

Consultation

on

MARKET MONITORING REVIEW
OF STRUCTURE, CONDUCT AND PERFORMANCE IN
THE WHOLESALE ELECTRICITY MARKET

Submission by Electric Power Optimization Centre

The University of Auckland

This document was prepared by
Professor Andy Philpott and Dr Anthony Downward
(with review by Dr Geoffrey Pritchard)

<https://www.epoc.org.nz>

December 22, 2021

Executive Summary

1. Confidence in the New Zealand Wholesale Electricity Market is affected by strategic offering by generators. The workable competition model that underlies the New Zealand wholesale market requires careful oversight to ensure that competition is seen to be happening. The Electricity Authority (EA) is investigating ways in which this oversight can be put into effect.
2. Researchers in the Electric Power Optimization Centre at the University of Auckland (EPOC) have been conducting studies on levels of competition in the New Zealand wholesale electricity market for 15 years. EPOC advocates the use of perfect competition benchmark models to assess levels of competition in the wholesale electricity market. We welcome the opportunity to comment on
 - a. the structure, conduct and performance approach to assessing competition in the market;
 - b. the indicators the EA has used under this approach and whether the EA have left out any important indicators;as well as a number of other issues we think the EA should consider.
3. Studies of market competition typically focus on identifying exercise of market power by generators. EPOC's models attempt to identify circumstances where market outcomes differ markedly from those expected from perfect competition. These can result from a number of sources:
 - a. Exercise of market power to increase revenue;
 - b. Acting to avoid risks when these cannot be hedged;
 - c. Responding to signals that are produced by imperfect auction mechanisms.
4. EPOC believes that identifying these situations in historical data provides some insights into how to improve the New Zealand wholesale electricity market design to reduce the differences between observed outcomes and perfectly competitive outcomes.
5. The DOASA package used by the EA in their study was developed by EPOC researchers. Although EPOC provide advice to the EA on the capabilities and operation of DOASA, the study was conducted independently of EPOC, using data collected and compiled by EA researchers.
6. EPOC has conducted past benchmark studies of the New Zealand Wholesale Electricity Market using the DOASA code. The most recent comprehensive study was of 2017. Results of this study are published on the EPOC website (see [9]).
7. The DOASA executable used in the EA study has been upgraded to an open-source Julia package called JADE. JADE leverages the JuMP modelling

language and Oscar Dowson's open-source Julia package SDDP.jl [1], which originated out of the DOASA implementation of stochastic dual dynamic programming, but now has a large international user community. The EA are transitioning their hydro modelling to use the JADE package.

8. JADE has extended the original DOASA model to enable several important features of the New Zealand system to be included in the model. These are:
 - a. Arbitrary numbers of load blocks;
 - b. Risk aversion;
 - c. Finer resolution of transmission constraints;
 - d. Steady-state water valuation;
 - e. Bespoke release constraints on some reservoirs
9. This submission describes the results of some experiments with JADE to illustrate the approach to market monitoring advocated by EPOC. Our approach is as follows:
 - a. Estimate end-of horizon water values for New Zealand reservoirs using JADE in infinite-horizon discounted-cost mode. We show how these water values vary with levels of social risk aversion and discount rate. They are also compared with the end-of horizon water values used by the EA in running DOASA.
 - b. Compute optimal release policies using JADE applied to 2020 with varying levels of social risk aversion to give perfectly competitive counterfactual outcomes.
 - c. Compare observed price levels in 2020 with counterfactual outcomes.
10. The submission also illustrates the application of the EPOC river-chain model HydrovSPD to a sample day from 2020. This yields a dispatch solution in which generation from stored water is offered at its marginal water value throughout the day. We compare this with historical dispatch and price outcomes. A more detailed discussion of HydrovSPD is available in [9].
11. Our submission is aimed at presenting the elements of a rigorous process for market monitoring. We do not propose market design solutions that might reduce differences between market and counterfactual outcomes. Any such improvements would require extensive computational testing with counterfactual models before adoption.

Section 1: Introduction

This document is a submission by The Electric Power Optimization Centre (EPOC) at the University of Auckland in response to a Market Monitoring Review published by the Electricity Authority [2].

EPOC is an independent research group and our response to this consultation document is made independently of any electricity market participant, the Electricity Authority, or the system operator. Our submission is focused on the methodology for market monitoring, and the conclusions that regulators can reasonably draw from market monitoring studies.

EPOC takes the view that wholesale electricity markets work best when they are as close to perfectly competitive as possible. Perfect competition in markets is often claimed to be an unrealistic standard by which to judge wholesale electricity markets, to be replaced by a standard of “workable” competition. The latter standard unfortunately is difficult to measure or assess and is open to interpretation. Perfect competition, although arguably unattainable in practice, is a computable benchmark against which outcomes resulting from market participant behaviour can be measured.

Wholesale electricity pool markets are highly structured market environments that are amenable to detailed analysis. The New Zealand market is unique in allowing public access to offer and bid data as well as access to dispatch and pricing software (vSPD) that can be used to investigate market outcomes against counterfactuals.

The document is laid out as follows. In the next section we describe what we see as the role of the Electricity Authority in monitoring the performance of the New Zealand wholesale electricity market. In Section 3 we describe conditions under which the outcomes of a central planning solution can be used as a perfectly competitive benchmark for measuring outcomes in the wholesale market. Section 4 discusses some additional factors that must be accounted for when considering benchmarks and proposes some approaches to handling these. In Section 5 we present the new features of JADE that enhance some aspects of DOASA. Section 6 provides an analysis of 2020 to illustrate the features of JADE. In Section 7 we present some results from HydrovSPD, and in Section 8 we conclude the submission with a general discussion.

Section 2: The New Zealand wholesale electricity market

The Electricity Authority provides regulation of the New Zealand electricity market with the aim of maximizing the long-term benefit of consumers. Consumers accrue value from using electricity that exceeds what they pay for it. The payment for electricity must be sufficient to cover the costs of building and operating the generation, transmission and distribution facilities that provide and transport the electricity.

The generation costs of providing electricity can be divided into variable costs that depend on the level of generation, and fixed costs that do not depend on this. Examples of variable costs are the cost of fuel (if the plant uses gas or coal) or operation and maintenance costs that depend on output (such as wear and tear). Fixed costs cover everything else including salaries and costs of capital.

Given fixed costs that are proportional to generation capacity and variable costs that are proportional to output, it is possible (in principle) to determine an optimal mix of dispatchable generation capacity to meet any given load profile. The simplest form of such an analysis uses a screening curve analysis applied to load duration curve to assign capacities to peaking plant, off-peak plant and baseload generation.

Uncertainty and intermittency complicate the screening curve approach: a simple graphical picture is insufficient. However, the optimal capacity mix can be determined using stochastic optimization models (such as EPOC's EMERALD model [3]). In other words, to maximize the long-term benefit of consumers, one should forecast demand, or a stochastic model of this, and solve a stochastic mixed integer program to determine the capacity expansion plan and how to operate this.

There are several perceived problems with this planning approach that we will not discuss in detail here. Planning models rely on forecasts of the future (or at least probability distributions governing future parameters). When there is no consensus on these, there is some doubt that government officials will choose the best numbers. The incentive to make the correct decisions is attenuated in comparison with a market of investors "putting their money where their mouths are".

The solution to this is to create a market for investors in generation capacity. Investors earn operating surpluses when the prices they are paid for energy exceed their short-run costs. These surpluses then incentivise investment in capacity. If these markets satisfy certain mathematical assumptions, namely they are convex, perfectly competitive and complete, then there will exist a competitive equilibrium that yields the optimal level of investments that maximizes the long-term benefit of consumers.

The New Zealand electricity market is based around a linear programming model (SPD) which (apart from some switching variables on discrete choices like direction of transmission flow) treats dispatch as a convex optimization problem. The remaining assumptions of perfect competition and completeness are not satisfied.

The key question facing the Electricity Authority is then “Can we use perfectly competitive complete market benchmarks to identify potential inefficiencies in the wholesale market?” EPOC researchers have been studying this question for a number of years and have developed models that represent perfectly competitive benchmarks under various assumptions. We argue that models of this form are an important ingredient in assessing levels of competition in markets that have hydroelectric storage.

This document describes the models, and how results should be interpreted. The use of the models is illustrated by applying them to the calendar year 2020.

Section 3: A perfectly competitive benchmark

In this section of our submission, we argue for the Electricity Authority monitoring the wholesale market using prices computed from a perfectly competitive benchmark model.

3.1 Perfect competition

Perfect competition is a mathematical abstraction. In a deregulated setting, it requires sufficiently many small generator agents so that none behaves as if they can affect the price by their actions. The levels of market concentration can be measured by indices such as the Herfindahl-Hirschman Index as discussed in the consultation document. Breaking large entities into smaller entities will improve these indices, but in a small market like New Zealand, with economies of scale in generation technologies, there is a limit to how much competition can be improved by these means. So perfect competition can be argued to be an unattainable ideal.

On the other hand, the New Zealand wholesale electricity market is organized around a very specific auction format that determines economic dispatch and prices. In this setting, perfect competition does not require infinitely many small economically rational agents but can be approached by the imposition of rules that require generators to behave as if they could not influence prices by their actions. Many markets for wholesale electricity (e.g. PJM in the United States¹) require generators to offer at their short-run marginal cost, i.e., the optimal offer curve for a generator that acts perfectly competitively.

Determining short-run marginal cost for generators with hydro storage is more complicated. Their short-run cost is the so-called *marginal water value* (MWV). Marginal water values are an opportunity cost, inasmuch as they depend on future circumstances. In this sense they are a (risk-adjusted) expectation of future prices for wholesale electricity which are uncertain, so they should more accurately be termed *expected risk-adjusted* marginal water values.

In some countries, e.g., Brazil, these expectations are computed by a central planner that decides on a probabilistic model for reservoir inflows and solves a large-scale dynamic programming model. The Brazilian version of this (called NEWAVE) implements a version of the stochastic dual dynamic programming algorithm (SDDP). In New Zealand the expectations are formed from the views of individual generators that might differ from each other. Even if the generators share the same probabilistic model of future inflows, they might have different assessments of the likelihood of dry winters that will have an effect on their

¹ PJM is not an energy-only market, and generators are paid a capacity payment to augment revenue from the energy market, so the comparison with New Zealand might be thought to be tenuous. We give the example merely to show that monitoring offers in deregulated electricity markets can be done relatively successfully.

assessment of marginal water values. Competition is supposed to provide some discipline on the estimation of water values, so that they approximate some rational consensus of the future. To what extent this discipline is successful in delivering prices that benefit consumers in the long run is the subject of the current consultation.

This submission does not advocate that New Zealand adopts a system like the Brazilian model, where prices are determined by a centrally administered stochastic dynamic program. This would lead to a number of difficulties that arise from such centrally-planned prices. The incentives to arrive at correct prices are absent from decision makers that have no money at stake. Moreover, the model used must be accepted by all parties, and possible litigation might occur over “incorrect” prices if an improved model were developed that delivered different prices as an output.

Although the Brazilian planning approach has its disadvantages, it points toward a way of measuring competition, so we advocate the use of dynamic programming software models as a market monitoring tool. They can be used to compute perfectly competitive benchmark market outcomes under explicitly stated assumptions about parameters. A common framework can then be used to debate assumptions about these parameters rather than relying on an article of faith that workable competition will deliver the right market outcomes.

3.2 Completeness

It might be argued that a central planning solution for the operation of an electricity system is essentially different from a perfectly competitive equilibrium where agents maximize their own profits. The relationship between them is expressed in terms of “welfare theorems”. These essentially state that perfect competition and central planning coincide when markets are convex and complete.

As mentioned in Section 2, the New Zealand electricity market essentially treats dispatch as a convex optimization problem. To ensure a correspondence with a central planning solution, the markets must also be complete, meaning that there are enough instruments traded in the market to enable prices to be formed for all commodities of interest.

The use of water in the Waitaki hydro catchment is a case in point. Water in Lake Tekapo is released for generation by Genesis. Releases that flow into Lake Pukaki provide potentially valuable stored energy for Meridian. A complete market would price this release of water from Tekapo to Pukaki to give transfer payments from Meridian to Genesis. In the absence of such a price or similar mechanism² the generation of Genesis and Meridian can be inefficient.

In the presence of uncertainty and risk aversion, completeness requires that every future random outcome can be hedged by some traded instrument. This is not true in practice simply by virtue of the uncountably many future outcomes that are

² In fact the transfer is governed by alternative contractual arrangements between the companies that are designed to improve the efficiency of the disaggregated river chain.

possible. However, it has been shown, at least in computational experiments [5], that the availability of a few risk-trading instruments such as contracts for difference can make the gap from incompleteness quite small. If the market for risk is complete and agents are endowed with coherent risk measures, then perfectly competitive partial equilibrium coincides with an optimal solution to a risk-averse social planning problem. The mathematics that proves these results can be found in [4], and they are discussed in the context of hydro storage and inflow risk in [7].

The welfare theorems established in [4] give us confidence to be able to integrate the risks of electricity market participants that use (possibly) different coherent risk measures into a social risk measure to use in a central planning solution. This enables the power of EPOC optimization software such as JADE to be applied to construct perfectly competitive risk-averse equilibria that can serve as suitable benchmarks to monitor observed market outcomes.

3.3 Repeated bidding

A potential impediment to competitive outcomes is the structure of the auction design. The wholesale market in New Zealand allows generators to adjust the prices and quantities of offers up until gate closure. The offers submitted are used by the system operator to publish a provisional dispatch and prices based on forecast demand. What results is a sequence of pre-dispatch outcomes that converges to a final dispatch and a real-time price³.

If generators were to offer a fixed supply function in each period, without the opportunity to revise this as demand changes, then the optimal offer to make under perfect competition would be at short-run marginal cost. Then as demand forecasts converge to their realized final values, pre-dispatch prices will converge to the perfectly competitive prices that correspond to these final loads.

In practice, generator offers can be revised after every pre-dispatch with forecast loads. Although a rational perfectly competitive generator would adhere to a short-run marginal cost offer throughout this process, it is natural for traders of energy to focus the tranches of energy offered around the predicted dispatch price. This is a form of bounded rationality that can lead to a sequence of pre-dispatch prices that do not converge on the final prices that would be obtained from perfectly competitive offers. The mathematics underlying this process is discussed in [10], where it is shown using a simple model that the discovery of prices through pre-dispatch can lead to inefficient outcomes.

It is difficult to say if the imperfection in pre-dispatch described above warrants a redesign of the New Zealand auction mechanism. It has always been taken as an article of faith in the New Zealand wholesale market that more frequent auctions lead to better outcomes, as they enable price discovery. This is certainly the case in auction settings where purchasers' valuations are private information, and the seller wishes to extract (and capture) as much of this value as she can. In an

³ We note that as renewable generation capacity increases, making accurate forecasts of demand will become more difficult, so the issues discussed here are likely to become more important.

electricity market that is regulated to maximize the long-term benefit of consumers, more frequent pre-dispatch auctions might decrease this benefit.

We advocate for an examination of this phenomenon by the Electricity Authority using pre-dispatch demand forecasts and historical offers made in response to these.

Section 4: River chains, transmission, contracting and portfolio effects

Generators in the wholesale electricity market take account of many factors when submitting energy offers, which might affect market outcomes. Such factors are often brought up as reasons for offering energy at prices that differ from a modelled short run marginal cost. We discuss each of these in turn.

4.1 River chains

Offer prices for generators with stored hydroelectricity vary across the trading day. These prices are used by generators to control the dispatch to levels that are feasible for river chain operation. Block dispatch provides some freedom in this respect in terms of allocating generation amongst different stations on a river chain. The overall quantity of dispatch for the chain is controlled by pricing the last quantities of available generation capacity at high prices.

Our JADE model has a weekly resolution (with several load blocks in each week) and so does not represent the intra-day river-chain optimization problem in the hour-by-hour detail required to ensure all the flow balance and environmental constraints are followed. In Section 7 we describe the HydrovSPD model that takes account of these.

4.2 Transmission

Gentailers in the wholesale electricity market typically have retail positions in different locations from their generation assets. For example, Contact Energy has some customers in Auckland and some generation assets in Otago. Transmission congestion between generators and loads produces price separation. Generators are paid less at their grid injection node than they must pay at the grid exit node serving their retail load. There is an incentive to offer in such a way as to reduce this price separation.

Evidence for this behaviour is widespread (see e.g. [9]) and it does not appear to be explicitly precluded by market rules. However, it would not occur in a perfectly competitive setting.

4.3 Contracting and portfolio effects

Generators in the wholesale electricity market hold contract positions with energy purchasers. Gentailers also purchase from the spot market to satisfy demand from retail customers. Both contracts and retail positions provide incentives to try and make wholesale prices different than they would be otherwise.

It is often claimed that contracts or retail load change the optimal offer to make even in a perfectly competitive setting, especially if generators are risk averse. The argument is as follows. If gentailer X is purchasing Q at wholesale price p and is generating Q at price p (assuming the same location) then this nets out and so wholesale price risk is removed. To ensure a generation level of Q to meet its retail load, it is deemed rational for X to offer the first Q units at $\$0/\text{MWh}$.

