Production inefficiency of electricity markets with hydro generation[*]

Andy Philpott, Ziming Guan, Javad Khazaei, Golbon Zakeri[†]

October 11, 2010

Abstract

Electricity market designs that decentralize decision making for participants can lead to inefficiencies in the presence of nonconvexity or missing markets. This has been shown in the case of unit-commitment problems that can make a decentralized market equilibrium less efficient than a centrally-planned solution. Less attention has been focused on systems with large amounts of hydro-electric generation. We describe the results of an empirical study of the New Zealand wholesale electricity market that attempts to quantify production efficiency losses by comparing market outcomes with a counterfactual central plan.

Keywords: electricity market, hydro-electricity, stochastic inflows, market power.

1   Introduction

In this paper we describe an empirical study of the extent of productive inefficiency of the New Zealand wholesale electricity market (NZEM). The NZEM is a nodal electricity pool market with a high proportion of hydro-electric generation. Unlike markets consisting solely of thermal plant, markets with hydro-electricity have an inter-temporal aspect arising from the fact that energy (water) can be

stored for later delivery. This complicates the decision making of hydro-generators as their computation of the marginal cost of releasing water must involve some modelling of opportunity cost and possible shortage costs. Ideally we would hope that all generators make decisions that are welfare maximizing for the system as a whole. In a system like NZEM that is dispatched every half hour based on generators' offers of energy that reflect their individual views of the future and other agent's actions, a suboptimal outcome is likely, leading to a loss in welfare.

One potential source of inefficiency concerns the dispatch of stations that are connected on the same river system and have a common operator. There are several such river chains in New Zealand. By default each generating station is treated by the market pool as a separate entity to be dispatched according to its offer of energy. This can lead to inefficiencies if the water releases are coupled or constrained by some physical limits that are not revealed to the dispatch process. In other words the operator of a group of stations on the same river chain can extract some efficiency by rearranging their dispatch within a trading period so as to extract the most energy from their stations while meeting river flow constraints.

The inefficiency arising from river-chain constraints is recognized in the New Zealand wholesale market design by including a feature called *block dispatch* for the two main river chains (Waitaki in the South Island) and Waikato (in the North Island), each of which are operated by a single generator[1]. Under block dispatch, offers are made by each generating station on the chain, and these stations are dispatched by the market software as if they were independent. Before implementing the dispatch, however, the owners of the generating stations can rearrange the dispatch amongst their stations on the chain, as long as the total energy delivered is the same as that required by the market[2].

The auction mechanism in the NZEM works in a single trading period, and so there is no allowance within the dispatch for intertemporal constraints (apart from thermal ramping limits). Generators with hydro stations have to choose their offers over a day or longer to respond to the dynamics of their river system as they unfold. Over the longer term they arrange to offer energy to the market based on their individual view of its future stocks of water accounting in some way for the risk of shortage. Although block dispatch affords some degree of flexibility, it is essentially an instantaneous process; the inter-temporal features of river chain operations are not represented in the single-period market clearing mechanism, and so may be a source of inefficiency.

---

[1]This is true at the time of writing. The 2010 Electricity Industry Bill provides provisions for the transfer of two generating stations in the Waitaki system from Meridian Energy to Genesis Energy.

[2]There are also some restrictions on rearranging dispatch between grid nodes with very different spot prices; since the stations are close geographically, in most circumstances the prices are similar enough to make this concern immaterial.

There are several contributions of this paper. The first is to establish a counterfactual model for wholesale electricity markets with large amounts of hydro electricity storage. A number of electricity systems (e.g. those in Brazil and Chile) with significant hydro-electric generation are operated as centrally planned systems. Comparisons of local market performance with that of a centrally planned system in a different country leaves much open to debate. Our work provides a framework to compare the outcomes of two institutional paradigms in the same physical setting. By quantifying the productive inefficiency of a pool market, we can examine the gains in allocative and dynamic efficiency and innovation that we are led to expect from a market, in light of the short-term efficiency sacrifices that are being made to achieve these.

The counterfactual model we describe extends earlier work in this area in several respects. Borenstein *et al* [2] develop a counterfactual model for hydro generation in the California market, some variations of which were applied to the New Zealand market by Wolak [14]. Their model uses a benchmark in which hydro generation is set to its historical levels and thermal plant is then offered in each period at its short-run marginal cost. They argue that the demand-averaged marginal cost computed from this experiment will be biased above the truly competitive marginal cost. Therefore, using this benchmark as a counterfactual will give conservative estimates of market power and productive inefficiency.

We do not agree with these assertions. They are based on an argument that compares different candidate plans for hydro release in the realization of inflows that actually occurred. The perfectly competitive offers of generators on hydro reservoirs do not come from a fixed plan over time - rather they are specified by a feedback *policy* that depends on observations of previous inflows and current reservoir levels. The realization of such a perfectly competitive policy might turn out to be very expensive in hindsight, indeed more expensive than the "optimal" policy computed with perfect information of the inflow sequence. This indicates that the historical release sequence may have low cost in hindsight, and thus produce a low demand-averaged marginal cost that is less than the demand-averaged marginal cost obtained by simulating the perfectly competitive hydro policy over the historical realization.

To illustrate this point with a concrete example, consider a situation of a summer with plentiful inflows and the chance of a dry winter (with very low inflows). If the winter is dry and reservoirs are low then some load must be shed with a high cost. The policy that minimizes expected generation cost is to burn some fuel in the summer to maintain winter reservoir levels in case the winter is dry. The optimal amount of thermal dispatch will depend on the relative costs of shortage and fuel, but it will be nonzero under most sensible choices of these parameters, and electricity prices in the summer i.e. the marginal thermal fuel cost will also be positive. Consider now an historical year in which generators were

3

persuaded that the coming winter would be wet, and so no thermal generation was used in the summer. In the same year, the winter had high inflows, as predicted, and so no thermal generation was needed in the winter. In this setting, the counterfactual of Borenstein *et al* has zero cost (since fixed hydro generation meets demand), which is claimed to be an upper bound on the demand-weighted competitive price in that year. This contradicts our assertion of positive prices in the summer under a minimum expected generation cost policy. To labour the point, the optimal policy *with hindsight* has zero cost, but no policy is perfectly clairvoyant, including the one corresponding to a perfectly competitive market.

In a recent study of the Nordpool, Kauppi and Liski [5] describe a different benchmark that computes a generation policy accounting for uncertainty in future inflows using a stochastic dynamic programming algorithm. This benchmark can be viewed as the policy that would be implemented by a central planner. If we assume that all agents are risk-neutral price takers and share the same probability distribution for future inflows, then the stochastic dynamic programming benchmark solution can be shown to be an equilibrium under perfect competition. This gives good grounds for using it as a model for outcomes in a perfectly competitive market. Our work extends the model of [5] by including several reservoirs and linking the model to a dispatch model within a constrained transmission system. Since there is some risk of electricity shortage in New Zealand in dry winters, our competitive benchmark accounts for shortages using estimated costs of demand reduction.

The second contribution of our approach is that it helps illuminate differences in behaviour in the real system as compared with perfectly competitive risk-neutral behaviour. Bushnell [3] presents a multi-period Cournot model of short-term exercise of market power that predicts the shifting of water from peak periods to off-peak periods, enabling generators to withhold in peak periods without spilling. Our model enables a comparison of observed market behaviour in this setting with a system optimum solution. In the short-term models (with up to a week in duration) we observe some evidence in the New Zealand setting of hydro generators shifting generation out of peaks. The model is not able to determine whether this is due to strategic or risk-averse behaviour.

Owing to limitations on data availability, our results are restricted to the productive efficiency of the NZEM in 2005, 2006, and 2007. Our centrally planned reservoir release policy is less risk averse than the market. It follows a trajectory that appears to be more extreme than the market and uses less thermal generation.

The layout of the paper is as follows. In the next section we describe the New Zealand wholesale electricity market. We then outline the features of a suite of optimization models that we use in our study. To make the computation time reasonable these models are approximations of the system models that are used to

dispatch and price the market. In section 4 we compare the short-term behaviour of a perfectly competitive model in comparison with the observed short-term behaviour of the market. Section 5 deals with optimization over a longer horizon. The final section makes some conclusions.

2 The wholesale electricity market

Since 2004, New Zealand has operated a compulsory pool market, in which the grid owner Transpower plays the role of Independent System Operator (ISO). In this market all generated and consumed electricity is traded[3]. Unlike most electricity markets in other parts of the world, the NZEM has no day-ahead power exchange. Bilateral and other hedge arrangements are still possible, but function as separate financial contracts. Trading develops by bids (purchaser/demand) and offers (generator/supply) for 48 half hour periods (called *trading periods*) over 244 pricing nodes on the national grid. (Although demand side bids are included in the official description of the ISO dispatch model, there is currently very little demand-side bidding in the NZEM, so we will omit them from further discussion.)

