a Stochastic Dynamic Programming Recursion

With two notable exceptions the conditions obtained in the previous section are identical to the conditions derived by [3]. These exceptions are the introduction of λ_t and the two end of horizon conditions. We may think of λ_t as the *marginal value of accumulated net revenue* or the marginal utility of an extra dollar in period t. The end horizon conditions are interesting as they are the only interaction the utility function has with the optimisation problem. This is a strength of the technique as a variety of utility functions can be used without changing the optimisation routines.

The optimality conditions above apply to a single risk-averse firm with hydro storage. Now a market consists of a number of firms all of which seek to optimise a similar problem. Each firm in the market will be responding to the same market price as given by the market demand curve. As each firm can control its own output g and hence influence the market price $p(.)$, they influence the behaviour of the other firms. Furthermore, this then feedbacks onto its own behaviour. This dilemma is resolved in [3] by assuming that each firm acts as a Cournot oligopolist. This assumption implies that each firm will determine an optimal generation level assuming the output of the other firms is fixed. The point where each firm finds its optimal generation level is refered to as the market equilibrium position.

As we are considering a market where only one firm has significant hydro storage and that firm is the only risk-averse participant the optimality conditions for the other firms are much simpler then those above. In particular, there are no intertemporal variables such as s_t and w_t for the other firms. To find the market equilibrium position for each company in a period t we simultaneously solve the optimality conditions for each firm. The market demand curve defines the price/generation relationship (dp/dg) which links the optimality equations of each player in the market. Given a market demand curve, contract position and marginal cost function for each firm and values for the two dual variables π_t and λ_t, finding the market equilibrium position is a straightforward process. A method is described in [3], which we have followed. The solution gives the generation quantities for each firm from which the market price can be inferred.

A key step in the above process is determining the optimal dual values π_t and λ_t, for $t = 1, ..., T$. Dual dynamic programming is a technique that has been successfully applied to problems of this type (see [3, 4]). Applied to our problem we begin in period T and discretise the two state variables s_T and w_T over their feasible ranges. For each pair of discrete quantities of s_T and w_T the end horizon conditions define a value for each dual variable. Finding the equilibrium position for these dual variables defines the optimal transition from period T to period $T-1$. After adjusting for the stochastic inflow we have π_{T-1} and λ_{T-1} defined over a range of s_{T-1} and w_{T-1}. Successively applying the state transition process and adjusting for hydrological uncertainty a picture of π_t and λ_t is built up for each period t.

4 End of Horizon Surface

The previous section provided a brief outline of the Stochastic Dual Dynamic Programming process that we have employed. Our method does not significantly differ from standard methods available. This is because, as noted previously, the dynamic programming recursion is independent of the utility function. The utility function only directly affects the dual variables for the final period.

The utility function itself must be carefully defined. As we have defined it $\alpha(w_T, s_T)$ is a function of both accumulated net revenue and storage level. We require $\alpha(.,.)$ to be a separable function of w_T and s_T i.e. the relationship between w_T and s_T cannot be a simple associative one. If such a utility function is used (for example $\alpha(w_T + V(s_T))$, where $V(s_T)$ gives the dollar value of a storage level s_T) then the end of horizon surfaces for π_T and λ_T are given as

$$\lambda_T = \alpha'(w_T + V(s_T)), \qquad \pi_T = \alpha'(w_T + V(s_T))V'(s_T).$$

when these results are substituted into the state transition equations for period T we get

$$\frac{\pi_T}{\lambda_T} = V'(s_T),$$

which is simply the marginal water value curve. Clearly this function is risk neutral as the marginal value at which water is scheduled is independent of the accumulated net revenue.

This behaviour occurs because the accumulated net revenue and value of water in storage hold equal utility value. An incremental dollar of net revenue holds the same utility as an incremental dollar value of storage, as given by the end of horizon water value curve, $V'(s)$. As the model percieves that water may be converted into net revenue at any time, the result is risk-neutral behaviour.

If the utility function is separated into two independent functions as follows

$$\alpha(w_T, s_T) = \alpha_1(w_T) + \alpha_2(V(s_T)),$$

then the outcome will be risk averse for both variables. An useful possibility is that $\alpha_2(x) = x$ i.e the model is risk neutral with respect to the value of stored water but could still be risk averse with respect to the accumulated net revenue. This allows us to isolate the effects of risk aversion on net revenue.