However, this might be suboptimal; the optimal offer under perfect competition in fact remains one at marginal cost. It is possible to prove this assertion using arguments from supply-function equilibrium (see [12]), but a little thought will convince the reader that X offering Q at \$0/MWh might be suboptimal.

Suppose demand is Q and a cheaper competing plant Y was offering at a (truthful) positive marginal cost that was less than the lowest marginal cost of X. If X offers at \$0/MWh then the energy Q dispatched to plant X instead of Y will incur a higher than optimal cost. In a competitive setting, X would offer at its marginal cost, dispatch Q would be to Y, and the price would be set at the marginal cost of Y at quantity Q. X would then buy at this price to meet its contract, which is cheaper than incurring costs by generating.

In summary, a gentailer acting in perfect competition would offer energy at its marginal cost irrespective of its contract or retail position. This can be shown to be true even if the gentailer is risk averse. Offering below marginal cost, ostensibly to reduce portfolio risk, is suboptimal in a perfectly competitive setting.

In contrast, a gentailer making a strategic offer would offer a supply function that lies below short-run marginal cost for $q < Q$ and above short-run marginal cost for $q > Q$ (see [12]). It is easy to find instances in the historical offer stacks of New Zealand generators that lie below marginal cost for low dispatch quantities. Some of these can be explained by constraints e.g., due to unit commitment. Other instances are evidence of offers that would not be optimal in a perfectly competitive setting.

Section 5: JADE

The DOASA executable used in the EA study has been upgraded to an open-source Julia package called JADE. JADE leverages the JuMP modelling language to represent the New Zealand electricity system, and Oscar Dowson's open-source Julia package SDDP.jl [1], which has a large international user community. Given the same inputs, JADE reproduces the same policies as DOASA, but has many more features that make it more suitable for market monitoring. The EA are transitioning their hydro modelling capability to use the JADE package⁴.

The new features are modelled in JADE using specific choices of parameters. The intention in this part of our submission is to show how these are modelled and illustrate the sensitivity of the outcomes to some choices of these parameters. The relevant parts of the model are laid out as follows⁵.

5.1 Transmission modelling

Although it can be expanded by providing more data, the DOASA code used by the EA represents the New Zealand transmission system by three nodes (NI, HAY, SI) and generators and loads are aggregated at these three locations. Transmission is approximated using a line joining the Haywards node (HAY) to the rest of the North island (NI) and another to the South Island (SI).

The JADE model currently has estimated transfer capacities and loss functions for an 11-node transmission network. This enables the capacity of transfers between e.g., Otago and Canterbury to be modelled. A more detailed transmission network in JADE enables a more accurate evaluation of water values that differ by location. As observed by [11], this results in lower cost policies when simulated in the full transmission network.

5.2 Demand modelling

The DOASA code uses three load blocks (PEAK, SHOULDER, OFFPEAK) to represent the load duration curve in any given week. The estimation of these load blocks using vSPD data is outlined in [9]. In JADE, the number of load blocks is arbitrary. We have developed a model with 8 load blocks in each week. The number of hours in each load block is determined using a separate EPOC optimization tool that yields blocks with a best fit to historical demand values. The increase in the number of blocks enables a more accurate representation of intermittency from wind generation that affects the shape of the weekly residual load duration curve. This change of shape and their fewer hours mean that the peak load blocks in the 8-node model have higher loads than those in the three-node model.

⁴ This is discussed in a presentation to the 2021 EPOC Winter Workshop by Dr Phil Bishop (see <https://www.epoc.org.nz/ww2021.html>).

⁵ More details of JADE can be found on www.epoc.org.nz

5.3 Demand response

A key ingredient in social planning models is modelling how demand responds to high prices. DOASA incorporates a load-shedding module that reduces some fraction of load at different locations at different price thresholds. This is much more detailed in JADE enabling specific demand response features (e.g., from the Tiwai point smelter) to be modelled.

5.4 Risk aversion

The main risk in a system such as the NZEM is that of a dry winter leading to energy shortages. This is the risk that is currently modelled in JADE using nested coherent risk measures. The approach to modelling this in SDDP is described in [8]. This is implemented in SDDP.jl which is the SDDP package used by JADE.

It is possible to model other risks in JADE by admitting uncertainty in demand or fuel costs which are currently assumed in the current version to be deterministic.

5.5 Boundary conditions

DOASA solves a stochastic optimization problem over a fixed time-horizon. It requires an end-of-horizon marginal water value surface to be defined. JADE can be applied in infinite-horizon discounted mode (a feature available in SDDP.jl). This produces a steady-state water value surface for each week of the year assuming that future costs are discounted at a given annual rate. These can then be used as end-of-horizon water values for a finite-horizon model.

Section 6: Illustrative examples of the use of JADE

We use this section to illustrate the approach that we propose for benchmarking competition in the wholesale electricity market. First, we present how we have used JADE to model a perfectly competitive benchmark, then we will detail the specific inputs for of JADE that we have used, and finally we will discuss the results.

6.1 Methodology

Given a model such as JADE, how should this be applied? Our approach is as follows:

- (1) We first train JADE in infinite-horizon mode for discount factor 0.925 (corresponding to an 8.1% discount rate) and a choice of different levels of risk aversion. This process gives end-of horizon water values for any week of the year.
- (2) Using parameters and starting reservoir levels for 2020, we train JADE over a 12-month horizon using the appropriate steady-state water values at the end. This gives an optimal water release policy for this period.
- (3) We then simulate the policy from (2) over 89 historical inflow sequences at the risk level chosen in (2).
- (4) For each level of risk aversion, the simulation can be compared with historical outcomes e.g., reservoir levels, prices, and marginal water values.

6.2 Calibration

There are a large number of inputs that need to be specified for JADE. Each of these needs to be calibrated to the conditions in the electricity sector for the period of analysis. Our implementation is JADE is calibrated for 2020.

The EA has posted its DOASA analysis online within the EA EMI site⁶. We have used the EA's input files as a basis for our calibration of the JADE model. Below we discuss the changes and enhancements that we have made.

Network Representation

We model New Zealand as a 11-node radial network with inter-regional transfer capacities that have been calibrated from observed flows. The nodes that we have used are shown in Figure 1.

⁶ The data are available at <https://emi.ea.govt.nz/Wholesale/Datasets/AdditionalInformation/SupportingInformationAndAnalysis/2021/>

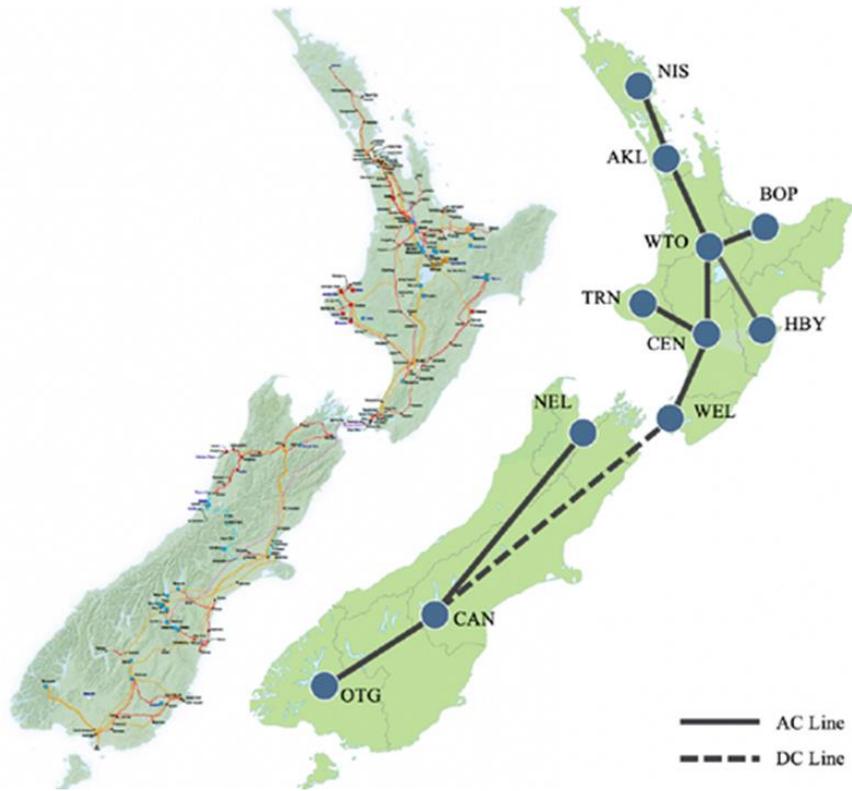


Figure 1: Transmission network (Transpower, 2021) and 11 node JADE approximation

Demand

As described in Section 5, we are using 8 load blocks. These have been computed from the residual load duration curves in the 11 regions that we model. Our approach clusters trading periods within each week based on similar national load. In Figure 2 we illustrate these loads for the peak block (B1) and an offpeak block (B7), for a subset of weeks in the year.

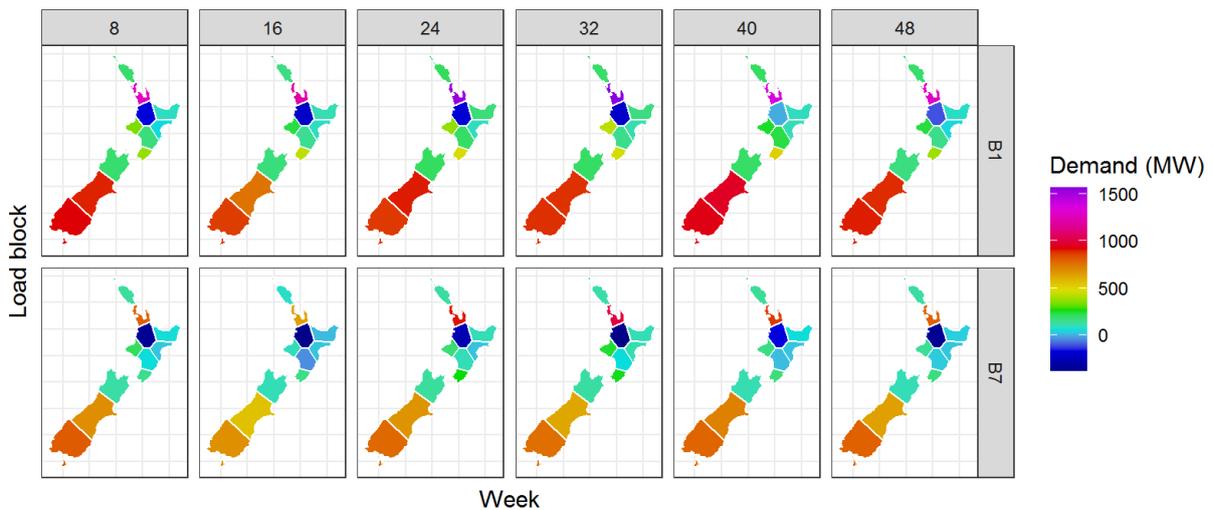


Figure 2: Subset of regional load blocks calibrated to 2020.

Overall, the load blocks ensure that the total GWh each week is the same as the sum of the half-hourly data, but there will be some approximation as, for example, the load for the peak block is an average of a number of trading periods.

Spatial correlation is maintained by clustering all regions together, based on national net-load. (This may mean, for example, that block B1 may not be the peak block for some smaller regions – but this ensures that the national peak load is approximated correctly.)

Steady-state water valuation

As stated in Section 5, DOASA uses a fixed terminal water value *curve*, which is a function of national storage. JADE can compute this endogenously, through solving a discounted infinite-horizon model to produce a steady-state policy. This policy is defined for each week of the year by a water-value *surface* (i.e., a function of the vector of reservoir levels).

This steady-state policy utilises a discount rate for convergence. Experiments have been undertaken with 6000 iterations for discount factors 0.95, 0.925 and 0.9, corresponding to discount rates of 5.3%, 8.1% and 11.1%, respectively. The terminal marginal water value curves are shown in Figure 3, below⁷.

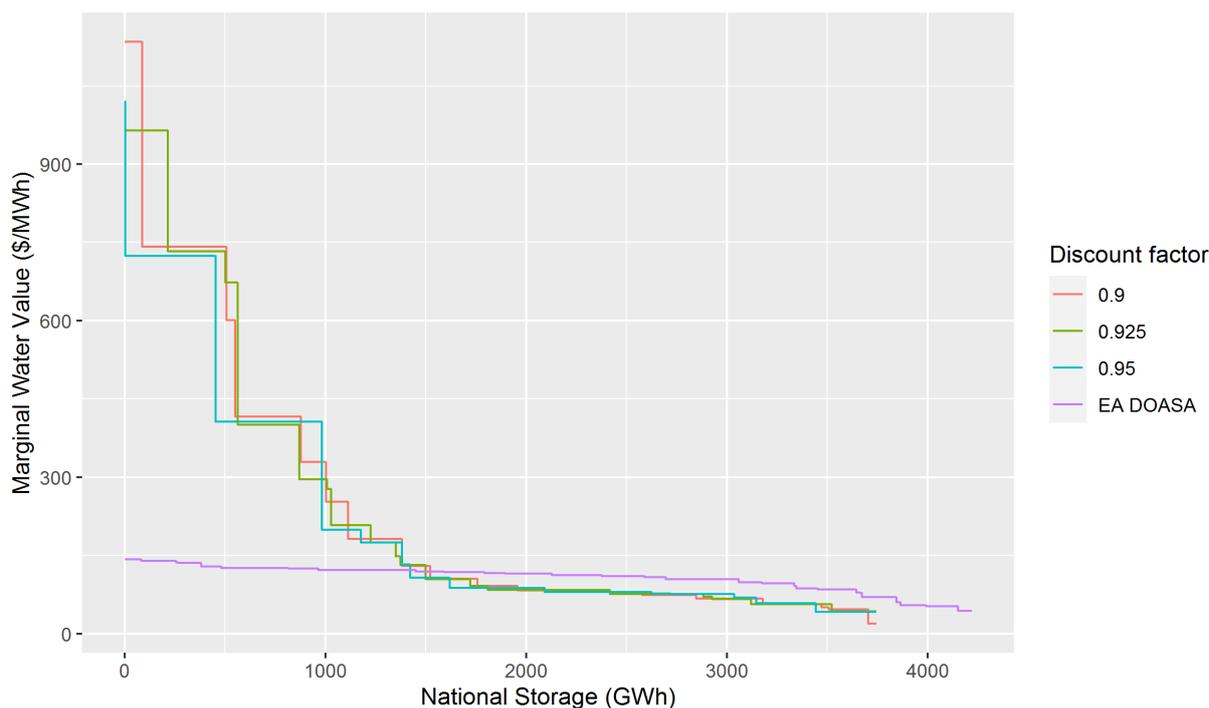


Figure 3: Projected terminal marginal water values from JADE with discounting.

As can be seen from these curves, the terminal water value is quite insensitive to the discount factor, over the range we have explored. Therefore, all our subsequent experiments have been undertaken using the terminal marginal water values corresponding to the discount factor 0.925. It should be noted that the flat region

⁷ The JADE curves are the projections of the JADE water-value surfaces onto a single dimension (national energy storage).

of the curve is typically linked to the costs of the (likely) marginal thermal generators, whereas the left-side is based on the cost and likelihood of load-shedding. The EA's curve is higher on the right and lower on the left than the marginal water value curves computed from JADE.

Thermal fuel costs

2020 thermal fuel costs are taken from the EA DOASA inputs, adjusted to account for the different ways that JADE and DOASA account for CO₂ prices.

Beyond 2020, thermal fuel costs are assumed to revert to lower levels in the medium term; we base these prices on the 2017 prices, with a carbon price of \$25/tCO_{2e}. Our choice of \$25 reflects what would be a conservative expectations of future CO₂ prices in 2020. With hindsight, given today's NZU prices, this estimate is probably too low⁸. This will be reflected in the level of the flat region of the marginal water value curves in Figure 3. Observe that the EA marginal water values are higher than ours in this part of the curve in Figure 3, reflecting a view that fuel costs will remain at 2020 levels in later years.

6.3 Results and Analysis from 2020

We now present some results of applying the methodology described above to the calendar year 2020. Our purpose here is to demonstrate how our software should be applied to monitor the performance of the wholesale market rather than to make definitive statements about its competitiveness in 2020.

In step (1) we trained JADE with levels of the risk aversion parameter $\lambda \in \{0.0, 0.1, 0.2\}$. Here $\lambda = 0.0$ corresponds to a risk-neutral policy. As λ is increased a higher weighting is (typically) placed upon the worst-case inflow from the set of training years.

Terminal marginal water value

For each level of risk aversion, we train an infinite-horizon model with an annual discount factor of 0.925 to create a set of terminal marginal water values for our 2020 benchmark models. The terminal marginal water values on December 31 for each policy are shown in Figure 4.

⁸ We observe that NZU prices on December 22, 2020 are \$68.50/t.