The offers of generation made by generators to the ISO take the form of *offer stacks*. These are piecewise constant functions defining the amount of power offered at up to five different prices that may be chosen by the generator m. We can represent the offer stack for generator m by the (step) function $C_m(x)$. In the New Zealand market the generator offer functions C_m are not publicly known at the time of dispatch, but are published two weeks after this time. These data are made available as part of a Centralized Data Set (CDS) distributed by the New Zealand Electricity Commision.

All the prices in the wholesale electricity market in New Zealand are computed by the ISO using a linear programming model called "Schedule Price and Dispatch" or SPD. This represents the New Zealand transmission network by a DC-load flow model. The full version of SPD includes constraints that ensure voltage support, $N - 1$ security for line failures, and meet requirements for spinning reserve that are dispatched at the same time (see [1]).

If we ignore these additional features then the problem solved every trading period by SPD can be described mathematically using the generic network model shown in Figure 1.

For each node i the set $\mathcal{O}(i)$ defines all the generators at node i, where generator m can supply any quantity $q_m \in Q_m$. The demand at node i is denoted D_i.

---

[3]Small generating stations with capacity of 10 MW or less are not required to make offers. From 1996-2004 a voluntary wholesale market existed, where approximately 80% of electricity was traded; the remaining 20% by bilateral contracts.
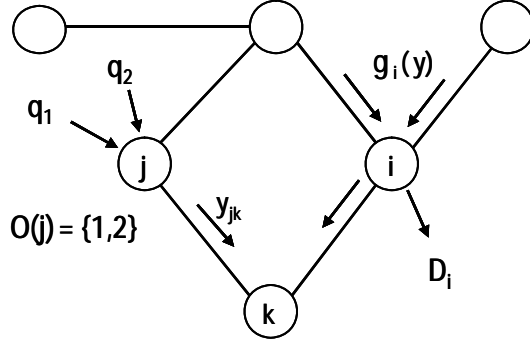
Figure 1: Generic network model illustrating notation

This gives the following market dispatch model:

$$
\begin{array}{rlrl}
\text{MP1:} \quad \text{minimize} & \sum_i \sum_{m \in \mathcal{O}(i)} \int_0^{q_m} C_m(x)dx & & \\
\text{s.t.} & g_i(y) + \sum_{m \in \mathcal{O}(i)} q_m & = & D_i, \quad i \in \mathcal{N}, \\
& q_m \in Q_m, \quad m \in \mathcal{O}(i), & & i \in \mathcal{N}, \\
& y \in Y. & &
\end{array}
$$

Here the components of the vector y measure the flow of power in each transmission line. We denote the flow in the directed line from i to k by y_{ik}, where by convention we assume $i < k$. (A negative value of y_{ik} denotes flow in the direction from k to i.) We require that this vector lies in the convex set Y, which means that each component satisfies the thermal limits on each line, and satisfies loop flow constraints that are required by Kirchhoff's Law. The function $g_i(y)$ defines the amount of power arriving at node i for a given choice of y. This notation enables different loss functions to be modelled. For example, if there are no line losses then we obtain

$$
g_i(y) = \sum_{k<i} y_{ki} - \sum_{k>i} y_{ik}.
$$

With quadratic losses we obtain

$$
g_i(y) = \sum_{k<i} y_{ki} - \sum_{k>i} y_{ik} - \sum_{k<i} \frac{1}{2} r_{ki} y_{ki}^2 - \sum_{k>i} \frac{1}{2} r_{ik} y_{ik}^2.
$$

In our model the quadratic losses are modelled as piecewise linear functions of arc flow which enables MP1 to be solved as a linear program (at least when losses are minimized by the optimal solution).

Bids and offers start 36 hours before the actual trading period. Up to 4 hours (pre-dispatch) before the trading period starts, a forecast price is calculated to

guide participants in the market. From 4 hours to the start of the trading period every half hour a *dispatch price* is calculated (and communicated). Two hours before the start of the trading period, bids and offers for the period in question are locked in. From that point onwards any new prices reflect the ISO's adjustments in load forecasts and system availability.

During the half hour period the ISO publishes a new real-time price every 5 minutes and a time-weighted 30-minute average price. The real-time prices are used by some large direct-connect consumers to adapt their demand. The above prices are a guide only, as the final prices are calculated ex-post (normally noon the following day, unless there are irregularities or disputes) using the offer prices as established 2 hours before the trading period, and volumes metered during the trading period.

3   The models

Our study will make use of a suite of models defined for reference as follows:

CP1:   A determininstic dispatch model solved over one trading period;

CP48:   A determininstic dispatch model solved over one day;

CP336: A determininstic dispatch model solved over one week;

YEAR: A stochastic planning model solved over one year;

INTER: A model that is used for estimating water release and spill over a day or week;

EP: A model that is used to calibrate nodal demand from historical dispatch and prices.

We examine a counterfactual proposal that supposes that the national electricity system is controlled centrally and is dispatched sequentially by the yearly and weekly models. This is compared with the actual dispatch in the wholesale market.

3.1   The dispatch model

Ideally, the dispatch model that should be used in our study is SPD, the full-scale version of MP1 with 244 nodes and reserve and security constraints. However we need to solve such a model many times in simulation, and so we have chosen to approximate this system with a model having only 18 nodes. The representation that we use is shown in Figure 2.

This approximation ignores constraints in the full model that arise from voltage support, $N-1$ security, spinning reserve and frequency keeping, and so it is likely to underestimate prices. Moreover since we aggregate electricity load into
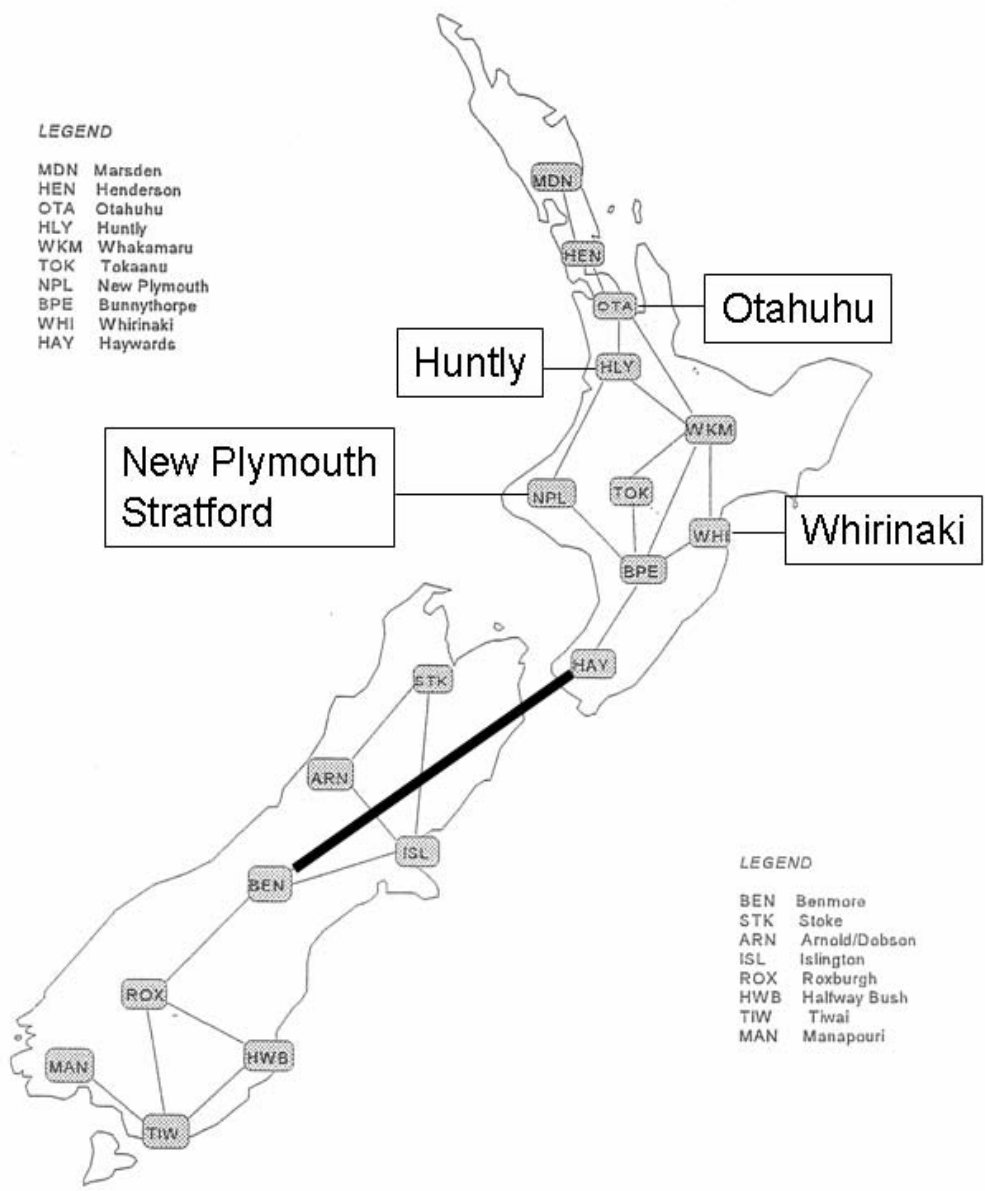
Figure 2: Approximation of New Zealand transmission network showing location of major thermal generators. The bold line represents a HVDC cable connecting the South and North islands.