The close relationship between π_t and λ_t leads to an interesting and useful observation. These two variables always appear in the optimality conditions in the ratio $\frac{\pi_t}{\lambda_t}$. Instead of dealing separately with each dual variable, we use the ratio instead. This is useful as it means that we can define a single marginal value surface. We define this end of horizon surface as $\frac{\partial\alpha}{\partial s_T}\frac{\partial w_T}{\partial\alpha}$ which is a function of the ending storage and the total accumulated net revenue. Figure 1 shows an example of a risk-averse end of horizon surface. This surface is risk neutral with respect to the value of water and has a quadratic utility function with respect to the total annual net revenue. The curvature along the storage axis is simply the marginal water value curve $V'(s)$ and the curvature along the net revenue axis is our attitude to risk. This surface is the starting point for our stochastic dual dynamic programming recursion.

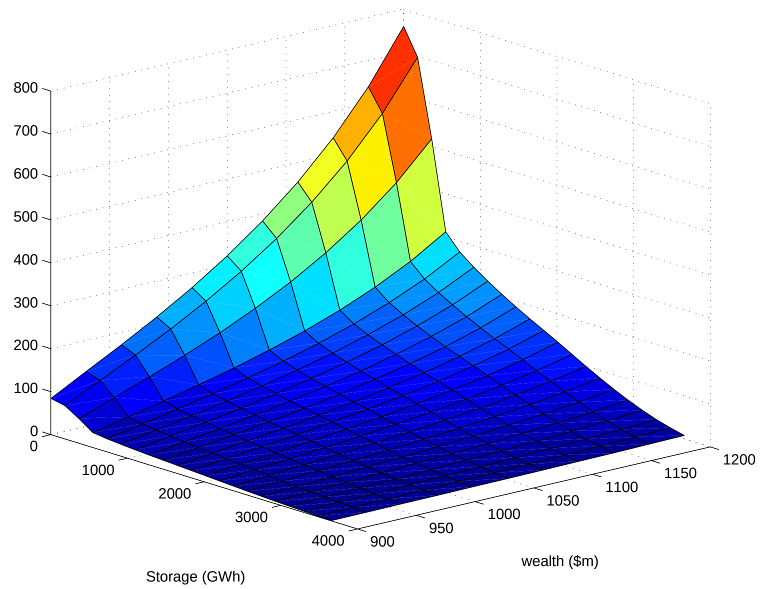

Figure 1: A quadratic end of horizon marginal value surface.

5 Model Results

The model that we have developed has been implemented mainly in MATLAB with a few key routines as C mex files. The model has been run on a number of test scenarios to confirm that it is behaving as would be expected. For a rather simplified model the computational time is approximately 6 minutes on a 200Mhz Pentium Pro.

For the test model we have a large risk-averse participant which owns all significant hydro storage and a large thermal station. A second smaller market participant holds a number of thermal stations at a range of marginal costs, and the third participant holds significant thermal generation at base load prices (i.e. price takers). Each of the market participants holds a contract position and offers into the same market. i.e. there is no account made of the transmission system.

The model has been set up over a year divided into weekly time periods which are further divided into three blocks representing peak, shoulder and off-peak load within the week. Station availablity and a market demand curve are defined for each block which is solved independently of the others.

In the following three sections we outline the model behaviour as we alter aspects of the model.

5.1 Demand Elasticity

Demand elasticity has been shown to have a strong effect on the form of the solution generated by the model. We have assumed that the market demand is described by a constant elasticity demand curve. The price elasticity of demand is defined as

$$\epsilon = \frac{\Delta g}{g} / \frac{\Delta p}{p},$$

so the percentage change of price to generation remains constant. The demand elasticity is negative as an increase in price always leads to a reduction in demand. For highly elastic demand the change in demand for electricity is disproportionally high in comparison to the change in price. A low demand elasticity has the opposite effect, large price changes occur in response to small changes in demand, which means that generators have excessive market power and can withhold generation capacity to force the price to unrealistically high levels. Between these two extremes a typical range for demand elasticity is from -0.3 to -0.8.