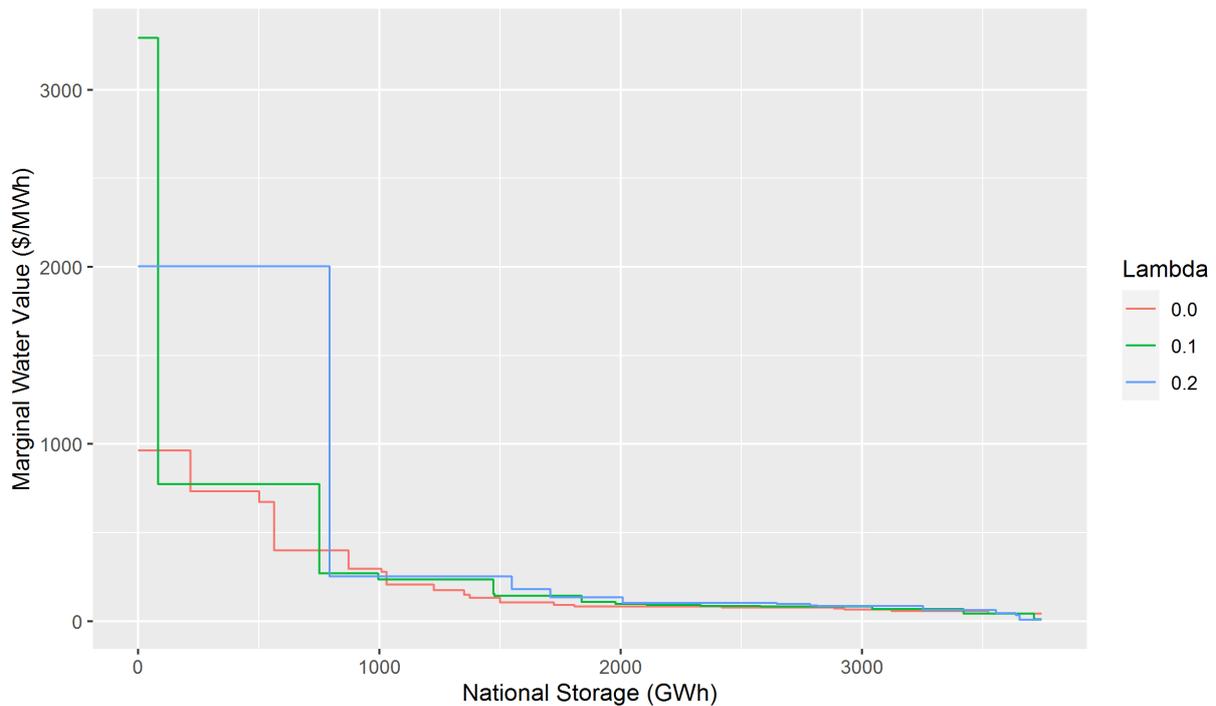


Figure 4: Projected terminal marginal water values from JADE with various risk measures.

We can see that the policies with higher risk aversion have higher marginal water values at end-of-horizon energy storage levels less than 2000GWh. Another point to note is that for energy storage less than 800GWh there are no more tranches for the most risk-averse policy. This is because these points would be rarely visited in steady state, given the policy.

Simulation of policies

After training JADE on the calendar year 2020 (Step 2) for each risk level, we simulate each policy on the 89 historical inflow sequences from 1932 to 2020, all starting with a single storage level at the beginning of 2020. Figure 5 shows the total energy storage (excluding contingent storage and the additional storage at Tekapo in the middle of the year). The light grey lines are the 89 trajectories of energy storage. On each plot there are two lines highlighted: in red is the trajectory for the policies applied to the 2020 inflow sequence, and in blue is the market energy storage trajectory (which is the same for all plots). The black line is the median trajectory.

It is interesting to note that the risk-neutral JADE policy follows the market trajectory very closely until week 36, at which point the market trajectory deviates towards the $\lambda = 0.1$ trajectory.

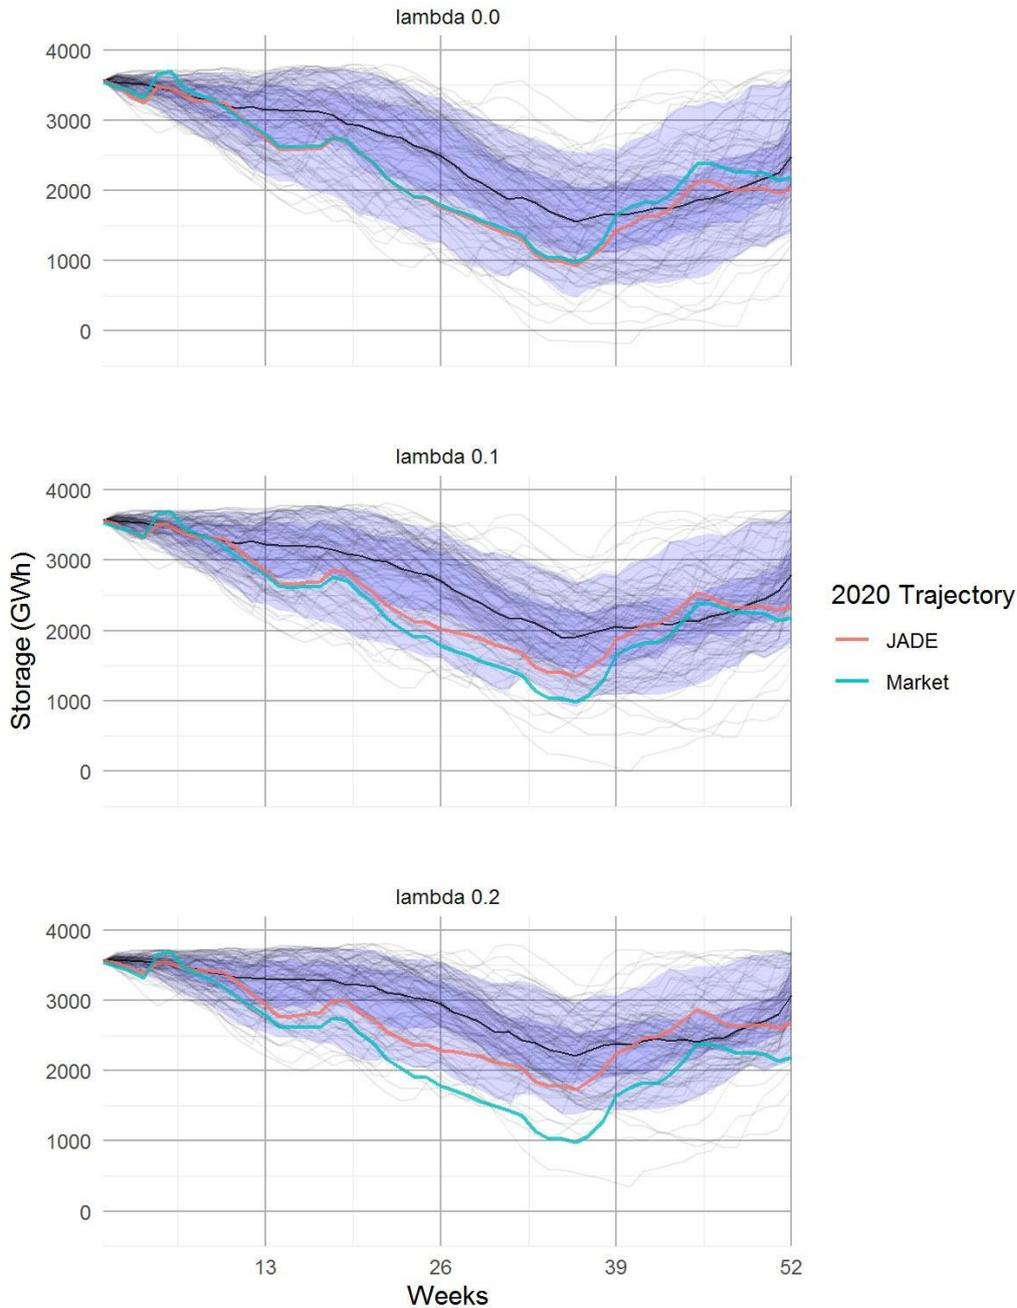


Figure 5: Storage trajectories for hydro inflows for years 1932-2020. Shaded regions are 5th-percentile to 95th-percentile (light blue) and 25th-percentage to 75th-percentile (dark blue).

The version of JADE we have implemented represents the individual main storage reservoirs separately (rather than aggregating). Figure 6 shows the equivalent figure to Figure 5, broken down into individual reservoirs. Here storage is measured in Mm3 rather than GWh.

Observe that the historical reservoir levels (blue curves) deviate from the JADE levels (red curves) except in the cases of Tekapo and Pukaki, where blue and red curves are quite close at least in the risk-neutral case. In the spirit of this submission, we are tempted to speculate on the causes of these deviations from system optimality, but we have no further evidence to support any of the following conjectures:

1. Contact Energy (the operators of the Clutha hydro scheme), and Mercury Energy (the operators of the Waikato hydro scheme) might have better information than JADE has used, resulting in different optimal policies.
2. As owners of gas plant in the North Island Contact Energy might dispatch Clyde and Roxburgh more aggressively than JADE, given the gas shortage in 2020. Offering hydro generation at low prices to meet contracts or retail load would draw down Hawea more than perfectly competitive offers (see Section 4.3)
3. The system optimal policies computed by JADE can diversify risk across all generation plant. It is possible that JADE is trading risk between Mercury and Contact that allows each generator to maintain lower storage in Taupo (Mercury) and higher storage in Hawea (Contact).

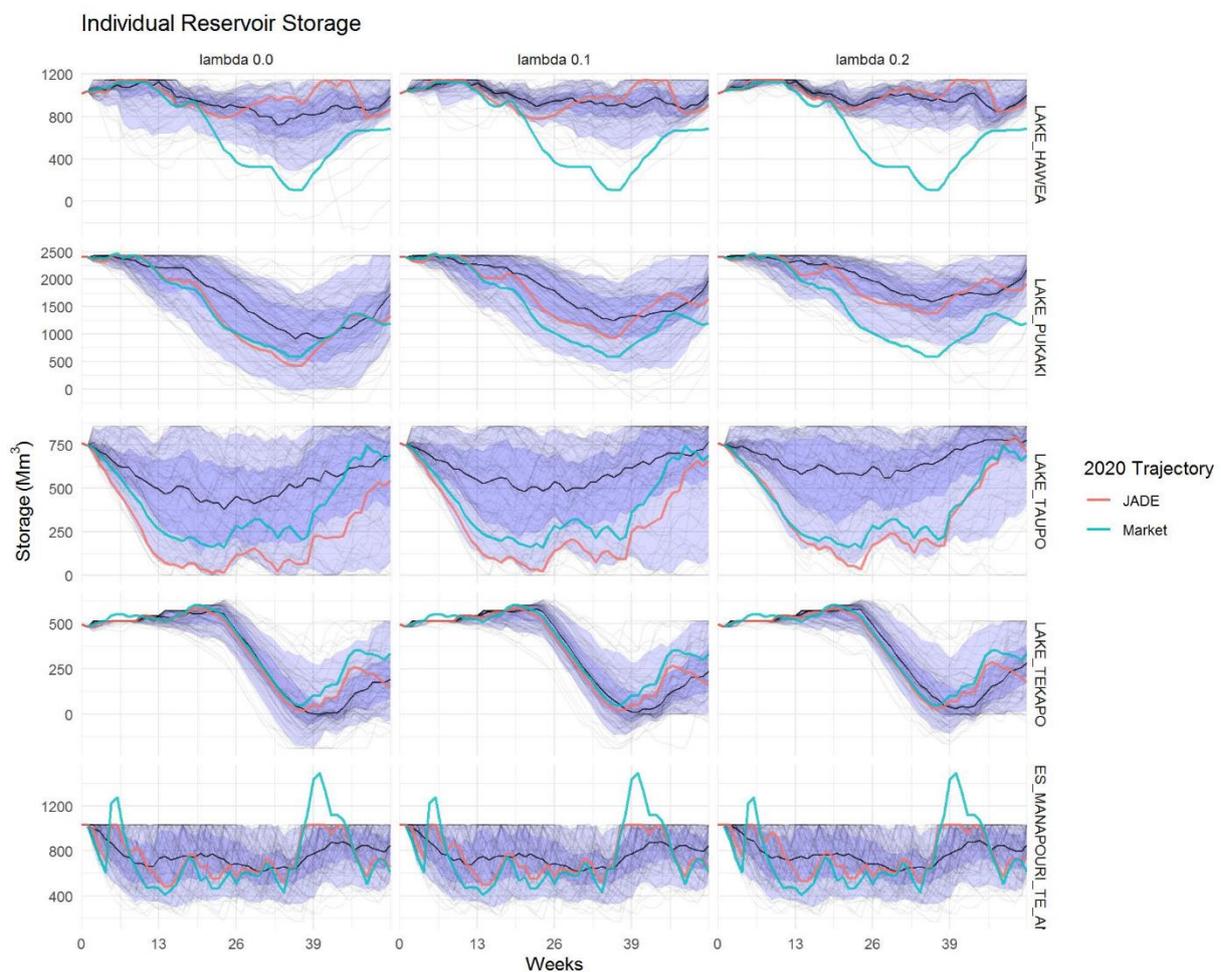


Figure 6: Storage trajectories for individual reservoirs for years 1932-2020.

Shaded regions are 5th-percentile to 95th-percentile (light blue) and 25th-percentage to 75th-percentile (dark blue).

JADE outputs marginal water values from the reservoir levels obtained from each simulated trajectory. These are plotted in Figure 7. There are some periods where water values are negative; these correspond to events where upper bounds on

water flows are exceeded (e.g., flood events)⁹. These violations are penalized in JADE and will lead to negative marginal water values, i.e., extra water incurs a cost.

For all reservoirs, the marginal water values are higher during the winter for more risk-averse policies (which assign low future inflows a higher probability than estimated from history). This reverses in the spring of 2020 where the risk-averse policies result in higher levels of storage (see Figures 5 and 6) thus decreasing the marginal water value (as demonstrated by Figure 4). Observe that all marginal water values in these simulations are below \$95/Mwh.

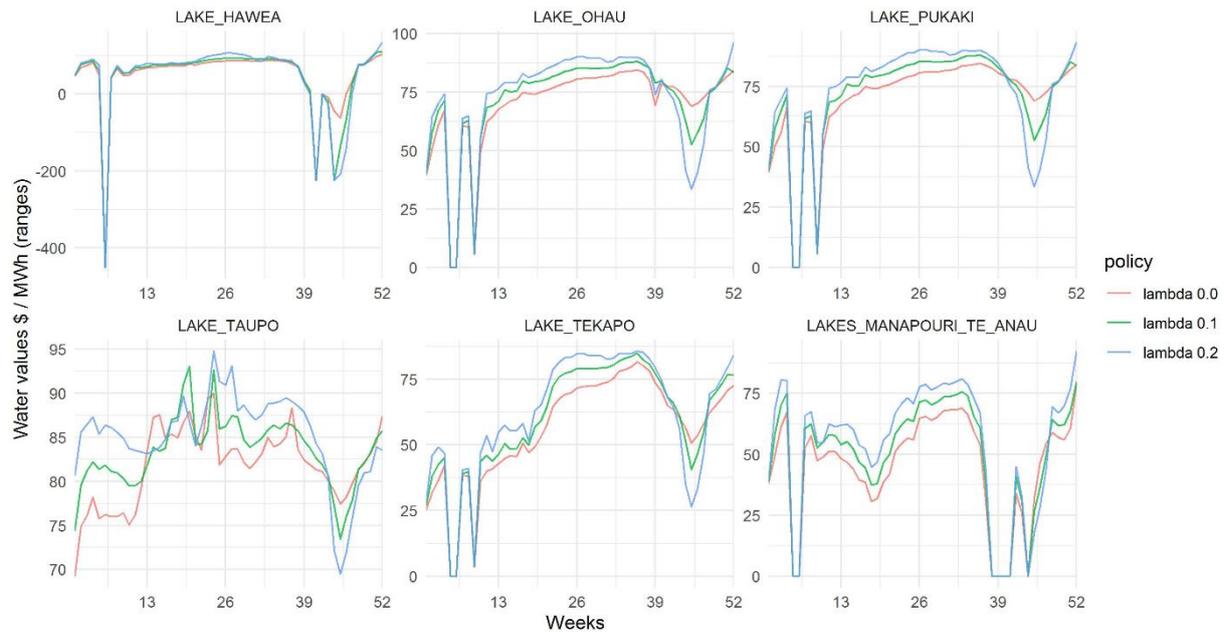


Figure 7: 2020 marginal water values.

JADE enables the user to extract prices for each region, week and load block in a simulation. These can be compared with corresponding values computed from historical prices at representative nodes in each region. One way of visualizing these data is a set of heat maps as shown in Figure 8 that shows weekly time-averaged prices in each region for each policy setting, with the prices for AKL and CAN written next to the corresponding region. Observe that before week 20, historical prices are below the JADE values, but become much higher after week 20.

⁹ In week 6 of this simulation Roxburgh and Clyde were spilling over their limits.