regions around each of the 18 nodes, we ignore in our approximation the thermal line losses that occur in lines that join the points within each region. This means that regional totals of historical demand will underestimate the true demand that we should use at each of the 18 nodes in the simplified network. One option is to apply a uniform scaling to demand, but this does not reflect the fact that demand is concentrated in some nodes (e.g. TIW which contains a large aluminium smelter) and not in others (e.g. WKM that meets the needs of a dispersed region in our model.) To overcome this, we estimate demand values for each node by solving the following model for each trading period being studied:

$$
\begin{array}{llll}
\text{EP:} & \text{maximize} & \sum_i \bar{\pi}_i g_i(y) \\
& \text{s.t.} & g_i(y) + \sum_{m \in \mathcal{O}(i)} \bar{q}_m & = & D_i, & i \in \mathcal{N}, \\
& & -\alpha \bar{D}_i \leq D_i - \bar{D}_i \leq \alpha \bar{D}_i, & & i \in \mathcal{N}, \\
& & y \in Y.
\end{array}
$$

As in the full-scale dispatch model SPD, the loss functions g_i in EP are modelled as piecewise linear functions of line flow y. Here $\bar{\pi}_i$, $\bar{D}_i$, and $\bar{q}_m$ are respectively the historical nodal price at the node i representing the region, aggregated demand for the region, and aggregated dispatch for the region, all summed from data recorded in the Centralized Data Set. It is known (see e.g. [11]) that the line flows from any given optimal dispatch maximize $\sum_i \bar{\pi}_i g_i(y)$ where $\bar{\pi}_i$ are the nodal prices. The problem EP seeks to scale aggregated demand $\bar{D}_i$ for each node i (by at most $1 \pm \alpha$) so that the demand estimates D_i obtained are consistent with the historical dispatch and historical prices according to EP. Here α is chosen to be 0.1[4].

In a centrally planned market, the offer curves C_m would be determined by the marginal cost of supply. For simplicity we assume that generator m can supply any quantity $q_m \in Q_m$ at cost ϕ_m per MWh. The single-period centrally planned economic dispatch model (CP1) then minimizes cost in a single trading period.

$$
\begin{array}{llll}
\text{CP1:} & \text{minimize} & \sum_i \sum_{m \in \mathcal{O}(i)} \phi_m q_m \\
& \text{s.t.} & g_i(y) + \sum_{m \in \mathcal{O}(i)} q_m & = & D_i, & i \in \mathcal{N}, \\
& & q_m \in Q_m, \quad m \in \mathcal{O}(i), & & i \in \mathcal{N}, \\
& & y \in Y.
\end{array}
$$

3.2 The short-term hydro model

To investigate the dispatch of hydro-electricity over the course of a day a national river-chain dispatch and nodal pricing model (CP48) combines offers with river

---

[4]Details of this estimation process, and all data used in this paper are available in the online companion downloadable from the data repository at http://www.epoc.org.nz.

scheduling constraints over 48 half-hour trading periods, $p = 1, 2, ..., 48$. The physical system for this model is shown in Figure 3.
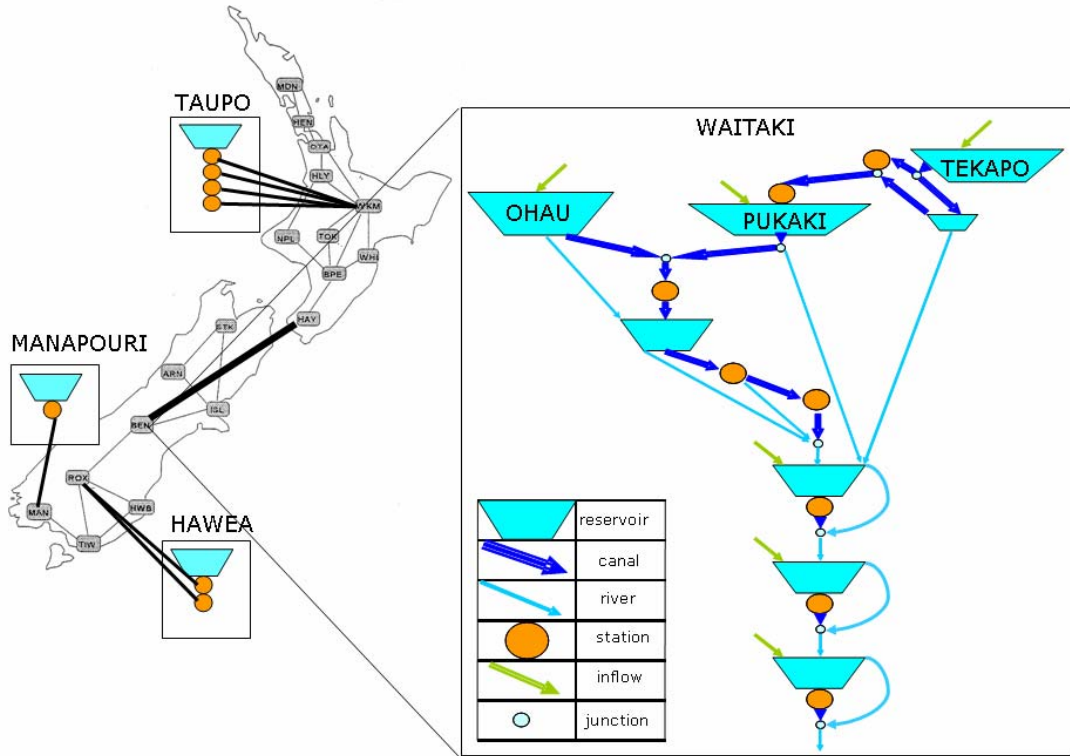

Figure 3: Approximate network representation of New Zealand electricity network showing main hydro-electricity generators

In the model we discriminate between thermal generation f_m, $m \in \mathcal{F}(i) \subseteq \mathcal{O}(i)$, and hydro generation $\gamma_m h_m$, $m \in \mathcal{H}(i) \subseteq \mathcal{O}(i)$. This gives the following

formulation:

CP48: min $\sum_{p=1}^{48} \sum_{i \in \mathcal{N}} \sum_{m \in \mathcal{F}(i)} \phi_m f_m(p)$

s.t. $g_i(y(p)) + \sum_{m \in \mathcal{F}(i)} f_m(p) + \sum_{m \in \mathcal{H}(i)} \gamma_m h_m(p) = D_i(p), \quad i \in \mathcal{N},$

$x(p+1) = x(p) - A(h(p) + s(p)) + \omega(p), \quad p = 1, 2, \ldots, 48,$

$x(1) = \bar{x}(1), \qquad x(49) = \bar{x}(49),$

$0 \leq f_m(p) \leq a_m, \quad m \in \mathcal{F}(i), \, i \in \mathcal{N},$

$0 \leq h_m(p) \leq b_m, \quad 0 \leq s_m(t) \leq c_m, \quad m \in \mathcal{H}(i),$

$0 \leq x_m(p) \leq r_m, \quad m \in \mathcal{H}(i), \, i \in \mathcal{N}, \quad y \in Y.$

Here the water balance constraints are represented by

$$x(p+1) = x(p) - A(h(p) + s(p)) + \omega(p)$$

where $x(p)$ is the reservoir storage at the start of period p, $s(p)$ denotes spill in period p, and $\omega(p)$ is the uncontrolled inflow into the reservoir in period p. All these are subject to capacity constraints. (In some cases we also have minimum flow constraints that are imposed by environmental resource consents.) The parameter γ_m, which varies by generating station m, converts flows of water $h_m(p)$ into electric power.

The node-arc incidence matrix A represents the river-valley network, and aggregates controlled flows that enter a reservoir from upstream and leave a reservoir by spilling or generating electricity. In other words row i of $A(h(p) + s(p))$ gives the total controlled flow into the reservoir (or river junction) represented by row i, this being the sum of any immediately upstream releases and spill minus the release and spill of reservoir i. With small modifications to the water balance constraints, our model can also represent transit times for flows in parts of the river system. These can be as long as 22 periods (11 hours). In this setting the boundary conditions are augmented to account for in-transit flows at the beginning and end of the day.