To summarise the model behaviour as the demand elasticity decreases

- Average spot market price rises.
- Storage levels rise, with consequently increased chance of spill.
- The smaller firm's generation rises as the larger firm withholds generation.
- The variation of the annual net revenue of the larger participant reduces.
- The effects are more pronounced if the larger participant has a lower hedge contract level.

Accurately estimating the deregulated market demand curve is a difficult task due to lack of market experience and differs according to time frame. In the short term the demand for electricity does not strongly respond to high prices, while in the longer term sustained high prices will lead to:

- Consumers switching to alternative energy sources.
- New competitors entering the market.
- Political intervention.

Hence, although using a low elasticity will realistically model market reactions in the weekly sub-model, the prices obtained are not sustainable in the medium to long term. Therefore a compromise elasticity, somewhere between the long and short run elasticity's might be more appropriate. Estimating such an appropriate value is an area where much more work remains to be done.

5.2 Contract levels

Hedge contracts reduce the exposure of a generator to the spot market price and therefore decreases the incentives to inflate the market price. With no contracts a dominant hydro participant has an strong incentive to withhold generation to force the market price up. As the contracted quantity is increased the incentive to withhold generation is diminished causing the market price to fall. In extreme scenarios when a market participant is very highly contracted the price can fall below the marginal cost of dispatched plant as the incentive to reduce the contract difference payment outweights the loss on generation.

In summary increasing the contracted quantity of the large hydro participant leads to:

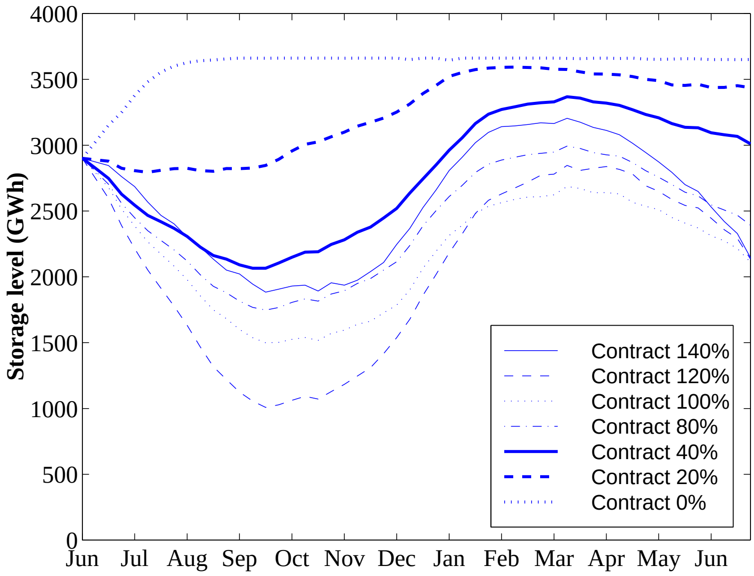

Figure 2: Average annual storage trajectories.

- Average spot market prices falling.
- Generation and net revenue of smaller market participants reduce as the hydro participant increases output.
- Storage levels lowering, although this trend reverses at extremely high contract levels.

At low contract levels the storage trajectories reflects the hydro participants incentive to withhold generation. These trajectories drop as the contract level increases in order to meet the requirement for more generation. This trend reverses at extremely high contract levels. At these levels the severe consequences of a dry year on the net revenue leads to a more conservative storage policy raising the average storage levels. The contract behaviour is illustrated in Figure 2 for a range of hedge contract levels. Here we have defined 100% contract cover to be when the contract quantity is equal to the quantity generated from this firm under a base scenario. In Figure 2 we can see that the lowest storage trajectory is produced by a contract level of 120% increasing this to 140% has resulted in a rise in the average storage levels.

5.3 Risk Aversion

A risk averse firm will adopt a strategy which reduces the probability of an extremely poor net revenue at the cost of reducing the probability of an extremely favourable net revenue. The risk attitude of the hydro participant is defined by the utility function $\alpha(w_T, s_T)$. Because the utility function enscapsulates the distributions of two variables into a single value a careful balance must be found between the final value of water stored and annual net revenue.