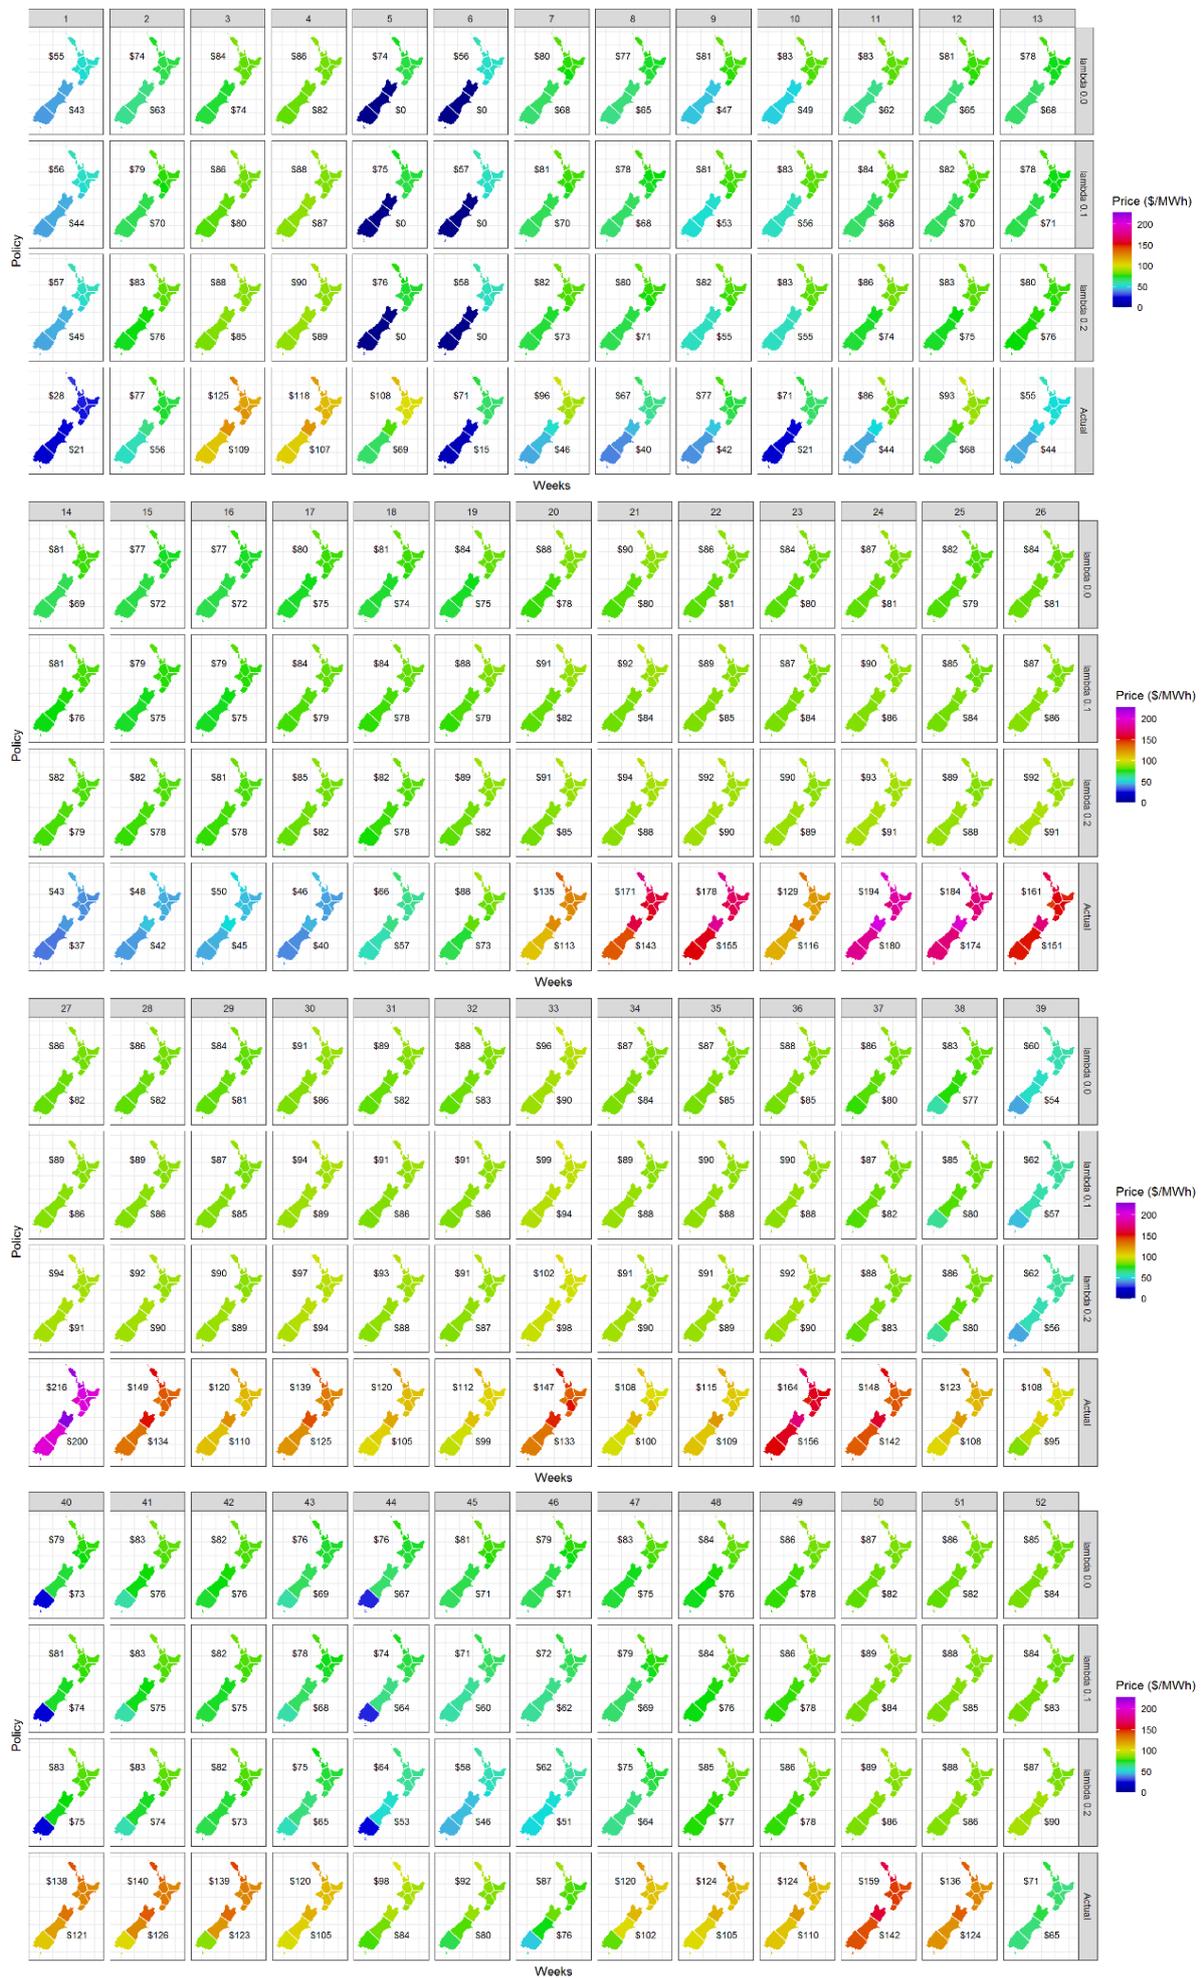


Figure 8: Weekly time-weighted average prices by region.

Figure 9 shows the weekly time-weighted average prices from JADE for regions AKL and CAN, and the corresponding values computed from historical prices at nodes OTA2201 and BEN2201. This shows that in weeks 1-19, prices from both JADE and the market fluctuate around similar values, with prices in the South Island consistently lower in the market from week 6 to week 19. After week 20, prices in the market were higher than those for all the JADE policies.

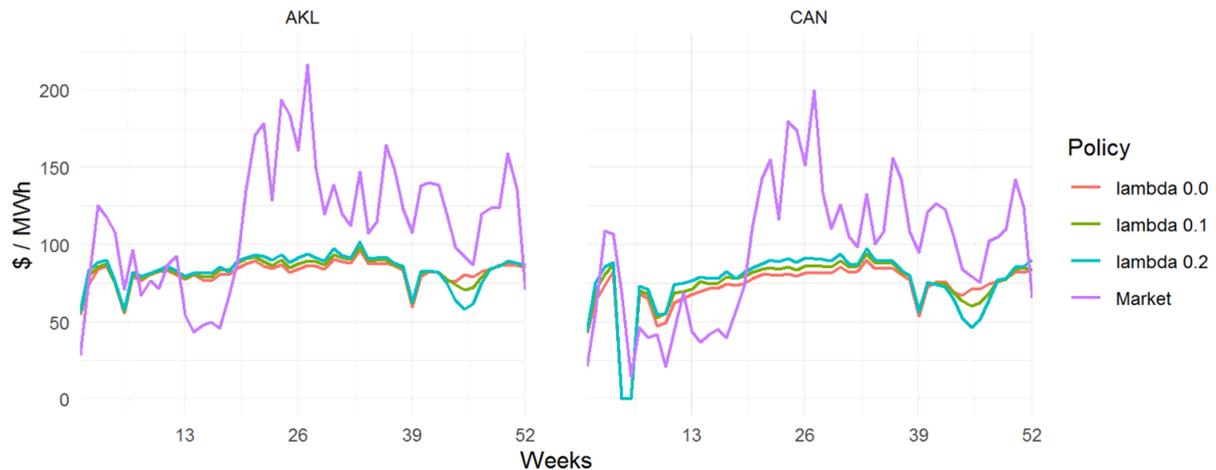


Figure 9: 2020 weekly time-weighted average prices comparison.

We are able to analyse price differences further by comparing prices in the peak (B1) and off-peak (B8) load blocks, as well as load blocks B3 and B6; this is shown in Figure 10. In each subfigure we are comparing weekly prices in the load block from each of the JADE policies simulated for 2020 with the market prices averaged over the set of periods corresponding to that load block. The left column shows prices for the AKL region, and the right column shows prices for the CAN region. For all policies the weekly prices increase as demand increases; this is much more pronounced in the market prices. However, for load block B1, all policies have price spikes in several weeks over the winter. Interestingly, after week 20 the off-peak (B8) market prices correspond quite well to the prices from the competitive benchmarks.

Another observation is that prices for different values of risk aversion are similar, but the storage trajectories deviate substantially. A key observation is that although with a risk-averse policy, we wish to maintain a higher level of storage, that doesn't require prices to stay high throughout the year – in fact, as we observed above, prices reverse towards the end of the year (with prices higher in the risk-neutral policy) as the reservoirs begin to fill up.

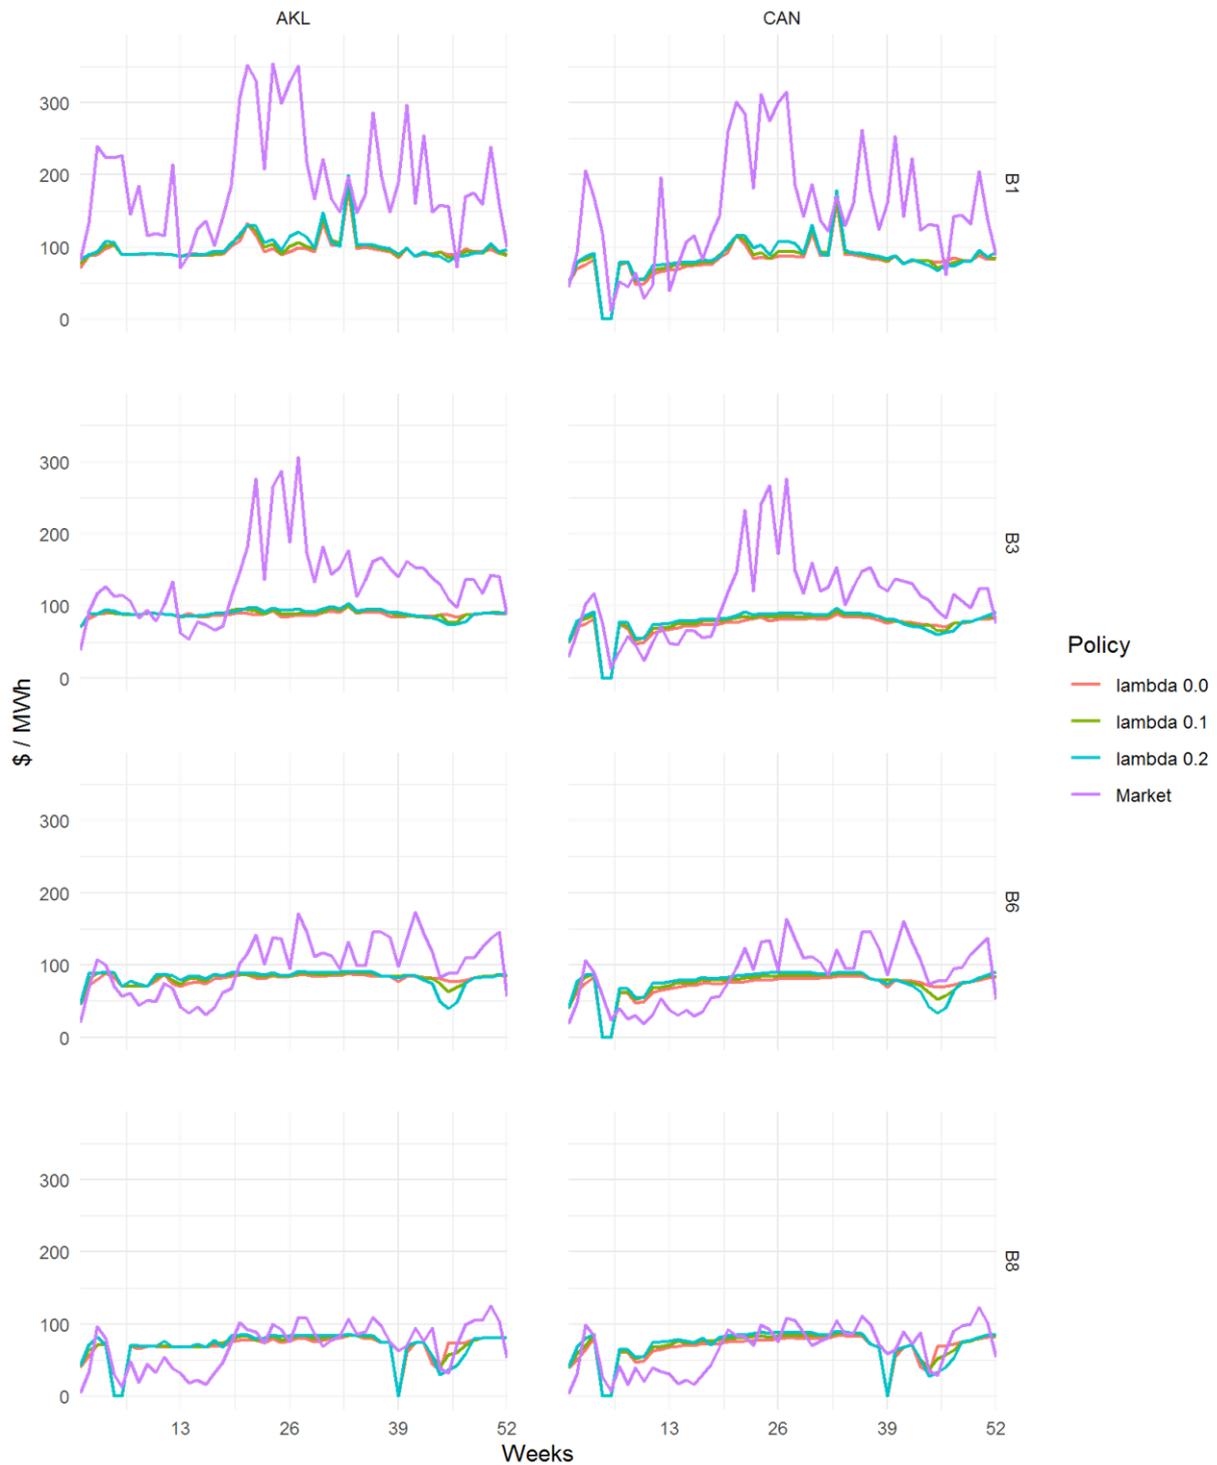


Figure 10: 2020 weekly prices for load blocks B1, B3, B6 and B8.

Shortage risk

We can use JADE to simulate forward from any storage level part-way through the year. In Figure 11, we demonstrate this feature by simulating from week 27. For each of the policies, the black curve follows the 2020 JADE trajectory up to week

26, after which we simulate forward using historical inflow sequences from 1932 to 2020¹⁰.