3.3 The medium-term hydro model

To investigate the dispatch of hydro-electricity over the course of a year, a hydro-thermal release policy must be determined. This involves the solution of a large-scale stochastic dynamic programming model which is defined as follows. Let $x(t)$

denote the reservoir storage at the beginning of week t, and let $C_t(x, \omega(t))$ be the minimum expected fuel cost to meet electricity demand in weeks $t, t+1, \ldots, T$, when reservoir storage $x(t) = x$ and week t's inflow is known to be $\omega(t)$. Here $C_t(x, \omega(t))$ is the optimal solution value of the mathematical program:

$$\mathrm{P}_t(x, \omega): \quad \min \quad \sum_{i \in \mathcal{N}} \sum_{m \in \mathcal{F}(i)} \phi_m f_m(t) + \mathbb{E}[C_{t+1}(x(t+1), \omega(t+1))]$$

$$\text{s.t.} \quad g_i(y(t)) + \sum_{m \in \mathcal{F}(i)} f_m(t) + \sum_{m \in \mathcal{H}(i)} \gamma_m h_m(t) = D_i(t), \quad i \in \mathcal{N},$$

$$x(t+1) = x - A(h(t) + s(t)) + \omega(t),$$

$$0 \leq f_m(t) \leq a_m, \quad m \in \mathcal{F}(i), \ i \in \mathcal{N},$$

$$0 \leq h_m(t) \leq b_m, \quad 0 \leq s_m(t) \leq c_m, \quad m \in \mathcal{H}(i),$$

$$0 \leq x_m(t) \leq r_m, \quad m \in \mathcal{H}(i), \ i \in \mathcal{N}, \quad y \in Y.$$

To solve this we have used the DOASA code [9] which is based on the SDDP technique of Pereira and Pinto [8]. This approximates $\mathbb{E}[C_{t+1}(x(t+1), \omega(t+1))]$ using a piecewise linear outer approximation that is updated using samples of the inflow process. Weekly demand is represented by a load duration curve with three blocks, and $\omega(t)$ is sampled from historical inflow observations. We use a simplified transmission network comprising three nodes: one for the South Island, one for the lower North Island and one for the upper North island. The hydro system assumes that six reservoirs, Manapouri, Hawea, Ohau, Pukaki, Tekapo and Taupo, can store water from week to week. The remaining reservoirs are treated as run-of-river plant with limited intra-week flexibility. Details of the DOASA model for this study can be found at the online companion [10] to this paper.

The solution to $\mathrm{P}_1(x_1, \omega(1))$ defines a set of thermal plants to run and a set of linear functions (or *cuts*) whose pointwise maximum approximates $\mathbb{E}[C_2(x(2), \omega(2))]$. Indeed the DOASA code yields an outer approximation to $\mathbb{E}[C_{t+1}(x(t+1), \omega(t+1))]$ at each stage t, and so this defines a policy at this stage by solving the

single-stage approximating problem:

$$\text{AP}_t(x,\omega): \quad \min \quad \sum_{i\in\mathcal{N}}\sum_{m\in\mathcal{F}(i)} \phi_m f_m(t) + \theta_{t+1}$$

$$\text{s.t.} \quad g_i(y(t)) + \sum_{m\in\mathcal{F}(i)} f_m(t) + \sum_{m\in\mathcal{H}(i)} \gamma_m h_m(t) = D_i(t), \quad i \in \mathcal{N},$$

$$x(t+1) = x - A(h(t) + s(t)) + \omega(t),$$

$$0 \le f_m(t) \le a_m, \quad m \in \mathcal{F}(i), \, i \in \mathcal{N},$$

$$0 \le h_m(t) \le b_m, \quad 0 \le s_m(t) \le c_m, \quad m \in \mathcal{H}(i),$$

$$0 \le x_m(t) \le r_m, \quad m \in \mathcal{H}(i), \, i \in \mathcal{N}, \quad y \in Y,$$

$$\alpha_{t+1}^k + \beta_{t=1}^k x(t+1) \le \theta_{t+1}, \, k \in \mathcal{C}(t+1).$$

3.4 The weekly simulation model

The DOASA model defines a policy by the cuts that give the outer approximation of $\mathbb{E}[C_{t+1}(x(t+1), \omega(t+1))]$. In other words we can represent the expected future cost of meeting demand when $x(t+1)$ remains in the reservoirs at the end of week t by θ_{t+1} where

$$\theta_{t+1} = \max\{\alpha_{t+1}^k + \beta_{t+1}^k x(t+1), \, k \in \mathcal{C}(t+1)\}.$$

The policy is then to release water to meet demand at least cost over the week while accounting for the future cost. Given a set of cuts this policy can be simulated over a given period using the observed sequence of inflows ω to the hydro reservoirs. In each week we solve a model that includes the 18-node transmission system. This involves the solution of the following problem with $P = 336$ trading

periods:

$$\text{SP}_t(x, \omega): \quad \min \quad \sum_{p=1}^{P} \sum_{i \in \mathcal{N}} \sum_{m \in \mathcal{F}(i)} \phi_m f_m(p) + \theta_{t+1}$$

$$\text{s.t.} \quad g_i(y(p)) + \sum_{m \in \mathcal{F}(i)} f_m(p) + \sum_{m \in \mathcal{H}(i)} \gamma_m h_m(p) = D_i(p), \quad i \in \mathcal{N},$$

$$x(p+1) = x(p) - A(h(p) + s(p)) + \omega(p), \; p = 1, 2, ..., P,$$

$$0 \le f_m(p) \le a_m, \quad m \in \mathcal{F}(i), \; i \in \mathcal{N}, \; p = 1, 2, ..., P,$$

$$0 \le h_m(p) \le b_m, \quad 0 \le s_m(p) \le c_m, \quad m \in \mathcal{H}(i), \; p = 1, 2, ..., P,$$

$$0 \le x_m(p) \le r_m, \quad m \in \mathcal{H}(i), \; i \in \mathcal{N}, \quad y \in Y, \; p = 1, 2, ..., P,$$

$$\alpha_{t+1}^k + \beta_{t+1}^k x(P+1) \le \theta_{t+1}, \; k \in \mathcal{C}(t+1).$$

Observe that $\text{SP}_t(x, \omega)$ assumes perfect information about demand and inflow over the week being simulated, but does not anticipate inflows beyond that. The simulation is therefore likely to give an optimistic estimate of the fuel cost from implementing this policy for the week.

4 Experiments

The Centralized Data Set (CDS) maintained by the New Zealand Electricity Commission [6] records the offer curves for every generator in the wholesale market. It also records the historical dispatch level of each generator and the daily reservoir inflows. Given costs per MWh of gas and coal generation it is therefore possible to compute the cost of fuel required to generate the electricity dispatched by the wholesale market in each half hour. This cost can be compared with the same cost as optimized by a central plan.

There are several difficulties with such an approach. The first of these concerns dispatch that has limited control. Examples of such dispatch is that from cogeneration, geothermal plant, run-of-river hydro and wind. Although these have low marginal cost, their availability is subject to the vagaries of inflows and wind, and so we cannot centrally dispatch these in a counterfactual. We choose to fix all cogeneration, geothermal generation wind generation, embedded generation, run-of-river generation and small hydro plant at their historical levels. This leaves the large hydro systems (Manapouri, Clutha, Waitaki and Waikato) available for control along with the major thermal plants (Huntly (4 units plus

e3p and P40), Otahuhu, New Plymouth, Stratford, and Whirinaki). These are the only generators that we allow to offer energy within our model. In reporting all costs, our measure will be the cost of fuel burned by these five plants as evaluated with the table of estimated fuel costs as shown in Table 1 and Table 2 (obtained from [7]). All costs are measured in 2008 dollars. Coal costs are assumed to be constant at \$4/GJ ([4]). The short-run marginal cost for any plant can be obtained by multiplying the heat rate by the fuel cost.

Station	Fuel	Heat rate
		(GJ/MWh)
Huntly	Coal	10.50
e3p	Gas	6.80
P40	Gas	9.50
Stratford	Gas	7.30
New Plymouth	Gas	11.00
Otahuhu	Gas	7.05
Whirinaki	Oil	11.00

Table 1: Heat rates and fuel types of thermal plant in New Zealand

		Gas	Diesel
2005	Mar	4.49	22.78
	Jun	4.21	24.60
	Sep	4.13	26.43
	Dec	5.14	25.68
2006	Mar	5.12	26.73
	Jun	5.07	31.64
	Sep	5.18	30.07
	Dec	5.67	25.23
2007	Mar	6.00	24.07
	Jun	5.97	25.39
	Sep	6.01	25.81
	Dec	5.57	28.76
2008	Mar	4.11	30.66
	Jun	5.13	37.22
	Sep	5.36	37.40
	Dec	5.77	28.07

Table 2: Quarterly real gas and diesel wholesale prices in (2008 NZ)\$/GJ.