We have defined $\alpha(w_T, s_T)$ to be risk neutral with respect to the final value of storage, and quadratic with respect to the annual net revenue. A useful way

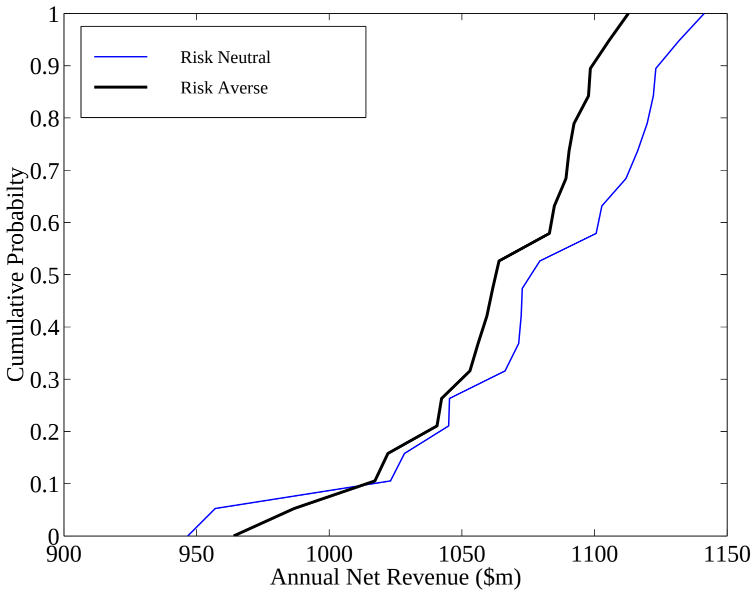

Figure 3: Risk averse and risk neutral cumulative probability curves.

of illustrating risk averse behaviour is on a cumulative probability curve of the annual net revenue. Figure 3 shows two example cumulative probability curves, one generated by our model with a risk neutral utility function, the second having a quadratic risk averse utility function. There is a clear reduction in the variation of the net revenue result, but at some cost to the median result. From a risk management viewpoint this may be the more desirable outcome as the probability of a poor return has been reduced.

The drawback of this behaviour is that the distribution of the final storage levels has widened. The model has traded off variability in the final storage levels against variability in the annual net revenue result. Obviously we can introduce a risk-averse utility function with respect to the final value of water stored, but this may lead to the variability of the net revenue increasing. This observation is important as it leads to the realisation that by focusing on achieving a good net revenue result the hydro participant may leave itself in a unpalatable storage position for the start of the next year. We have observed this in our model where the hydro reservior can be completely emptied in order at the expense of a net revenue result. Understanding the balance between net revenue objectives and final storage objectives is a key area where a greater understanding is required.

6 Further Work

There are a number of areas of possible development for this model. An improvement in the technical detail of the model is necessary. In particular the transmission system can have a significant influence upon the market. In the New Zealand system the HVDC link between the North and South Island has a significant influence upon the spot market. When the HVDC link is constrained there are essentially two separate markets which combine into a single market when the constraint is eased. Current research is investigating the theorectical implications of the link. Another technical enhancement required of this model is that the single national reservior should be divided into two independent reserviors in each island. This would require a number of modifications to the model to incorporate the extra dimension.

The most important point about the model we have described is that it represents a considerable advance over the current models. However it is not sufficent simply to continue with the technical development. The implications of the enhancements need to be throughly tested. Understanding the effects of demand elasticity, contract cover and risk aversion is of equal importance as the technical enhancements. Additionally in light of the continuing reform to the electricity industry in New Zealand, investigation of the effect of increasing the number of players is of increasing importance. Properly tuned the model that we have developed gives us a tool which can used to give insight into the operation of the electricity market.

References

[1] D K Hildebrand and L Ott. *Statistical Thinking for Managers*. Wadsworth. Inc, Belmont, 1991.

[2] A L Kerr, E G Read, and R J Kaye. Risk averse stochastic dynamic programming applied to medium-term reservoir scheduling. Technical Report EMRG-WP-97-02, Department of Management, University of Canterbury, 1997.

[3] T J Scott. *Hydro Reservoir Management for an Electricity Market with Long-term Contracts*. PhD thesis, University of Canterbury, 1997.

[4] M Yang and E G Read. A dual dynamic programming approach to reservoir management with correlated inflows. In *Proceedings of the 26th Conference of the ORSNZ*, pages 21–26. ORSNZ, 1990.