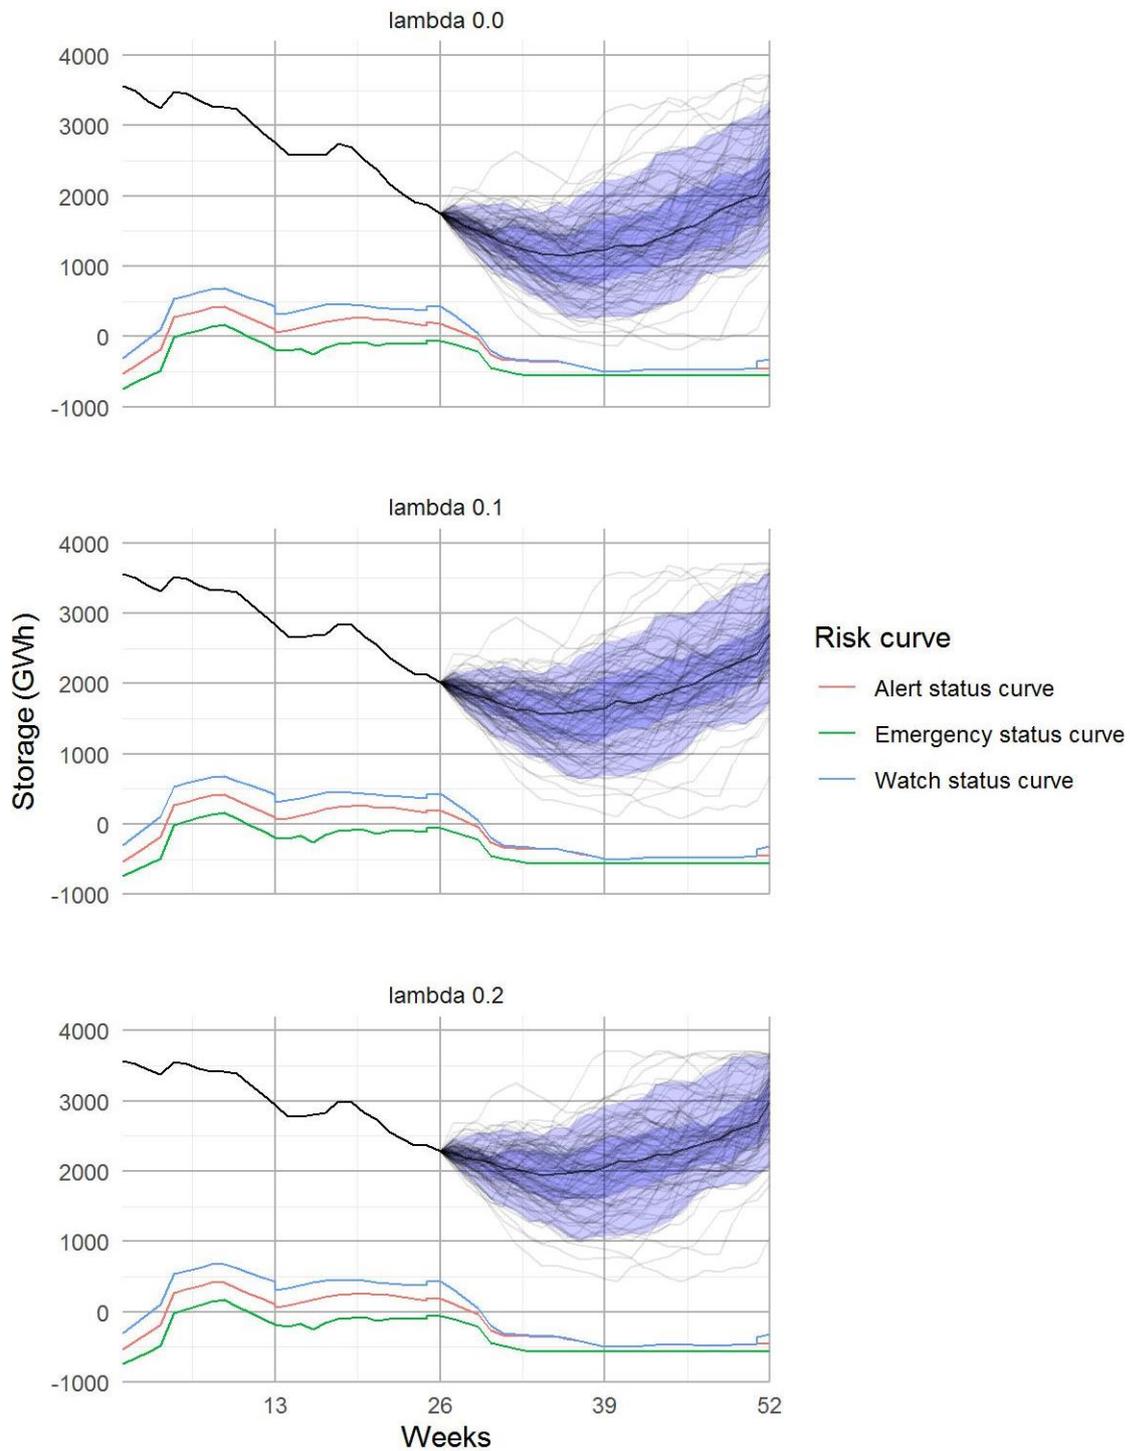


Figure 11: 2020 weekly stored energy trajectories up to week 26 then projections on 1932-2020 inflows.

¹⁰ The national storage level is measured in GWh and does not include the energy content of contingent storage, which means that the storage levels in Figure 11 and Figure 12 should be increased by the energy content of contingent storage if they are to match Transpower’s risk curves.

Figure 12 reproduces the simulation shown in Figure 11 but starting from the historical level (shown by the blue trajectory).

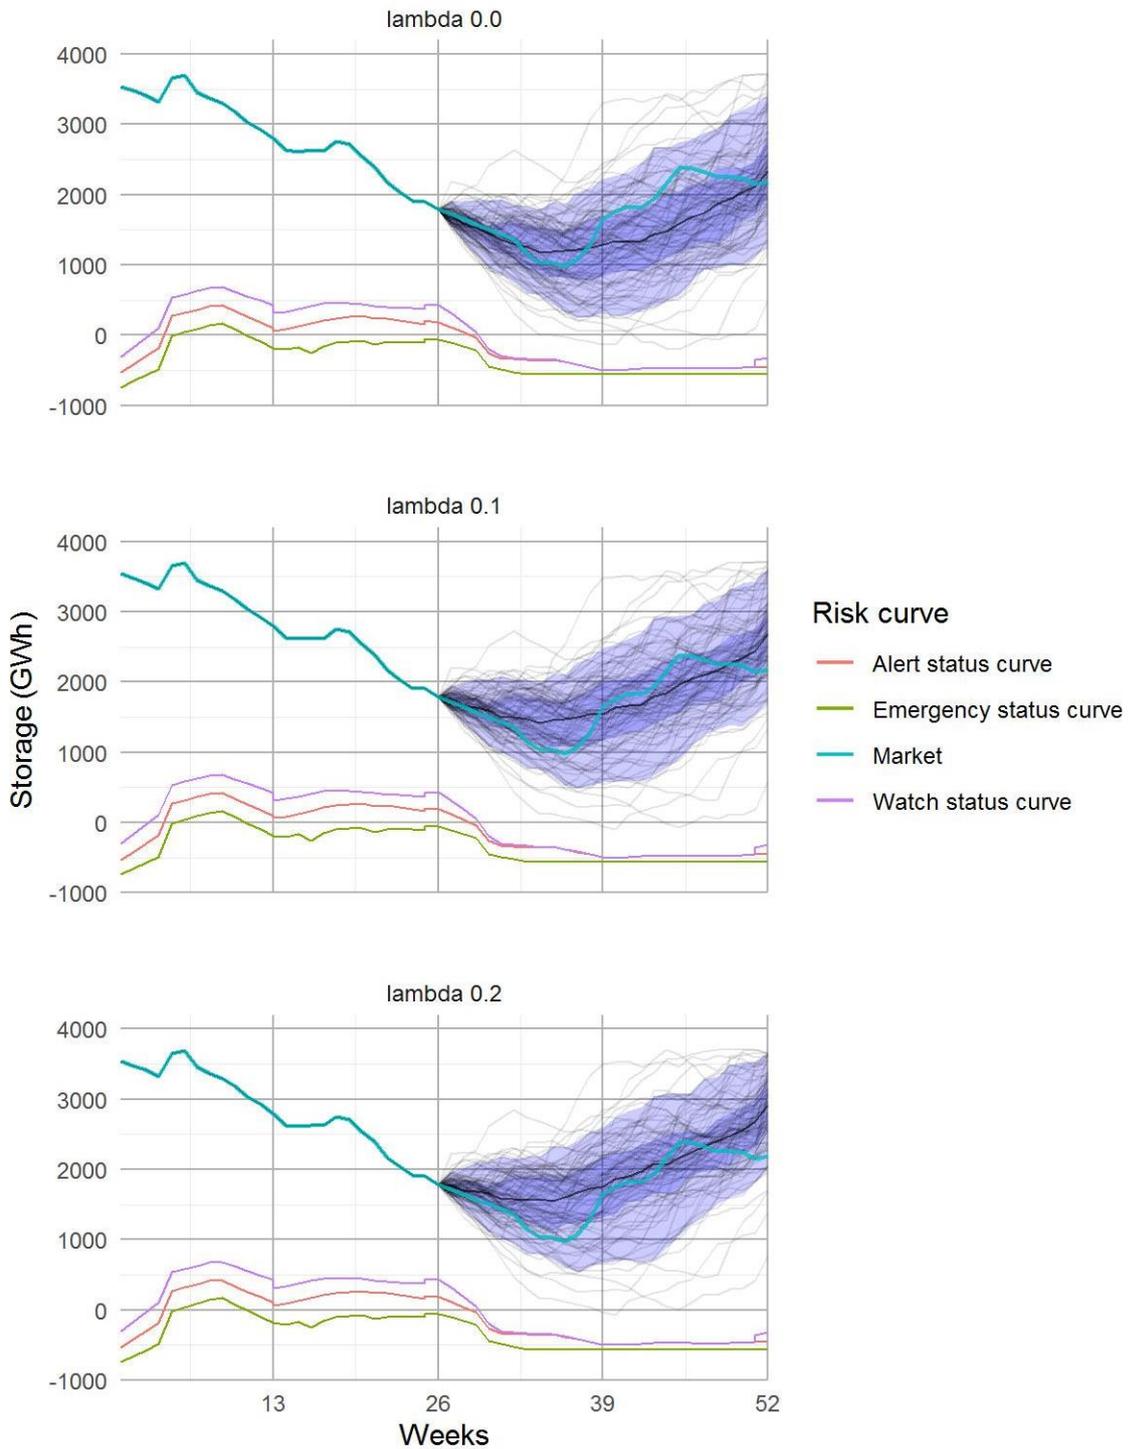


Figure 12: Actual 2020 weekly stored energy trajectories up to week 26 then projections on 1932-2020 inflows.

Figure 13 shows the prices for 2020 when simulating three different policies from week 26 using 2020 inflows, and historical reservoir levels (as in Figure 12). The

prices in off-peak periods match reasonably closely, but prices in peak periods differ considerably. We see that in week 33, peak prices (B1) spike higher than in Figure 10. In fact, prices spike higher than the market in that week with $\lambda = 0.2$. It is possible that these differences are caused by constraints not modelled in JADE. We attempt to account for these in Section 7.

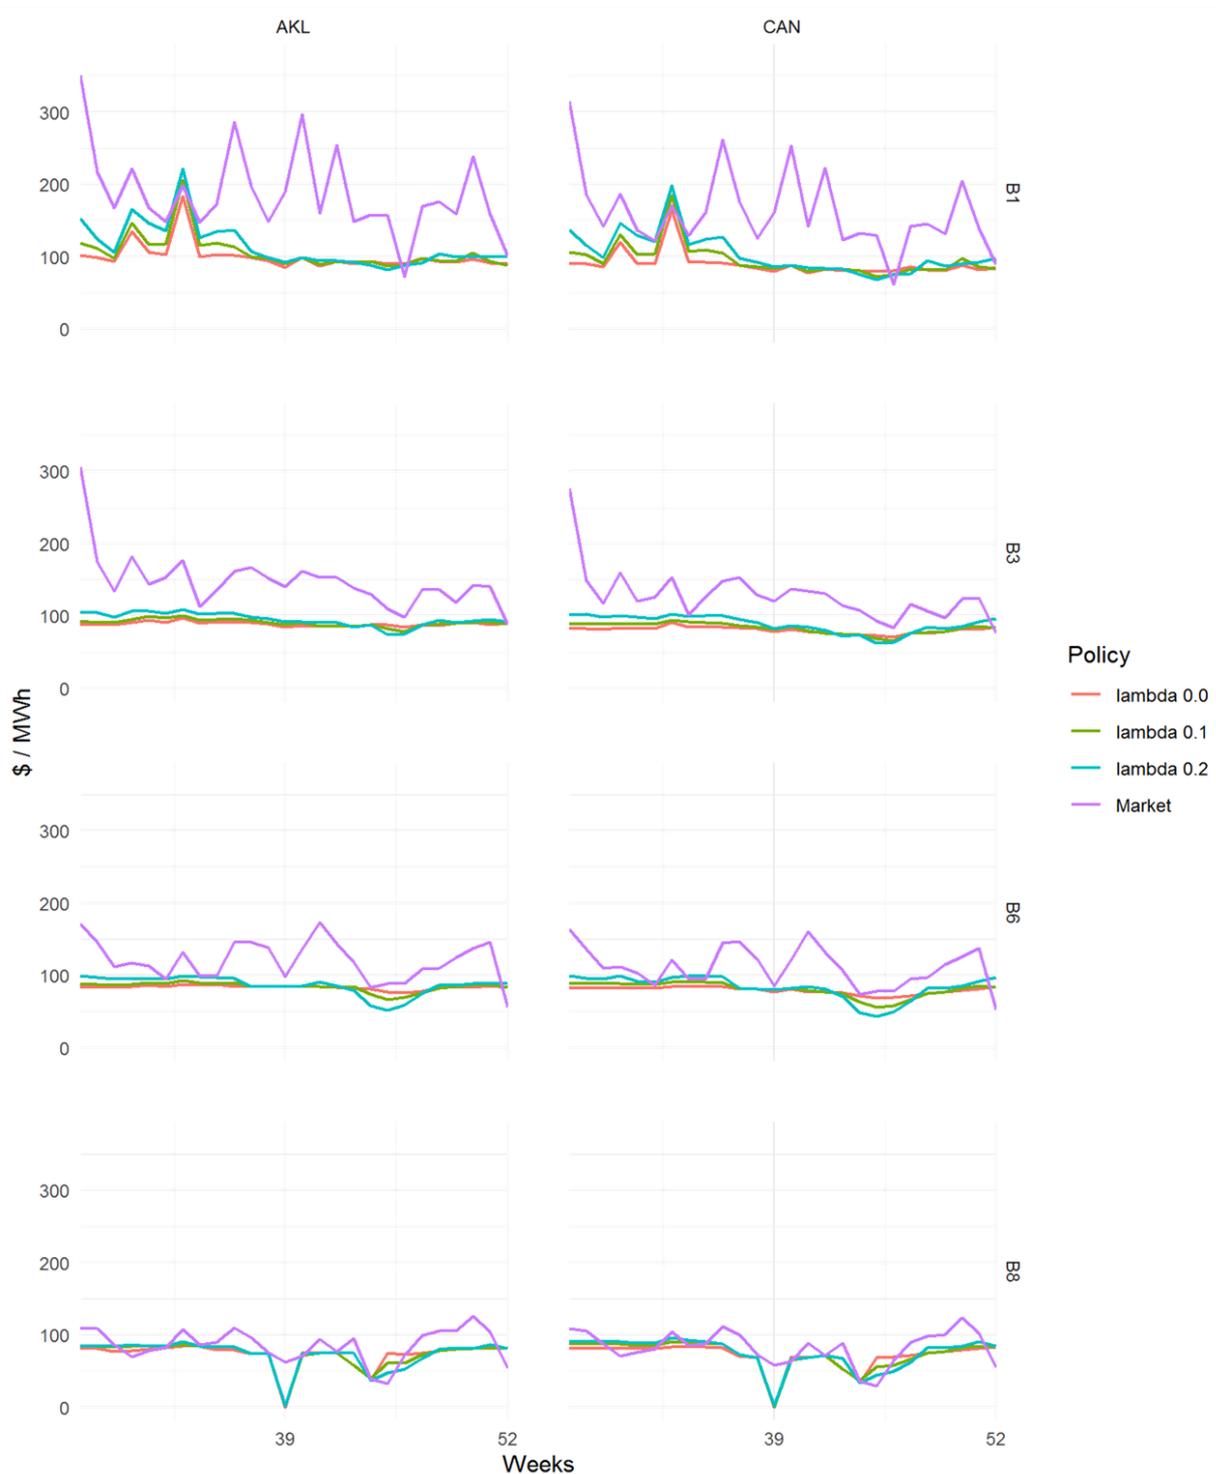


Figure 13: Simulated prices from week 27 to week 52, with storage matching market storage for 2020 at the end of week 26.

Section 7: HydrovSPD

To assess competition using counterfactual models requires a robust method for computing the opportunity value of stored water. This requires some form of stochastic equilibrium model. Under assumptions of convexity, completeness and perfect competition this can be shown to be equivalent to a stochastic optimization model. The previous sections have demonstrated how this can be used to compute perfectly competitive marginal water values under different assumptions about levels of risk aversion.

As mentioned above, the outcomes of the New Zealand wholesale electricity market can be replicated exactly using the vSPD software maintained by the Electricity Authority on the EMI site. This software can be used to simulate the outcomes that would result from different energy offers from those made historically.

The question then remains: how do marginal water values computed using JADE translate into perfectly competitive energy offers that are inputs to vSPD? This translation is a key ingredient to producing a perfectly competitive dispatch solution using vSPD and comparing prices and operating surpluses with historical values.