We note that the fuel prices here are estimated average values assuming fuel can be purchased on demand. Natural gas is typically acquired under a take-or-pay contract that gives a different operating imperative from that faced by a purchaser with more flexibility. Similarly coal is typically used from a stockpile that is periodically restocked; in this setting, supply shortages can lead to high opportunity costs. We argue, however, that a central planner might avoid many of the contractual problems in obtaining thermal fuel that a number of competing generators might face, which would make our assumption less important. In any

case including these effects leads to a more complicated optimization problem than we would want to study here, so we ignore them.

To enable a fair comparison with the market outcomes, we have de-rated stations at which plant have been removed for planned maintenance, as outlined in the POCP database [12]. The schedule in POCP defines the starting and end time of scheduled maintenance for generators, which includes the offering generators and three others (Tokaanu, Rangipo and Waikaremoana) that we treat as fixed. We have observed that in some declared maintenance periods, the generators still offered or were dispatched energy. For the offering generators, we define the capacity loss due to maintenance to be the nominal capacity minus the maximum of their total offer and their dispatch (which can sometimes exceed the nominal capacity by a small amount) in the period of interest. For the three fixed generators, for which only dispatch data are available, the capacity loss is defined to be the nominal capacity minus their dispatch. Note that since a generator may not offer or dispatch at its maximum available capacity in any period, our model overestimates the capacity loss due to maintenance.

We also assume that transmission outages are known at the time of dispatch. The time periods of HDVC line outage, and the HVDC flow are available in the CDS. We define the HVDC capacity loss in each such trading period to be the nominal capacity minus the HVDC flow. Other line outages can be detected by examining historical nodal prices. For each line, the ratio of historical nodal prices at the ends of the line and the ratio of power sent and received along this line are computed. If the former exceeds the latter, then this line is deemed to be constrained by some contingency. Some care is needed in treating lines in loops, as a contingency in one line can affect price differences around the loop. If this is the case, then the line in the loop with the highest ratio of nodal prices between its endpoints is assumed to be the one with the contingency. The capacity loss is then defined to be the difference between the nominal capacity and the power sent in the EP model.

In order to compare the dispatch of a central plan with that of the market we need to ensure that they both have the same boundary conditions. In other words, a market solution may burn more fuel than a central plan, while leaving all reservoirs with more water in them at the end of the day. In our daily experiments we wish to impose the same boundary conditions on both models in order to compare the efficiency of the dispatch.

The first task in doing this is to estimate what these boundary conditions should be in our models, since they do not represent all details of the river systems involved. For example the CDS contains only daily averages of tributary inflows, many of which vary over the course of a day. Even the averages are sometimes approximations of the true values, and reservoir levels at the beginning and end of a day fail to capture water flows that may be in transit between days. Moreover

the conversion factors γ can vary with reservoir head level, and so nominal values of these might not correspond with the water releases associated with a given dispatch level.

The two largest river chains (Waikato at node WKM and Waitaki at node BEN) are block dispatched by the market. This means that the actual dispatches at nodes WKM and BEN respectively were allocated over the stations on their respective river chains so as to give a desirable hydrological flow pattern. We do not have information on how this reallocation was done. We use the historical dispatches $d_m(p)$ obtained from CDS data and solve the following model (called INTER) to find a block dispatch of each hydro chain i where $\mathcal{H}(i) \neq \emptyset$, that is consistent with $d_m(p)$ while meeting hydrology constraints and boundary conditions.

INTER: min $\quad \sum_{p=1}^{T} \sum_{m \in \mathcal{H}(i)} \lambda s_m(p) + \mu \left\| x(T+1) - \bar{x}(T+1) \right\|_1$

s.t. $\qquad \sum_{m \in \mathcal{H}(i)} u_m(p) = \sum_{m \in \mathcal{H}(i)} d_m(p), \quad i \in \mathcal{N},$

$\qquad (1-\alpha)\gamma_m h_m(p) \leq u_m(p) \leq_m (1+\alpha)\gamma_m h_m(p),$

$\qquad x(p+1) = x(p) - A(u(p) - s(p)) + \omega(p),$

$\qquad x(1) = \bar{x}(1),$

$\qquad 0 \leq f_m(p) \leq a_m, \quad m \in \mathcal{F}(i), i \in \mathcal{N},$

$\qquad 0 \leq h_m(p) \leq b_m, \quad 0 \leq s_m(t) \leq c_m, \quad m \in \mathcal{H}(i),$

$\qquad 0 \leq x_m(p) \leq r_m, \quad m \in \mathcal{H}(i), i \in \mathcal{N}, \quad y \in Y.$

The objective of INTER is to minimize spill of water and absolute deviations from historical end storage levels using penalties λ and μ respectively. To minimize these deviations the conversion factor for each station m is allowed to vary between $(1-\alpha)\gamma_m$ and $(1+\alpha)\gamma_m$ (where $\alpha = 0.1$ is a typical choice and γ_m denotes the nominal conversion factor with average headpond levels). The solution from INTER gives a (possibly new) set of hydrological boundary conditions for which the dispatches computed from MP1 for this day are feasible, while allowing for block dispatch. These final boundary conditions are the ones used in our comparisons of the central plan with the market. Since we have computed a market solution already, it remains to solve the model CP48 with the final boundary conditions imposed, to give a centrally planned solution.

It is important here to make some observations about the use of INTER. If we consider it more important to match historical reservoir levels then we should

penalize deviations from these levels more severely than spill. The difficulty in doing this is that the historical solution might appear to spill large volumes of water to match what might be erroneous daily storage or inflow observations. Recall that the lake inflow values are daily averages, and tributary inflows are estimates, and so variations in historical generation might be accommodating changes in these inflow values that we do not have recorded.

Since spill appears to be relatively rare in practice, our approach is to penalize spill more heavily than matching the end conditions. This means that the market solution in each trading period is not forced to spill past generating stations so as to match an historical end-of-day boundary condition. Instead we replace the boundary condition by a synthetic one, i.e. a set of storage levels that the market would have attained with average inflows, minimal spill and its historical dispatch. As a comparison the central plan is then computed using the synthetic boundary conditions.

4.1 Experiment 1: Daily model

The first experiment we conducted compared a daily dispatch under the market model with a dispatch under a central plan. Consider the results for Monday June 20, 2005. Given the costs of thermal fuel in Table 2, the thermal cost in the central plan here is $1,547,273 as compared to $1,580,918 for the market model, giving an inefficiency estimate of (NZ)$33,645 (or about 2.1%). The key difference in thermal dispatch between the solutions is that the market model dispatches a high-cost thermal plant in New Plymouth that is not dispatched in the central model. The inefficiency estimate does not change by much when computed for other days in this week varying between 1.9% and 2.9% (and averaging 2.3%).

It is interesting to examine the hydro dispatch in each of the two main river systems on June 20. The main South Island system, Waitaki, is shown in Figure 4, and the main North Island system, Waikato (i.e. Lake Taupo), is shown in Figure 5. In both these figures one can see that when compared with a central plan, the hydro generation observed in the market is shifted from the peak periods to off-peak periods. This has at least two interpretations. If we assume that generators have perfect foresight of future demand and inflows over the day then the result is consistent with the observations made by Bushnell in his study [3] of hydro generation in California. In particular, Bushnell demonstrated that in comparison with a perfectly competitive solution, strategic hydro generators have incentives to withhold electricity during peak times when demand is relatively inelastic and generate more in off-peak times.