The situation would be simple if all hydroelectric generators operated a single plant fed by its own reservoir. Then the counterfactual offers to be submitted to vSPD would be the plant capacity offered at the current marginal water value. In reality, reservoirs typically form part of a river chain where water released from one reservoir may flow through several stations, possibly filling reservoirs downstream. Furthermore, there are often environmental constraints limiting the releases of water at different points on the river chain. In JADE we represent the Clutha scheme, the Waitaki scheme and the Waikato scheme as river chains.

Block dispatch enables operators of river chains some freedom to reallocate water after dispatch, but even if a river-chain is block dispatched, the dynamics of river flows require some linking of offers in different periods. This linking is implicit in the offers submitted by river chains to the system operator.

The solution adopted by the New Zealand wholesale market to this intertemporal linkage is to repeatedly compute and publish provisional dispatch solutions over the next 36 hours. This enables generators to revise their energy offers up to gate closure as they observe the published outcomes of their previous offers and those of their competitors. The sequence of offers made must depend on the marginal water value, but also on a lot of other factors. How does one account for these to determine if the offers made are close to perfectly competitive?

Our approach is the model HydrovSPD described in [9]. This is a 48-period version of vSPD that incorporates reservoirs and river flows in the Waitaki, Clutha and Waikato systems, and links consecutive periods with water transfers between station headponds. Inflows are assumed to be constant throughout the day and known at the start of the day, and are set to be their daily average levels.

HydrovSPD solves all 48 trading periods in a day in one model. This means that the model assumes perfect information of all parameters throughout the day. In other words, HydrovSPD is *clairvoyant*. As a result, the optimizer can guarantee that all river-chain constraints are satisfied throughout the day¹¹.

The results of any counterfactual experiments carried out using HydrovSPD depend on a number of assumptions regarding the many inputs to vSPD. We assume that these inputs are identical to those used historically except for generation and reserve offer prices. Thermal generators offer the same quantities as historical offers but offer at their short-run marginal cost (computed from gas and coal costs and other variable costs), and hydro generators pay for daily use of reservoir water at its risk-adjusted opportunity value computed using JADE. Reserve is offered at zero cost but must satisfy the reserve requirements that pertain to the historical period.

There are several ways that HydrovSPD can be used to monitor market outcomes. As shown above we can identify weeks in the year where time-averaged prices computed using JADE differ from historical values. Week 20 in 2020 is an example. HydrovSPD can then be applied to each day in this week to see if the JADE prices are low because JADE is missing key constraints that might materially affect prices.

To illustrate this application, consider Tuesday, May 19, 2020, which is the last day in week 20. The cuts that represent the marginal water value (with $\lambda = 0.1$) at the end of this week were input to HydrovSPD, which produces prices at nodes BEN2201, HAY2201, and OTA2201 shown in Figure 14. These are compared with the historical prices. The prices in the competitive solution peak fluctuate as shown driven by the dispatch of the Whirinaki peaking plant. BEN2201 prices remain low throughout, so there is substantial price separation between BEN2201 and HAY2201 in some periods in contrast to the historical solution.

¹¹ The clairvoyance implies that solutions will anticipate future events in the day, so the optimal costs will be lower than what can be achieved without this knowledge. Although this gap has been shown to be small in historical studies (Porter, N.R., Intra-day uncertainty and efficiency in electricity markets, <https://www.epoc.org.nz/theses/NickPorterFinalMasters.pdf>, 2015) it will become more important as wind and solar capacity increases.

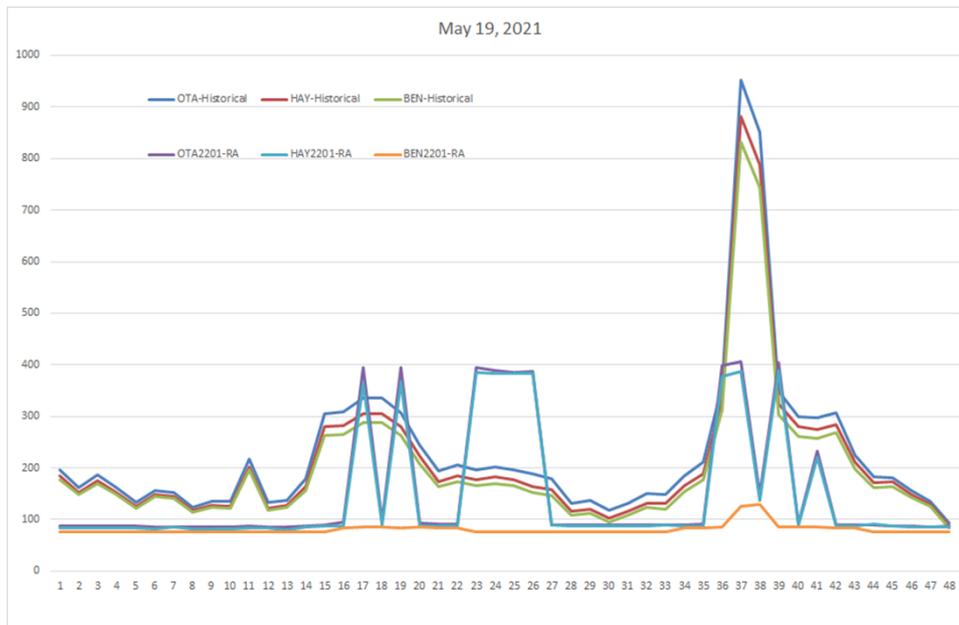


Figure 14: Historical prices from May19, 2020, compared with HydrovSPD prices computed from risk averse cuts

Although the focus of benchmarks is naturally on the wholesale prices they deliver, as these are what consumers pay either directly, or averaged and risk-adjusted through contracts or retail prices, HydrovSPD can be used to compare other performance indicators.

The water released by the historical solution and by HydrovSPD will be different and so it is important to keep track of this, since savings in HydrovSPD thermal fuel cost from releasing more water must be somehow paid for by releasing less later on.

Although we have not done this yet for 2020, we report in [9] on an experiment that applies HydrovSPD to every day in 2017¹². In this study, HydrovSPD and DOASA are alternated. DOASA is run every two weeks with updated reservoir storages produced by a sequence of 14 HydrovSPD runs, one for each day of the two weeks. This gives 365 consecutive days of dispatch and corresponding storage trajectories.

The detailed outputs from this experiment are presented in [9]. The results show similar prices to average historical North Island prices (although there are significant differences over time), but lower average South Island Prices than those observed historically. The competitive counterfactual model generates more energy from the Waitaki system and constrains the HVDC line more often. One can contrast the resulting counterfactual prices with historical values as shown in Figure 15 and Figure 16.

¹² We present these results to illustrate the approach rather than draw conclusions about market performance in 2017. The improvements in JADE that we describe above might yield different results from DOASA when combined with HydrovSPD.

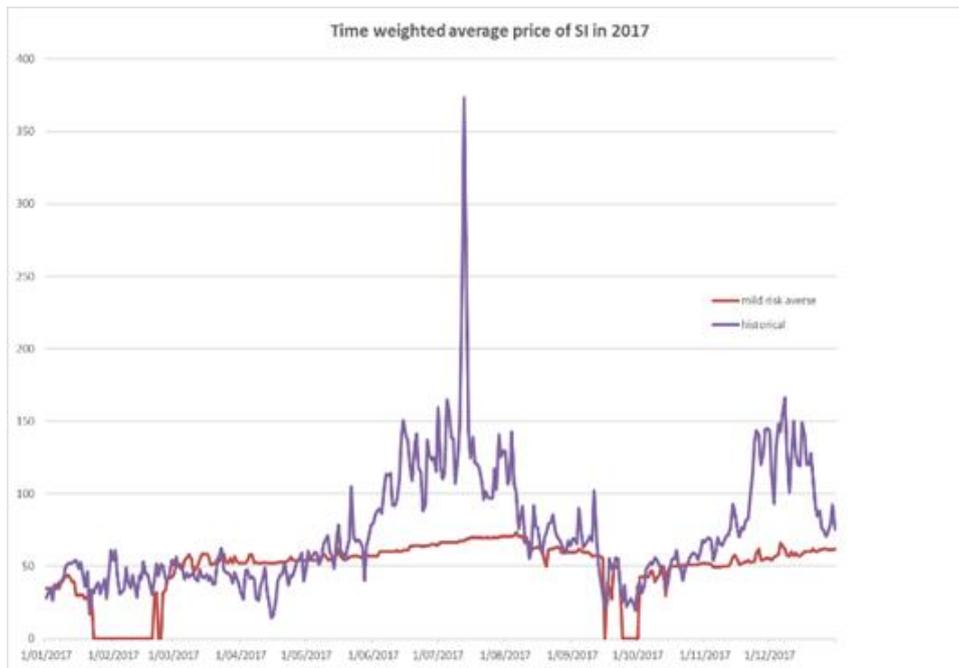


Figure 15: Historical time-weighted daily prices at Benmore (purple) compared with counterfactual prices (red). Historical prices are \$23/MWh higher on average (reproduced from [9]).

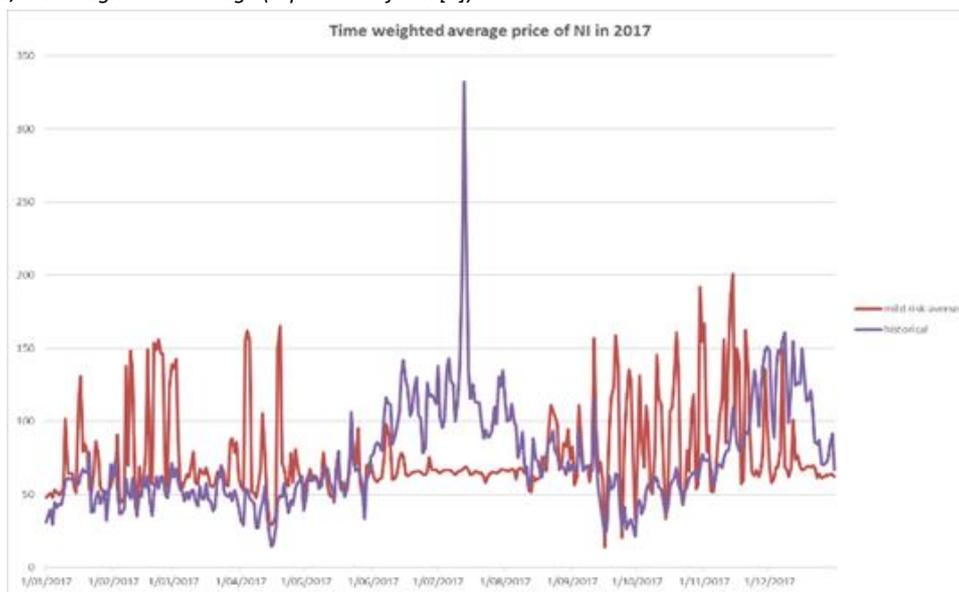


Figure 16: Historical time-weighted daily prices at Otahuhu (purple) compared with counterfactual prices (red) (reproduced from [9]).

Historical HVDC rentals in 2017 computed using vSPD prices were about \$4.3M. The rentals from the competitive counterfactual model were approximately \$20 M.

The paper [9] also has comparisons of Ricardian rent for historical and perfectly competitive counterfactual solutions for 2017, and comparisons of historical and counterfactual storage trajectories for 2017. The former can be used to assess to what extent a perfectly competitive model results in missing money for generators. One might use the latter to assess whether the level of shortage risk in the competitive counterfactual model is acceptable. We note that 2017 is one sample path so care must be taken in drawing conclusions about risk from such a

simulation; the simulations carried out using JADE in Section 7 of this report over 89 historical years provide a more robust assessment of shortage risk.

Section 8: Discussion

This document has focused on models that we believe should be used to monitor outcomes in the New Zealand wholesale electricity market. The purpose of monitoring is to identify episodes where market outcomes diverge from perfectly competitive ones. These can then be examined in much greater detail to discover why this is happening.

It is important to recognize that the outcomes from counterfactual models depend on the assumed inputs. JADE relies heavily on cost estimates to produce the marginal water value curves in Figure 3. The flat part of each curve varies in height depending on assumptions of thermal fuel costs and CO₂ prices. Identifying the sensitivity of outcomes to these cost estimates is important. This is discussed to some extent in [9], where two different cost estimates are modelled. The assumption that natural gas is a readily available resource (at a known price) is also open to question in the light of recent gas shortages. It is possible that including a gas reservoir model in JADE (that behaves similarly to a hydro reservoir) might provide more realistic estimates of the cost of gas for electricity generation.

The second important assumption is the estimated cost of demand response and involuntary load shedding. As the electricity system reduces its dependence on coal and gas, demand response costs will play a greater role in determining electricity prices, and so some care and attention must be devoted to estimating them. Although JADE has the flexibility to include sophisticated demand reduction mechanisms, in this submission we have used the same load-shedding costs as the EA assumed in their review document. The values chosen for load-shedding costs affect the left-hand end of the marginal water value curves in Figure 3 (and to some extent the slope of the flatter parts). One contribution of this submission is showing how JADE's infinite horizon model yields arguably more realistic end-of-horizon marginal water value curves (at least at the left-hand end) compared with those assumed by DOASA¹³.

Some caution is needed in drawing conclusions about the motives of market participants from counterfactual models. Exercise of market power is commonly blamed for deviations from perfectly competitive outcomes, but other factors might have caused these. For example, a generator might be acting to avoid risks that she cannot hedge using available market instruments. Having said this, risk-averse counterfactual models like JADE can help identify situations where missing markets for risk trading might have contributed to this difference.

HydroSPD shows that there are significant deviations between a clairvoyant central plan for daily dispatch and the current market solution. The outcomes of the repeated pre-dispatch discovery process can result in productive inefficiencies as reported in [9]. These can arise from generators attempting to meet inter-temporal

¹³ Recall that DOASA assumes an exogenous end-of-horizon marginal water value curve, whereas this is computed in JADE by training on an infinite horizon.

constraints by altering offer prices. In many other wholesale electricity markets intertemporal constraints (such as ramping constraints on large thermal units) are handled with two-settlement markets, a day-ahead market and a balancing market. Perhaps this is worth considering for river chain optimization.

References

- [1] Dowson, O. and Kapelevich, L., (2021) SDDP. jl: a Julia package for stochastic dual dynamic programming. *INFORMS Journal on Computing*, 33(1), pp.27-33.
- [2] Electricity Authority, (2021) Market Monitoring Review of Structure Conduct and Performance in the Wholesale Electricity Market - Information Paper, <https://www.ea.govt.nz/assets/dms-assets/29/Monitoring-Review-of-structure-conduct-and-performance-in-the-wholesale-electricity-market-updated-paper.pdf>.
- [3] EMERALD, downloadable as <https://www.epoc.org.nz/papers/EMERALD.pdf>
- [4] Ferris, M.C. and Philpott, A.B., (2021) Dynamic risked equilibrium, *Operations Research*, <https://doi.org/10.1287/opre.2019.1958>, (downloadable from www.epoc.org.nz).
- [5] Kok, C., Philpott, A. and Zakeri, G., (2018) Value of transmission capacity in electricity markets with risk averse agents, (downloadable from www.epoc.org.nz).
- [6] Pereira, M.V. and Pinto, L.M., (1991) Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, 52(1-3), pp.359-375.
- [7] Philpott, A.B., Ferris, M.C. and Wets, R., (2016) Equilibrium, uncertainty and risk in hydro-thermal electricity systems. *Mathematical Programming*, 157(2), pp.483-513, (downloadable from www.epoc.org.nz).
- [8] Philpott, A.B., de Matos, V.L. and Finardi, E.C., (2013) On solving multistage stochastic programs with coherent risk measures. *Operations Research*, 61(4) pp. 957-970. (downloadable from www.epoc.org.nz).
- [9] Philpott, A.B and Guan, Z., (2020) Benchmarking wholesale hydroelectricity markets: 2017 update, (downloadable from www.epoc.org.nz).
- [10] Philpott, A.B. (2021) Price discovery can be inefficient, (downloadable from www.epoc.org.nz).
- [11] Rosemberg, A., Street, A., Garcia, J.D., Valladao, D.M., Silva, T. and Dowson, O., (2021) Assessing the Cost of Network Simplifications in Long-Term Hydrothermal Dispatch Planning Models. *IEEE Transactions on Sustainable Energy*.
- [12] Ruddell, K., Downward, A. and Philpott, A.B., (2018) Market power and forward prices. *Economics Letters*, 166, pp.6-9, (downloadable from www.epoc.org.nz).

Benchmarking wholesale hydroelectricity markets: 2017 update*

Andy Philpott[†] Ziming Guan[‡]

November 27, 2020

Abstract

Using a suite of empirical models we study the effects of risk aversion on wholesale electricity prices that are affected by uncertainty in hydroelectric reservoir inflows. Our models combine stochastic dual dynamic programming with a high fidelity simulation model of the New Zealand wholesale electricity market. Results of the models are given for the calendar year 2017, and compared with wholesale market outcomes in that year.

Keywords: electricity market, hydroelectricity, stochastic inflows, risk-aversion, market power.

1 Introduction

In this paper we revisit an earlier study by the same authors [27] that investigated the effects of uncertainty on outcomes in wholesale electricity markets that are dominated by hydroelectric generation. The paper [27] applied risk-averse stochastic optimization models to a counterfactual model of the New Zealand wholesale market over the calendar year 2012. The current study incorporates a more sophisticated long-term model, and like the previous study links this with fully detailed model of the New Zealand dispatch model, improved to represent

*We would like to acknowledge support of the New Zealand Marsden fund under contract UOA1520. We would also like to acknowledge discussions and suggestions from Dr Anthony Downward, Dr Geoff Pritchard, Professor Danny Ralph and Professor David Newbery. The views and opinions expressed in this paper are, however, solely attributed to the authors.

[†]Electric Power Optimization Centre (EPOC), Department of Engineering Science, University of Auckland: a.philpott@auckland.ac.nz.

[‡]Electric Power Optimization Centre (EPOC), Department of Engineering Science, University of Auckland: z.guan@auckland.ac.nz.

the national instantaneous reserve market that came into force in late 2016. The current study is confined to 2017; we leave detailed analysis of other years to a forthcoming companion paper [28] that presents the results of applying our model to data from the ten historical years 2008-2017.

The perfectly competitive benchmarks that we compute are the result of simulating a system optimal policy. If electricity market participants maximize productive and allocative efficiency, then they will maximize social welfare for the system as a whole, at least in the short term. It is worthwhile to discuss this assertion in the context of the New Zealand wholesale electricity market consisting of five large gentailers who compete to supply electricity to consumers in an energy-only market.

We begin with an account of the classical theory of energy-only markets (see e.g. [33]). In a single-period setting with no uncertainty and convex costs the locational marginal prices emerging from an economic dispatch in which all energy is offered at its short-run marginal cost will correspond to a perfectly competitive partial equilibrium. If demand is elastic in the short term then in equilibrium these prices will be sufficient to sustain the optimal level of generation capacity. If demand is inelastic then some form of involuntary load curtailment is needed if capacity is insufficient. In theory this occurs at a value of lost load (VOLL) which is determined by the annual acceptable frequency of load curtailment. If prices go to VOLL enough then the rents earned by generators are theoretically enough to cover their fixed costs on average.

The high values of VOLL needed in order to have low shortage probabilities are not always acceptable to regulators, and so many jurisdictions impose lower price caps. To ensure that shortages do not become more frequent from underinvestment, they institute other mechanisms (capacity markets) to remunerate generators for their fixed and capital costs, which would otherwise be “missing money”. By contrast, in an energy-only market, it is theoretically possible to produce a competitive equilibrium without any missing money by requiring offers at short-run marginal cost and either paying VOLL during shortages or introducing some elasticity in demand using operating reserve margins (see [13]).

There are a number of factors that serve to complicate this simple theory:

1. Indivisible decisions (e.g. from unit commitment) destroy convexity in the economic dispatch problem and so true locational marginal prices are unavailable, and must be approximated. Because New Zealand has very few large thermal plants this is not a major effect. The economic dispatch model “Scheduling Pricing and Dispatch” (SPD) used by the system operator is essentially a linear program, so locational marginal prices are readily obtained from dual variables.
2. Markets are incomplete, meaning that some instruments are not priced in

the market. The classic case occurs when agents compete for transmission assets in the absence of a system operator. Suppose South Island generators could specify some fraction of the HVDC line that they wished to use. Then in competition there are many potential equilibria, in which they end up sharing the line in different proportions. To ensure the optimal proportions requires a single price on the transmission asset, i.e. the difference in nodal prices delivered by the optimal solution of SPD.

A less obvious incompleteness arises from the transfer of water between Lake Tekapo and Lake Pukaki, releases from which are controlled by different companies (Meridian and Genesis). This form of market incompleteness was first identified in the paper by Lino et al. [15], who demonstrated this using a computational example based in Brazil. If all agents are located on different river systems then a centrally planned solution gives rise to system marginal prices that will clear the market with an optimal dispatch if each agent optimizes its own objective using these prices, but inefficiency can result in equilibrium from different agents operating hydro stations on the same cascaded river system. The same issue of different agents sharing a common resource is explored in [9] in their treatment of competitive equilibrium. They show that a market instrument pricing the transfer of water between generating stations is required in order for a market equilibrium to yield the economically efficient solution¹. We shall return to this issue in the discussion of our results.

3. Uncertainty affects equilibrium in several ways. In theory, under very special circumstances it is possible for a Walrasian equilibrium to give a stochastic process of prices with respect to which every agent optimizes its own expected benefit (revenue minus variable cost) with the outcome of maximizing total expected welfare. However, as shown by the examples in [4], the stochastic process of prices that yields an equilibrium might be very complicated with none of the stagewise independence properties that make computing optimal policies relatively easy for generators.

In practice, price formation is more complex. In the short term, the New Zealand system operator provides forecast prices by solving SPD sequentially over future trading periods. Astute generators should not base their dispatch plans entirely on this forecast, but should account for random variation in the price process. How this is done in practice varies, since agents

¹In New Zealand, Genesis and Meridian operate different generating stations on the Waitaki system. Our counterfactual models assume that the generation of these stations is coordinated by contractual arrangements so that they yield the collective benefits that would be delivered by one (perfectly competitive) owner.

have different views of the future. In principle, a market equilibrium delivers prices that integrate these views by matching supply and demand from agents who base commercial decisions on their beliefs. We will assume in this paper, however, that market participants' beliefs about the probability distribution of future demand are the same, and do not change in response to market interactions. (This common view could emerge, for example, from them using the same data and models.) When agents maximize expected benefits this yields an equilibrium stochastic process for prices that can be computed from a social planning model.

Longer-term uncertainty in the New Zealand setting is dominated by so-called "dry-year risk" that occurs when inflows into the hydro lakes are lower than expected, risking an energy shortage. The risks of such a shortage are communicated to the market via an expected marginal water value that is assigned by hydro generators to reflect their opportunity cost of releasing water now for electricity dispatch. This opportunity cost can be estimated from historical price data, or computed under an assumption of perfect competition by solving a large-scale stochastic dynamic program that minimizes expected social cost.

4. Risk aversion changes generators' views of occasional high payoffs. The classical model for risk aversion [10] uses concave utility functions that are applied to random payoffs before taking expectations, so high payoffs are discounted at the margin compared with low payoffs. If agents are risk averse then occasional high payoffs from infrequent shortages might be insufficient to cover the fixed costs of peaking generators that are only dispatched as these episodes are approached. Some of this risk can be traded through derivative contracts such as the swaption contract between Genesis and Meridian that contributes to Huntly's fixed costs when water is plentiful and thermal plant are not dispatched.
5. Perfectly competitive actions of risk-averse agents might not yield a social optimum if risks cannot be traded. Agents in such an equilibrium (if it exists) will estimate marginal water values based on their risk-adjusted view of the future, and their actions in aggregate will yield equilibrium prices that are then used to form these views. Determining these prices for a multistage equilibrium would be very difficult. Furthermore, if we were to seek an equivalent socially optimal plan then it is necessary to integrate individual risk measures into a system risk measure to be optimized. As identified by Newbery and Stiglitz [19] this is not possible for general nonlinear utility functions.

As can be seen from the above discussion uncertainty and risk make it difficult

to equate social welfare maximization with competitive equilibrium. However all is not lost. In the last twenty years a “modern theory of risk” has emerged from the theory of finance, based around the concept of *convex risk measures* and their subclass *coherent risk measures* [2]. Coherent risk measures have very attractive properties for optimization. Unlike most utility functions, coherent risk measures have a translation equivariant property, which means that they can be expressed in nested form, a useful property for stochastic dynamic programming models. This property is exploited in a series of papers [22], [23], [8] that implement nested dynamic coherent risk measures in the SDDP algorithm [21].

Coherent risk measures also give a principled methodology for connecting the optimization of social welfare and competitive equilibrium in a dynamic setting. The two welfare theorems we require are proved in [9] under the assumption that perfectly competitive agents solve risk averse convex optimization problems using (possibly) different coherent risk measures that are reasonably similar (in a sense made precise in the paper), and agents are able to trade risk with each other through a complete market of contingent securities. Under these circumstances, the risk measures of the agents can be integrated into a social risk measure. Essentially the agents swap risky payoffs with each other to remove individual risks until they all agree on what are the worst scenarios, the risk of which cannot be removed. The social planner’s solution then optimizes using this risk measure, so it tries to protect agents against these bad scenarios. The welfare theorems (precisely stated in [9]) are paraphrased as follows:

Theorem 1 *Given an integrated social risk measure, the optimal risk-averse social plan gives a set of actions for each agent and shadow prices, so that each agent’s actions optimize their own risk-adjusted operating surplus at these prices after trading risk with other agents, and the prices clear the markets for energy and risk (i.e. there is a Walrasian partial equilibrium).*

Theorem 2 *Suppose for some stochastic process of prices that each agent’s actions optimize their own risk-adjusted operating surpluses at these prices after trading risk with other agents. If prices for energy and risk clear these respective markets (i.e. there is a Walrasian partial equilibrium) then the actions of all the agents yield an optimal risk-averse social plan for the integrated risk measure.*

These theorems form the foundation for our computational study. Using DOASA (the implementation of SDDP developed by researchers at the Electric Power Optimization Centre as described in [25]) we solve a risk-averse social planning problem using nested coherent risk measures with different levels of risk aversion. According to the theorems above, this risk measure can be interpreted as an integrated risk measure corresponding to electricity market participants who

can trade risk in a complete market of derivative instruments. In this setting the dispatch solution and prices we compute will maximize the risk-adjusted rents of each agent (accounting for returns from risk derivative trading) in a perfectly competitive Walrasian equilibrium.

The assumption of complete markets for risk is arguably unrealistic as it requires that every possible random payoff can be replicated by some portfolio of traded instruments. In practice, however, it is often enough for them to be almost complete. For example, Walrasian partial equilibria in markets with liquidity in contracts for differences appear to be close to socially optimal [14]. So, although we cannot expect such complete markets in practice, we claim that the formulation and solution of risk-averse centrally planned hydrothermal models in our work provides a compelling competitive equilibrium benchmark of the New Zealand wholesale electricity market.

The layout of the paper is as follows. In the next section we describe the physical and institutional details of the New Zealand wholesale electricity market. We then outline the features of the suite of stochastic optimization models that we use in our study. The details of the models and the data that they use are provided in an online companion [12]. In section 4 we apply the models to every trading period in 2017, and compare results between models under different assumptions on participants' levels of risk aversion. The final section draws 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 System Operator (SO). In this market all generated and consumed electricity is traded². 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 several hundred pricing nodes on the national grid. (Although demand side bids are included in the official description of the SO 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 SO take the form of *offer stacks*. These are piecewise-constant functions defining the amount of power

²Small 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.

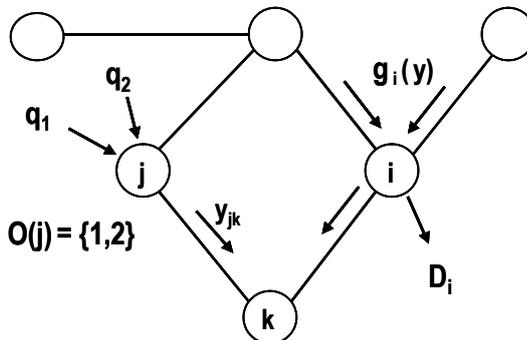


Figure 1: Generic network model illustrating notation

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 $C_m(x)$, where x is the amount offered in MW. In the New Zealand market the generator offer functions C_m are not publicly known at the time of dispatch, but are published the following day. These data are made available as part of the Electricity Market Information (EMI) system supported by the New Zealand Electricity Authority [16].

All the prices in the wholesale electricity market in New Zealand are computed by the SO using a linear programming model called “Scheduling, Pricing, 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]). The New Zealand Electricity Authority supports a publicly available GAMS model called vSPD [18]. The EMI system archives historical input files that when run on the model will reproduce historical prices and dispatch exactly. This enables very precise historical simulations to be run using counterfactual assumptions.

The essential features of SPD can be described mathematically using a DC-Load flow model formulated in 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 . This gives the following market dispatch model:

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

At the optimal solution to MP1, the shadow price π_i on the flow balance constraint at node i defines the locational marginal price. This is the price at

which energy is traded at this location. 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

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

where y_{ik} measures the average flow in the line from node i to node k . In SPD 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 SO's adjustments in load forecasts and system availability.

During the half hour period the SO 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.

As mentioned above SPD (and vSPD) include constraints that ensure voltage support, $N - 1$ security for line failures, and meet requirements for spinning reserve that are dispatched at the same time as energy offers. We assume that voltage support and security constraints are relatively unaffected by the dispatch, and so in any simulation of a historical trading period we assume the constraints that applied at the time.

Spinning reserve does depend on the dispatch, and can have a large effect on prices, so we attempt to account for this in our counterfactual models. Spinning reserve protects the system from a frequency collapse if a large thermal unit or transmission line fails. At the beginning of every run the system operator takes the current dispatch and runs an AC simulation (called Reserve Management Tool or RMT) to estimate the levels of fast response (6-second) and sustained response (60-second) spinning reserve that would be required should a large unit (or the inter-island HVDC link) fail. The outputs of RMT are levels of freely available reserve and automatic load shedding, and the level of extra reserve that must be supplied by market participants in each island, who offer quantities of reserve to the market at prices of their own choosing. Details can be found in [1].

3 The models

Our study makes use of a similar suite of models as defined in [29]. The major difference in comparison with previous work is that we now use a full representation of the New Zealand high voltage transmission system as represented by vSPD. The models we use are:

vSPD:	Dispatch model solved over one trading period;
HydrovSPD:	Daily dispatch model including river chains;
DOASA:	A stochastic planning model solved over one year;

We examine a counterfactual proposal that supposes that the national electricity system is controlled centrally by a system planner who solves DOASA every two weeks in a rolling horizon fashion with updated data. The output from DOASA is used to determine water values for the model HydrovSPD that is solved sequentially over 14 days between solves of DOASA. The outcomes of the model HydrovSPD are then compared with observed outcomes in the wholesale market as computed in vSPD. The details of this process are defined in publications that can be downloaded from the online companion [12]. We digress only briefly here to give an overview of the process.

3.1 HydrovSPD

To investigate the dispatch of hydroelectricity over the course of a day, a national river-chain dispatch and nodal pricing model (HydrovSPD) combines offers from generation plant with river scheduling constraints over 48 half-hour trading periods, $p = 1, 2, \dots, 48$. A diagram of the location of the river chains modelled is shown in Figure 2.

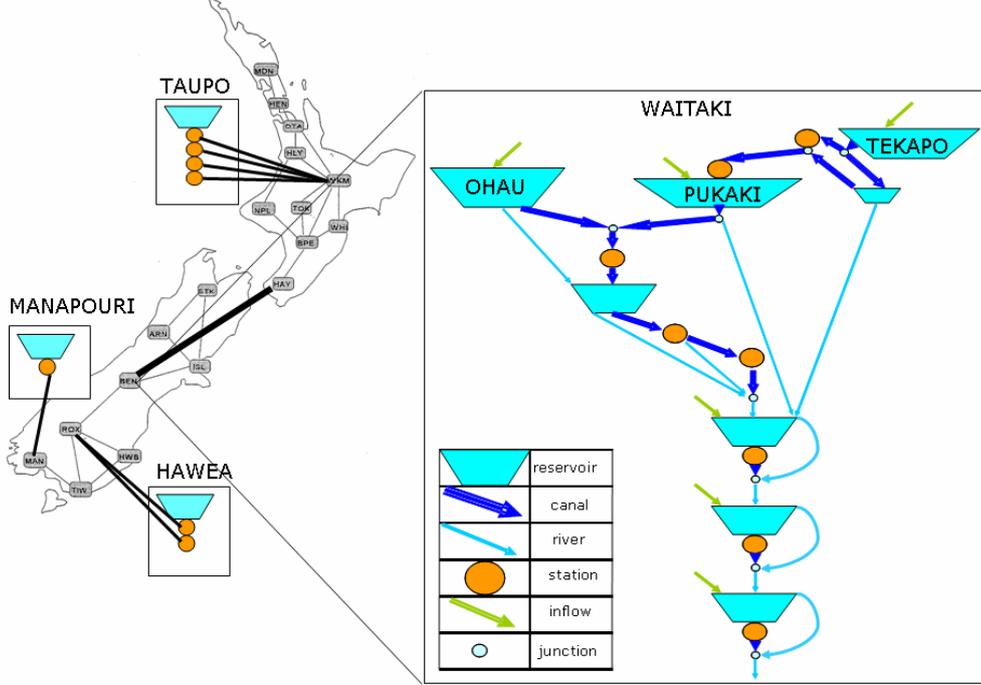


Figure 2: Approximate network representation of New Zealand electricity network showing main hydro-electricity generators

HydroSPD is based on a model developed by Nicholas Porter in a Masters thesis [32]. In the model we discriminate between thermal generation f_m , $m \in \mathcal{F}(i) \subseteq \mathcal{O}(i)$, and hydro generation $\gamma_m h_m$, $m \in \mathcal{H}(i) \subseteq \mathcal{O}(i)$. The parameter γ_m , which varies by generating station m , converts flows of water $h_m(p)$ into electric power. We denote the set of trading periods by $\mathcal{P} = \{1, 2, \dots, 48\}$ ³, and add the argument p to all variables.