An alternative view is that the difference in the plots shows the effect of uncertainty. In other words, a perfectly competitive generator might act in a less extreme manner than shown by the CENTRAL plot in Figure 4 and Figure 5 if
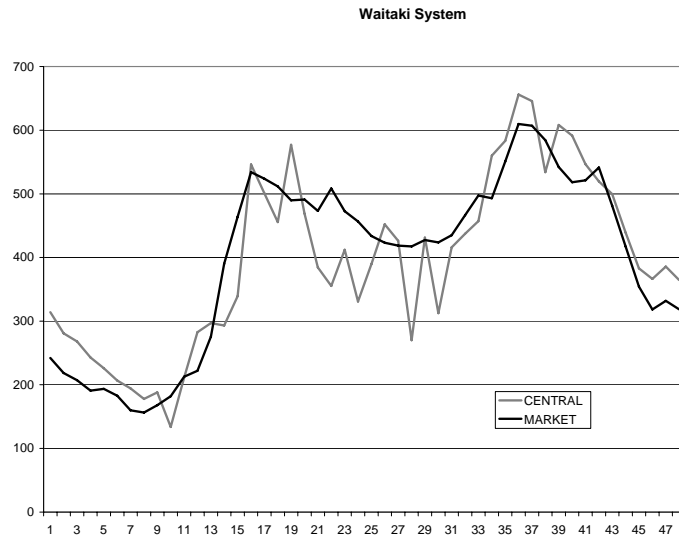
Waitaki System

Figure 4: Total MWh generated in each half hour trading period in Waitaki river system (June 20, 2005).
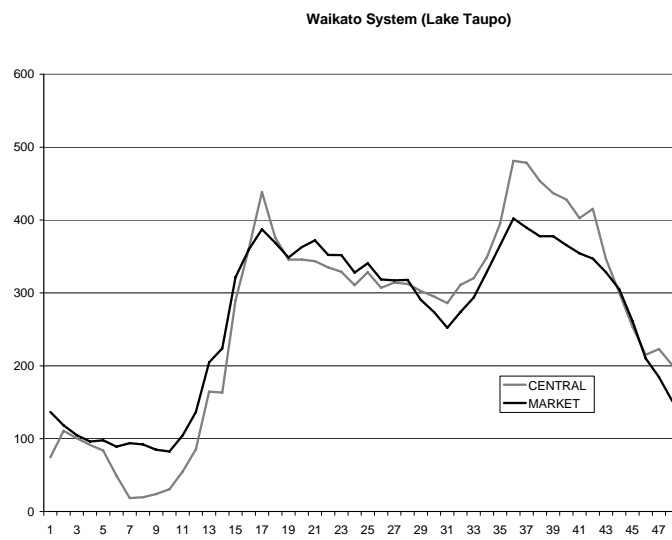
Waikato System (Lake Taupo)

Figure 5: Total MWh generated in each half hour trading period in Waikato river system (June 20, 2005).

there was some uncertainty about future inflows and demand over the course of a day. This interpretation is supported by the less extreme MARKET plot showing the total generated in each half hour trading period.

The output of the central-plan solution shown in Figure 4 and Figure 5 appears to be more variable than that of the market solution, raising the possibility that thermal startup/ramping constraints may be limiting the latter, and accounting for much of the loss in value. However, our tests show that this is not the case, and solving a mixed-integer program for the central plan with unit commitment constraints imposed on the thermal units gives a solution with practically the same value.

To estimate whether the differences in fuel cost between the centrally planned and market solutions for this day were reflected in other days, we repeated the experiment for all days in 2006 and 2007. The total difference in fuel cost for each year is given by Table 3.

YEAR	MARKET	CENTRAL	DIFFERENCE	PERCENTAGE
2005	$451,896,771	$426,720,979	$25,175,792	5.6%
2006	$491,720,636	$471,071,279	$20,649,358	4.2%
2007	$487,837,719	$473,346,254	$14,491,466	3.0%

Table 3: Differences in annual total fuel cost (in $NZ) between the market dispatch (MARKET) and central planning models (CENTRAL) solved each day.

The decline in the difference over the three years is accompanied by a decrease in the number of trading periods that units from the New Plymouth plant are offered into the NZEM from 10,800 in 2005 to 5796 in 2007. Since September 26, 2007 the New Plymouth plant has not offered to the market, and it was decommissioned in December, 2007.

4.2   Experiment 2: Weekly model

In the daily model, the boundary conditions effectively constrain the total hydro generation to be the same under the market assumptions and the central plan. We now examine a central planning model (CP336) of 336 trading periods over one week. Here we compare a central dispatch for all river chains assuming perfect information about demand and inflows in all periods in the week with the cost of the dispatch in the market dispatch. To produce a set of target reservoir levels for the end of the week that both models share, we solve an instance of INTER with 336 trading periods.

We first compare the results for week 25 in 2005, that is June 18-June 24, 2005. These are shown in Table 4, the first part of which shows the result of running CP48 for each day in week 25. The initial reservoir levels for CP48 for

each day are given by historical observations, and the estimated final levels are computed using the INTER model over 48 periods as described above. Observe that the final levels used might not match the initial levels used for the following day, and so the sequence of seven daily solutions is not feasible for the weekly model CP336.

	MARKET	CENTRAL	DIFFERENCE
Saturday, June 18, 2005	$ 1,444,762	$ 1,415,376	$ 29,387
Sunday, June 19, 2005	$ 1,385,803	$ 1,358,982	$ 26,822
Monday, June 20, 2005	$ 1,580,918	$ 1,547,273	$ 33,645
Tuesday, June 21, 2005	$ 1,587,786	$ 1,541,573	$ 46,213
Wednesday, June 22, 2005	$ 1,546,380	$ 1,504,886	$ 41,494
Thursday, June 23, 2005	$ 1,547,246	$ 1,511,118	$ 36,128
Friday, June 24, 2005	$ 1,576,457	$ 1,541,685	$ 34,773
Total	$ 10,669,353	$ 10,420,892	$ 248,461
Weekly model	$ 10,669,353	$ 9,939,920	$ 729,433

Table 4: Weekly fuel costs (in $NZ) for the central plan as compared with the market dispatch. The weekly model has a lower cost as it shifts water between days to make savings.

The last line of Table 4 shows the fuel cost of the optimal solution to CP336, which has boundary conditions imposed at the end of the week only. The central plan can therefore exploit some flexibility in transferring water between days as shown by the difference between the total daily savings $248,461 (2.1%) and the overall savings observed in this week of $729,433, which is 6.8% of market fuel cost (i.e. the fuel cost of stations Huntly, Otahuhu, New Plymouth, Stratford and Whirinaki).

We can examine the dispatch of hydro plants over this week for several generators. For the Waitaki system we obtain the plot shown in Figure 6. A similar plot is shown for the Waikato river in Figure 7. The Waikato system shows some of the same load-shifting behaviour for the market solution as discussed above.

To see if these weekly savings are seen in other weeks, we repeated the computation for each week in 2005, 2006, and 2007, and computed the total weekly cost savings for each year as shown in Table 5.

YEAR	MARKET	CENTRAL	DIFFERENCE	PERCENTAGE
2005	$451,896,771	$407,369,223	$44,527,548	9.9%
2006	$491,720,636	$450,223,740	$41,496,896	8.4%
2007	$487,837,719	$455,319,555	$32,518,164	6.7%

Table 5: Total fuel cost savings (in $NZ) from weekly central plan dispatch compared with market dispatch

In the weekly central plan we observe that it rarely dispatches the New Plymouth plant in 2005, whereas this is often dispatched by the market. For example,
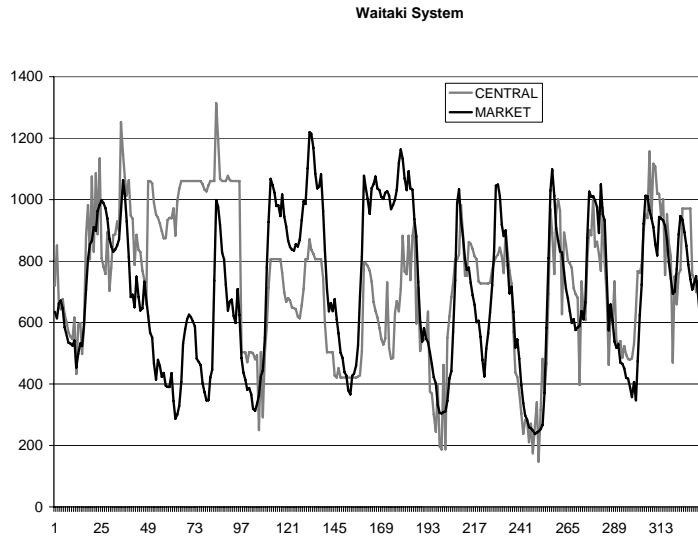
Waitaki System

Figure 6: Total MWh generated in each half hour trading period in Waitaki river system (June 18-24, 2005).
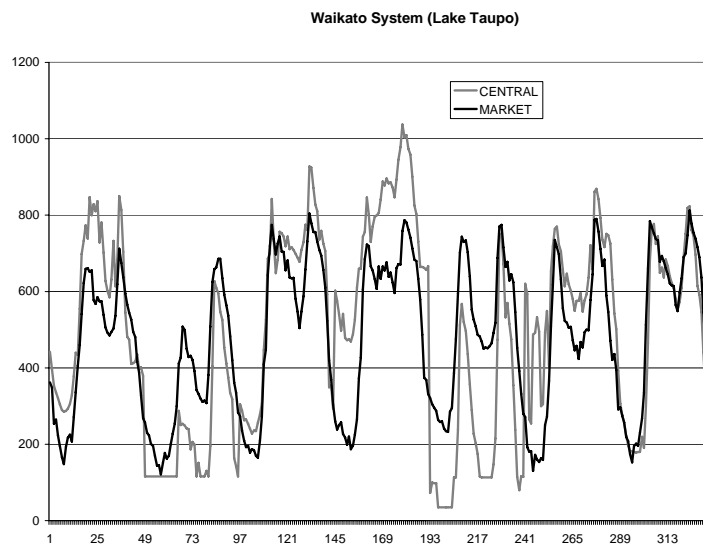
Waikato System (Lake Taupo)

Figure 7: Total MWh generated in each half hour trading period in Waikato river system (June 18-24, 2005).

in week 25, the observed nodal price is often above \$70/MWh, and the historical offers at New Plymouth in this week (apart from periods when it is ramping up to 45 MW) are as shown in Table 6.

Quantity (MW)	45	56	10
Offer price g1 (\$/MWh)	0	65	85
Offer price g2 (\$/MWh)	0	65	85

Table 6: Historical offers of New Plymouth generating units g1 and g2 in week 25, 2005.

This means that in many periods in this week the New Plymouth plant is dispatched between 100 MW and 210 MW at a nodal price above its marginal fuel cost (of \$46/MWh). In contrast, the central plan does not dispatch New Plymouth at all in this week as its fuel cost is more expensive than alternative generation.

The experiment just described requires the storage of water in each reservoir at the end of each week to match an end-of-week target chosen to be close to the historical level. These targets are decided by generators using estimates of their own opportunity costs of generation, which might differ significantly from the system opportunity cost of releasing water from a particular reservoir. The next experiment tests the conjecture that this difference will lead to substantial differences in thermal fuel cost when a central planning policy is tested against the market over a longer planning horizon.

4.3 Experiment 3: Yearly model

In this section we investigate inefficiencies that might arise in the market from generators moving water from week to week in a suboptimal fashion. Studying this requires some care as perfect foresight here can generate large gains for a central plan that would not be realizable in practice. As observed by [5], the appropriate benchmark is the solution to a stochastic dynamic program that seeks to minimize expected fuel cost. As discussed above, we use a sampling-based model with six reservoirs. This gives water-value surfaces that can be used in a weekly dispatch model $\mathrm{SP}_t(x, \omega)$ to determine the optimal dispatch in each trading period of the week to minimize the thermal fuel cost in that week plus the expected future cost of using fuel from the end of the week. Since fuel costs are measured in 2008 dollars we do not discount future costs in our stochastic model.

As discussed above we also make use of costs for unserved load. These depend on the type of customer and the amount of load reduction.

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

Table 7: Load reduction costs ($/MWh) and proportions of each load that is industrial, commercial, and residential load.

The last two columns of Table 7 show the proportion of load of each type in each island. In simulations we assume for simplicity that these proportions are the same at every node of the transmission grid. The costs (in NZ$/MWh) of shedding load are also shown in Table 7. We assume that up to 10% reduction in load can be achieved at a relatively low cost, but the value of unplanned interruption (or reduction above this level) is very high ($10,000/MWh).

The experiment compares the market dispatch with a simulation of a centrally planned dispatch. This is computed using a rolling horizon model. Given initial reservoir levels, the DOASA model is run over a 12 month planning horizon (with a terminal water value of $50/MWh) in every reservoir. We use 35 historical reservoir inflow sequences from 1970 to 2004 as an empirical distribution of inflows, and assume that these are stagewise independent. This means that in its forward pass DOASA samples from a set of 35 inflow vectors, and in its backward pass each cut involves an expectation over 35 discrete outcomes. Before DOASA is run the historical inflow sequences are adjusted using

$$I_t = (\frac{I_0}{h_0})^{\alpha^t} h_t$$

where $\alpha < 1$ is a parameter to be estimated ($\alpha = 0.44$ in our model), h_t is the historical inflow observation, and I_t is the inflow sample we use. This means that every scenario of inflows I_t starts with I_0, and eventually becomes h_t as the term $(\frac{I_0}{h_0})^{\alpha^t} \rightarrow 1$.

The DOASA model minimizes expected fuel cost and so it is risk neutral. This means that shortages might be more frequent than is considered desirable. In practice the Electricity Commission in New Zealand has a mandate to intervene in the electricity market when reservoir levels become very low. This intervention usually involves a public electricity savings campaign. It can be triggered when total reservoir levels fall below what is called the national *minzone*. We have implemented this in DOASA and any violation of the minzone incurs a penalty of $9000/MWh. This means that up to 10% of load reduction (that has penalty costs of at most $4000/MWh) will occur in preference to minzone violation.

We have also added two minzone constraints that reflect reservoir levels in each island. The South Island minzone is 250 GWh less than the national minzone and has the same violation penalty. The North Island minzone is computed to

enable environmental minimum flow constraints at Lake Karapiro to be met with high probability. Violations of this minzone are penalized at a lower value of $500/MWh, as breaches in these flow constraints would typically be allowed in preference to load shedding.

The solution to DOASA yields a set of cuts for each week that define the future cost of meeting demand. Using the cuts from DOASA we simulate a central-plan policy by solving $SP_t(x, \omega)$ for each of 13 weeks, using starting reservoir levels obtained from the end of the previous week. At the end of 13 weeks, we re-solve DOASA with the computed reservoir levels and a new 12 month planning horizon to obtain new cuts. This process then repeats.

The results of our experiment are shown in Figure 8 which compares the historically observed storage with that computed by the simulation of the central plan. Observe that the trajectory of water storage for the central plan in Figure
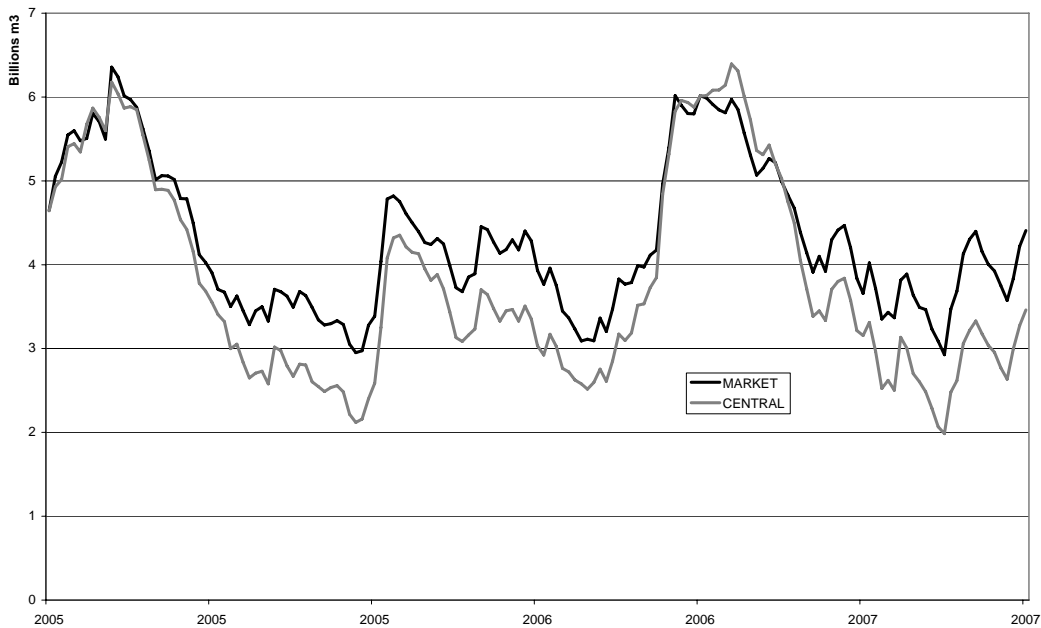

Figure 8: Total New Zealand reservoir storage levels (m^3) from the central planning simulation compared with historical (MARKET) levels over calendar years 2005-2007.

8 is more extreme than the market trajectory. In the winters of 2005, 2006 and 2007 the central plan does not use the more expensive thermal plant as much

25

as the market, and then recovers reservoir storage in the spring by augmenting hydro with base-load coal generation.

YEAR	MARKET	CENTRAL	DIFFERENCE	PERCENTAGE
2005	$451,896,771	$379,400,887	$72,495,884	16.0%
2006	$491,720,636	$425,691,207	$66,029,429	13.4%
2007	$487,837,719	$397,555,859	$90,281,861	18.5%

Table 8: Annual fuel costs (in $NZ) incurred by a central plan that uses DOASA to plan water releases.

Table 8 shows how much extra value is extracted from the DOASA policy. The difference of $90,281,861 in 2007 must be weighed up against the extra 386 GWh of water stored by the market at the end of 2007 as compared with the central plan. Assuming a future value of $50/MWh, this amounts to about $19.3 million, which gives the central plan a lower net savings in 2007 of $70,988,583 or about 14.6%.

5   Discussion

When simulated over 2005, 2006 and 2007, the centrally planned policy incurs less fuel cost than the market. Some of this comes from differences in the merit order arising from generators marking up their short-run marginal costs unevenly in the market (e.g. resulting in the dispatch of New Plymouth rather than less expensive stations). Based on our limited examination of the offer stacks, the hydro plants and thermal plants appear to alter their offer stacks over the day, in contrast to a marginal cost offer that should not vary in the short term.

Such behaviour could be interpreted either as generators bidding strategically, or bidding to meet short-term intertemporal constraints. Generators make offers and are dispatched each half hour, and so the offer stacks they submit in each half hour are the only mechanism that they have to sculpt a varying generation plan to comply with their own constraints over the day. This requirement to "self-dispatch" can be expensive compared with a dispatch optimized over a whole day by the ISO (see e.g. [13] for a discussion of this issue in connection with unit commitment). In this respect, some of the inefficiency could arise as an artifact of the half-hourly market design, rather than reflecting strategic offering behaviour. As we have seen, the daily central planning model integrates most of the short-term intertemporal constraints into a system optimization, and on average saves fuel cost of 5.6% in 2005, 4.2% in 2006, and 3.0% in 2007.

Of course the central planning model does not include all constraints and incentives faced by individual generators. In particular we do not model spinning reserve offers made by generators. In New Zealand these are optimized simultaneously with energy offers (see [1]). This means that a generating unit that

offers reserve might be partially dispatched in the market even though the price at its node exceeds the energy offer price. This effect will lead to some estimation errors, for example in the model EP that does not include reserve, and some upward bias in our savings estimation since the central planning model has fewer constraints than the market faces. Similar observations can be made about security constraints and the provision of frequency keeping services all of which serve to make the actual dispatch deviate from the unconstrained social optimum.

The yearly central planning model saves fuel cost of 16% in 2005, 13.4% in 2006, and 14.6% in 2007, but some of this comes from savings in each day and some from savings made by being able to anticipate inflows within a week and shift hydro generation to avoid spilling or shortages, a clairvoyance that is not enjoyed by the market solution. We can estimate a bound on this estimated value of perfect information (EVPI) by the difference in total weekly savings and total daily savings. This is shown in Table 9 as estimated dollar figures and in Table 10 as percentages.

YEAR	SAVINGS	EVPI	DAILY	ANNUAL
2005	$72,495,884	$19,351,755	$25,175,792	$27,968,336
2006	$66,029,429	$20,847,539	$20,649,358	$24,532,533
2007	$70,988,583	$18,026,699	$14,491,466	$38,470,419

Table 9: Estimate (in $NZ) of the annual value of perfect information (EVPI) accrued in each week of the simulation. Savings in column 1 are optimistic and must be reduced by EVPI values. EVPI is the weekly savings minus the daily savings, and is subtracted from our simulated savings estimates to correct for anticipating inflows.

YEAR	SAVINGS	EVPI	DAILY	ANNUAL
2005	16.0%	4.3%	5.6%	6.2%
2006	13.4%	4.2%	4.2%	5.0%
2007	14.6%	3.7%	3.0%	7.9%

Table 10: Estimated savings broken down as percentages of market fuel cost in each year. Savings in column 1 are optimistic and must be reduced by EVPI values.

The new estimated savings from the simulated central plan correspond to 11.8%, 9.2% and 10.9% of market discretionary fuel cost respectively of which about half is due to a rearranged dispatch in each day and half due to a more efficient use of water from week to week. These savings may appear large, but can be put in the context of a national average annual electricity consumption in 2005, 2006 and 2007 of about 42,000 GWh, which amounts to over $2 billion at $50/MWh.

The centrally planned reservoir release policy minimizes expected costs and so it is risk neutral. In contrast, the market storage trajectory shows some evidence of risk aversion, where generators (who typically have contracts and/or retail customers) appear to avoid the risk of being short of water. Furthermore, the individual market trajectories of all reservoirs are reasonably conservative, while the central plan pools the risk of a shortage (in as much as it seeks to avoid the penalties this imposes) over all of the six reservoirs in the model.

The construction of an optimal policy for water release has been a demanding test of our DOASA code. Although this optimization technique converges almost surely in theory [9], we terminate each run early, and so our optimal central planning policies are approximations, albeit the best available given the state of the art in multi-stage stochastic programming. Including data for 2008 (a notably dry year) and 2009 when these become available through the CDS will test the code even further. It is also possible that our assumption of stagewise independence of inflows leads to suboptimal policies in an environment where low inflows can persist for several weeks. However, despite this assumption the code appears to give good policies, at least in simulation.

We have not tried to estimate the sensitivity of our results to the modelling assumptions used. Our estimates of hydroelectric plant efficiencies and hourly reservoir inflows are made to match historical dispatch. We assume that these estimates remain valid for the central plan, even though it follows a different storage trajectory. This might appear to be contentious, but the main differences between the centrally planned policy and the market are in the levels of the large reservoirs Taupo, Pukaki and Tekapo that have a minimal affect on efficiency for most of the stations on the hydro chains they supply.

We have also not considered the extent or effects of allocative inefficiency in our model, or computed prices that correspond to the central plan. Apart from allowing some load shedding at high prices, we assume an inelastic demand in the short run. (In the simulations of the central plan, no load shedding is observed.) When water is in short supply in New Zealand hydro lakes, wholesale electricity prices peak. In a recent study Wolak [14] identifies extensive exercise of market power in these periods as a source of these price increases, an effect that has also been observed in the Nordic markets [5]. Further work using our model will help to identify the extent that the markups in these prices reflect shortage costs or risk premia as possible alternative explanations.

Finally, this study has made us realize that the backward-looking benchmarks we compute are useful only in a limited sense. Minimizing cost in expectation might prove to be a good policy only in the long run. Our assumption of risk neutrality for constructing a central plan is also open to question when electricity shortages are regarded as serious system flaws by consumers. These issues will motivate improvements to our code in the hope of computing a benchmark policy

that incorporates different levels of risk aversion.

Whether a simulation over three years accurately estimates market productive inefficiency is open to debate with such a small sample. As remarked upon in the introduction, a sequence of years with high inflows will make an ex-ante optimal solution look too conservative. An ideal test of a stochastic optimization model should compare solutions over thousands of out-of-sample simulations. However this is not possible unless one has a believable model of electricity market participant behaviour. In the absence of such a model, our analysis is a step towards some discussion of what-if questions relating to hydro-electric market performance in past years.

References

[1] Alvey, T., Goodwin, D. Ma, X., Sun, D., A security-constrained bid-clearing system for the New Zealand wholesale electricity market, *IEEE Transactions on Power Systems*, Vol. 13, No 2, pp. 340-346, May 1998.

[2] Borenstein, S., Bushnell, J.B., Wolak, F.A., Measuring market inefficiencies in California's restructured wholesale electricity market, *American Economic Review*, 92, 5, 2002, 1376-1405.

[3] Bushnell, J.B., A mixed complementarity model of hydrothermal electricity competition in the western United States, *Operations Research*, 51, 1, 2003, 80-93.

[4] Denne, T., Bond-Smith, S., Hennessy, W., Coal prices in New Zealand markets, COVEC report for New Zealand Ministry of Economic Development, 2009 (downloadable from http://www.med.govt.nz/upload/68784/coal-price-report.pdf).

[5] Kauppi, O., Liski, M., An empirical model of imperfect dynamic competition and application to hydroelectricity storage, HECER discussion paper 232, September, 2008.

[6] New Zealand Electricity Commission database (downloadable from http://www.electricitycommission.govt.nz/opdev/modelling/gpas/May2007/-Generation).

[7] New Zealand Ministry of Economic Development - Energy Data, June 2009, (downloadable from http://www.med.govt.nz/templates/MultipageDocumentTOC___21660.aspx).

[8] Pereira M.V., Pinto L.M., Multi-stage stochastic optimization applied to energy planning, *Mathematical Programming* 52 (1991) 359-375.

[9] Philpott, A.B., Guan, Z., On the convergence of stochastic dual dynamic programming and related methods, *Operations Research Letters* 36 (2008) 450-455.

[10] Philpott, A.B., Guan, Z., Online companion for "Production inefficiency of electricity markets with hydro generation", (downloadable from http://www.epoc.org.nz).

[11] Philpott, A.B., Pritchard, G., Financial transmission rights in convex pool markets, *Operations Research Letters* 32 (2004) 109-113.

[12] POCP database of planned outages (downloadable from http://pocp.redspider.co.nz/).

[13] Sioshansi, R., Oren, S., O'Neill, R., The cost of anarchy in self-commitment-based electricity markets, in *Competitive Electricity Markets: Design, Implementation, Performance*, Sioshansi, F.P. (ed), Elsevier, Oxford, 2008.

[14] Wolak, F.A., An assessment of the performance of the New Zealand wholesale electricity market, Report for New Zealand Commerce Commission, May 2009.