The storage in a reservoir or headpond r is denoted by x_r . The initial storage at the start of period $p = 1$ is given by the vector \bar{x} . The water balance constraints in each period are represented by

$$x_r(p+1) = x_r(p) - A_{rm}(h_m(p) + s_m(p)) + \omega_r(p),$$

where $x_r(p)$ is the storage in reservoir r at the start of period p , $s_m(p)$ denotes the spill from above station m in period p , and $\omega_r(p)$ 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

³ \mathcal{P} can have 46 or 50 trading periods on days in which daylight saving changes.

resource consents.) The node-arc incidence matrix A represents all the river-valley networks, and subtracts controlled flows that enter a reservoir from upstream from those that leave a reservoir by spilling or generating electricity. In other words row r of $A(h(p) + s(p))$ gives the total controlled flow out of the reservoir (or river junction) represented by row r , this being the release and spill of reservoir r minus the sum of any immediately upstream releases and spill.

We differentiate between large storage reservoirs $r \in \mathcal{R}$ and small headponds $r \in \mathcal{S}$. We require small headponds to start and end the day 50% full, while the marginal value of water in the large storage reservoirs is calculated using a piecewise linear convex cost-to-go function $\theta(x)$, defined by cutting planes

$$\theta(x) = \max_{k \in \mathcal{K}} \{ \alpha^k + \sum_{r \in \mathcal{R}} \beta_r^k x_r(49) \}.$$

Here the values of α^k and β_r^k are determined from the output of a longer term model. If at some x we have $\theta(x) = \alpha^\ell + \sum_{r \in \mathcal{R}} \beta_r^\ell x_r(49)$ then $-\beta_r^\ell$ defines the marginal value of water in reservoir r at the end of the day. Marginal water values defined by these cutting planes will, after being adjusted by γ , define the counterfactual energy prices $\pi_i(p)$ determined for each trading period and location by solving $\text{HydrovSPD}(\bar{x})$. This is formulated as:

$$\begin{aligned} \min \quad & \sum_{p \in \mathcal{P}} \sum_{i \in \mathcal{N}} \sum_{m \in \mathcal{F}(i)} \phi_m f_m(p) + \theta \\ \text{s.t.} \quad & g_i(y(p)) + \sum_{m \in \mathcal{F}(i)} f_m(p) + \sum_{m \in \mathcal{H}(i)} \gamma_m h_m(p) = D_i(p), \quad [\pi_i(p)] \\ & \quad \quad \quad i \in \mathcal{N}, p \in \mathcal{P}, \\ & \text{vSPD constraints: e.g. security, spinning reserve,} \\ & \quad \quad \text{voltage support, loop flow, ramping} \\ & 0 \leq f_m(p) \leq a_m, \quad m \in \mathcal{F}(i), i \in \mathcal{N}, p \in \mathcal{P}, \\ & x_r(p+1) = x_r(p) - A_{rm}(h_m(p) + s_m(p)) + \omega_r(p), \quad r \in \mathcal{R} \cup \mathcal{S}, p \in \mathcal{P}, \\ & 0 \leq h_m(p) \leq b_m, \quad 0 \leq s_m(p) \leq c_m, \quad m \in \mathcal{H}(i), p \in \mathcal{P}, \\ & 0 \leq x_r(p) \leq w_r, \quad r \in \mathcal{R} \cup \mathcal{S}, \quad p \in \mathcal{P}, \\ & \alpha^k + \sum_{r \in \mathcal{R}} \beta_r^k x_r(49) \leq \theta, \quad k \in \mathcal{K}, \\ & x_r(49) \geq 0.5w_r(1), \quad r \in \mathcal{S}, \quad x_r(1) = \bar{x}_r, \quad r \in \mathcal{R} \cup \mathcal{S}. \end{aligned}$$

Observe that we include spinning reserve offers and constraints in HydrovSPD . A difficulty here is that we do not have access to the appropriate reserve parameters

(from RMT) for the counterfactual models, so we simply choose reserve requirements to be the same as historical levels, and make the same historical reserve offers in HydrovSPD at zero cost.

3.2 The medium-term hydro model: DOASA

To investigate the dispatch of hydroelectricity over the course of a year, a hydrothermal 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 $C_t(x, \omega)$ be the minimum expected fuel cost to meet electricity demand in weeks $t, t+1, \dots, T$, when reservoir storage $x(t) = x$, week t 's inflow is known to be ω . Here $C_t(x, \omega)$ is the optimal solution value of the mathematical program:

$$\begin{aligned}
P_t(x, \omega): \quad & \min \quad \sum_{i \in \mathcal{N}} \sum_{m \in \mathcal{F}(i)} \phi_m f_m(t) + \mathbb{E}_\eta[C_{t+1}(x(t+1), \eta)] \\
\text{s.t.} \quad & g_i(y(t)) + \sum_{m \in \mathcal{F}(i)} f_m(t) + \sum_{m \in \mathcal{H}(i)} \gamma_m h_m(t) = D_i(t), \quad i \in \mathcal{N}, \\
& x(t+1) = x - A(h(t) + s(t)) + \omega, \\
& 0 \leq f_m(t) \leq a_m, \quad m \in \mathcal{F}(i), i \in \mathcal{N}, \\
& 0 \leq h_m(t) \leq b_m, \quad 0 \leq s_m(t) \leq c_m, \quad m \in \mathcal{H}(i), \\
& 0 \leq x_r(t) \leq w_r, \quad r \in \mathcal{R}, \\
& y \in Y,
\end{aligned}$$

where η represents the random inflow in week $t+1$. To solve P_t , we use the DOASA code [25] which is based on the SDDP technique of Pereira and Pinto [21]. This approximates $\mathbb{E}_\eta[C_{t+1}(x(t+1), \eta)]$ using a polyhedral function defined by cutting planes that is updated using samples of the inflow process. Since each week t has its own distribution of inflows we will henceforth denote this random variable $\omega(t)$, so $\mathbb{E}_\eta[C_{t+1}(x(t+1), \eta)] = \mathbb{E}[C_{t+1}(x(t+1), \omega(t+1))]$.

The DOASA model uses weekly stages. A calendar year is divided into 52 weeks. A plan year is typically a year of 52 weeks with the starting week chosen to be a particular week in the calendar year. Historical inflows are sampled from a file that records weekly inflows as described below. The New Zealand electricity system is represented as shown in Figure 3.

Weekly demand is represented by a load duration curve with three blocks. These are called peak, off-peak and shoulder. We have chosen peak hours to

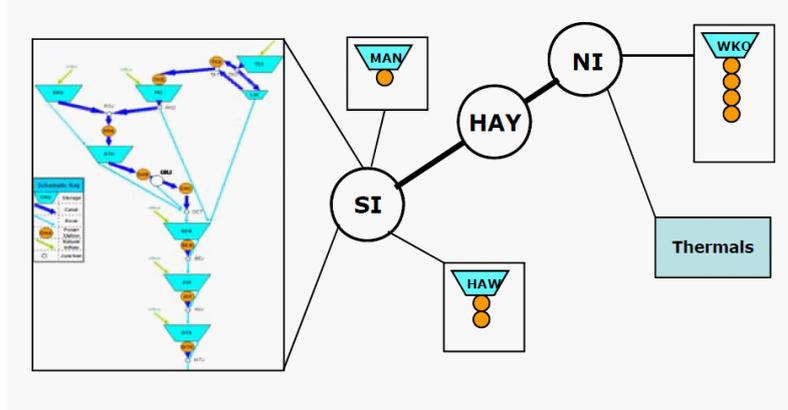


Figure 3: The 3 node transmission network and major generators in DOASA.

be 6am-8am and 6pm-8pm weekdays, shoulder hours to be 8am-6pm and 8pm-10pm weekdays, and offpeak hours to be the other hours in the week. The total demand in MW in each node i is then averaged over these trading periods to get a total demand rate $D_i(b, t)$ for each block. The energy requirement in node i for each block b in week t will be its duration $T(b, t)$ times the average demand rate $D_i(b, t)$ for this block.

The choice of what data to include in demand is a delicate matter. Publicly available demand figures (e.g. those in the EMI Data Set [16]) make various assumptions about what embedded generation and demand is included. These must be carefully studied to ensure that demand is not overlooked or double counted. The DOASA model aggregates demand to three locations (SI, HAY, NI) representing the South Island, the lower North Island, and the upper North Island, and allows transfers between these regions limited by line capacities. This means that aggregating demand in each region will ignore the intra-regional losses, implying that the regional totals of historical demand will underestimate the true demand to be met by generation and net imports to SI, HAY, and NI. So some inflation of total demand is needed in DOASA to account for these losses. If the aggregation is carried out geographically then ignoring line losses might also bias the generation mix in the dispatch towards geographically close (yet possibly electrically distant) plant.

The aggregation (into regions SI, HAY, and NI) of historical dispatch of the large generators can be used as a proxy for the demand adjusted for losses. We ignore the generation supplied by small generators as long as demand is adjusted for this, and compute the total generation of large generators in each region (SI, HAY, and NI) in each trading period using vSPD, and then add the net import minus export of power through transmission lines joining the region to its adjacent regions. The result will give the half-hourly demand in the region satisfied by

large generators and transfers between the regions. This is then transformed into load blocks for each week and used as the demand to be met by large generators. Of course this means that small fixed generators should have their generation fixed at zero in the DOASA input data (as demand has effectively been reduced by these values).

A precise description of the demand calculation is as follows. Index all the generators represented in DOASA by $g \in \mathcal{G}$. Denote the regions (SI, HAY, and NI) by indices S , H and N . The first step in computing demand is to determine what vSPD nodes lie in each region. The boundaries of these regions can be defined somewhat arbitrarily, although Cook Strait is one obvious boundary between HAY and SI.

Let \mathcal{N}_S , \mathcal{N}_H , and \mathcal{N}_N define the vSPD nodes corresponding to each region. Based on the transmission lines in vSPD, we define four sets of transmission variables

$$\begin{aligned}\mathcal{L}_{SH} &= \{\text{transmission lines directed from a node in } \mathcal{N}_S \text{ to a node in } \mathcal{N}_H\} \\ \mathcal{L}_{HS} &= \{\text{transmission lines directed from a node in } \mathcal{N}_H \text{ to a node in } \mathcal{N}_S\} \\ \mathcal{L}_{HN} &= \{\text{transmission lines directed from a node in } \mathcal{N}_H \text{ to a node in } \mathcal{N}_N\} \\ \mathcal{L}_{NH} &= \{\text{transmission lines directed from a node in } \mathcal{N}_N \text{ to a node in } \mathcal{N}_H\}.\end{aligned}$$

The generators $g \in \mathcal{O}(i)$ in each region $i \in \{S, H, N\}$ are treated as if they are in a single location. We compute the total generation in each region in each trading period in each day of the plan year, by running vSPD with the GDX file for the trading periods in that day. This gives us 365 days, each containing 48 periods (except for leap years and daylight savings days.). Consider a particular period p on a specific day in the plan year. The DOASA demand for this period is estimated as follows.

For each generator g let

$q_g(p)$ = generation in MWh returned by vSPD using historical data for period p

and

$f_l(p)$ = the line flow variable (MWh) for line l as computed by vSPD for period p .

Here $f_l(p)$ does not include the losses incurred, half at each endpoint of the line. For these we define the flow leaving the node at the start of the line to be $f_l^+(p)$ and the flow arriving at the end of the line to be $f_l^-(p)$. Irrespective of the sign of $f_l(p)$ it follows that

$$\begin{aligned}f_l^+(p) &= f_l(p) + \frac{\alpha}{2}(\text{flow loss}) \\ f_l^-(p) &= f_l(p) - \frac{\alpha}{2}(\text{flow loss}),\end{aligned}$$

where we choose $\alpha = 0$ for a model with lossless flow and $\alpha = 1$ to represent losses. Then we compute the total generation in a region i to be

$$q_i(p) = \sum_{g \in \mathcal{O}(i)} q_g(p).$$

Now adding flows we get

$$\begin{aligned} f_{SH}(p) &= \sum_{l \in \mathcal{L}_{SH}} f_l^+(p) - \sum_{l \in \mathcal{L}_{HS}} f_l^-(p) \\ f_{HS}(p) &= \sum_{l \in \mathcal{L}_{HS}} f_l^+(p) - \sum_{l \in \mathcal{L}_{SH}} f_l^-(p) \\ f_{HN}(p) &= \sum_{l \in \mathcal{L}_{HN}} f_l^+(p) - \sum_{l \in \mathcal{L}_{NH}} f_l^-(p) \\ f_{NH}(p) &= \sum_{l \in \mathcal{L}_{NH}} f_l^+(p) - \sum_{l \in \mathcal{L}_{HN}} f_l^-(p). \end{aligned}$$

Observe that

$$\begin{aligned} & f_{SH}(p) + f_{HS}(p) + f_{HN}(p) + f_{NH}(p) \\ &= \sum_{l \in \mathcal{L}_{SH} \cup \mathcal{L}_{HS} \cup \mathcal{L}_{HN} \cup \mathcal{L}_{NH}} (f_l^+(p) - f_l^-(p)) \\ &= \alpha \sum_{l \in \mathcal{L}_{SH} \cup \mathcal{L}_{HS} \cup \mathcal{L}_{HN} \cup \mathcal{L}_{NH}} \text{flow loss}. \end{aligned}$$

The regional demand in period p is now estimated for each region to be

$$\begin{aligned} d_S(p) &= q_S(p) - f_{SH}(p) \\ d_H(p) &= q_H(p) - f_{HS}(p) - f_{HN}(p) \\ d_N(p) &= q_N(p) - f_{NH}(p). \end{aligned}$$

Observe that with this definition, the national demand in period p is

$$\begin{aligned} \sum_{k \in \{S,H,N\}} d_k(p) &= \sum_{k \in \{S,H,N\}} q_k(p) - f_{SH}(p) - f_{HS}(p) - f_{HN}(p) - f_{NH}(p) \\ &= \sum_{k \in \{S,H,N\}} q_k(p) - \alpha (\text{total flow loss from inter-regional flow}). \end{aligned}$$

The rationale behind this choice is that our backtest is intended to compute an optimal dispatch to meet some observed demand. The demand used in vSPD makes various assumptions about embedded generation and wind that are often

difficult to verify, especially for past years. One approach is to backtest the allocation of generation amongst the large generators. If transmission losses are ignored in DOASA (the default) then we choose $\alpha = 0$ in the above analysis which yields national demand in period p equal to $\sum_{k \in \{S,H,N\}} q_k(p)$. DOASA will possibly reallocate generation amongst the generators to be cheaper to meet this demand. This might give different transfers between regions than those observed in vSPD.

If transmission losses between regions are modelled in DOASA then we choose $\alpha = 1$, which yields national demand in period p equal to $\sum_{k \in \{S,H,N\}} q_k(p)$ minus national losses. Now DOASA will possibly reallocate generation to be cheaper to meet this demand. This might give different transfers between regions than those observed in vSPD, to not only allow cheaper generation but potentially give lower transmission losses. Once $d_k(p)$ has been computed for the three regions for every trading period, we take $\sum_{k \in \{S,H,N\}} d_k(p)$ for each p in this week, and allocate the demand to one of three blocks: peak, shoulder, and offpeak. The assignment of trading period to block is made a priori and fixed. Thus computing the energy (MWh) in each block in each week is obtained by summing $d_k(p)$ over periods p corresponding to the block.

In meeting demand, in case of supply shortages, load shedding (in MW) is allowed at high costs. The costs depend on the type of customers and amount of reduction (in \$/MWh). Load in each node is divided into three sectors to represent different types of customers, which are industrial, commercial and residential, and each sector has some distribution in each island. The default proportions are the proportions of consumption in 2015 adjusted to higher commercial and residential proportions in the North Island due to a denser population, and to a higher industrial proportion in the South Island due to an aluminium smelter. Each sector is then divided into three segments of load reduction. We assume the first 5% of load reduction can be made at a modest cost, while the next 5% of load reduction incurs a higher price, and the remaining 90% represents unplanned interruption of power supply at VOLL. The VOLL values for the industrial sector are set to be lower than the other two and the VOLL values increase over segments in each sector. We assume that up to 10% reduction in load can be achieved at a relatively low cost, but the value of unplanned interruption is very high (\$10,000/MWh)⁴.

The DOASA model assumes that six reservoirs, Manapouri, Hawea, Ohau, Pukaki, Tekapo and Taupo, can store water from week to week. The release of

⁴The value of \$10000/MWh is open to some debate. The NEM in Australia applies a VOLL that is indexed to inflation. In 2018-19 the value was \$14,500 [34]. Our choice of \$10,000 is based on the capital cost of approximately \$1m/MW for peaking plant [20] that would be required 5 hours per year over 20 years.

this water through generating stations is controlled⁵. The hydroelectric stations in other parts of the system are treated as run-of-river plant with limited intra-week flexibility.

Each DOASA model we solve has 52 weekly stages. We specify a cost-to-go function for the end of week 52, based on price levels and total storage observed in historical years at this time of the year. This assumption places a caveat on the prices we compute in our counterfactual simulations. A more principled (though computationally intensive) approach would be to compute steady-state water values using an infinite horizon version of DOASA. It is important also to note that we assume inflows to the main catchments are stagewise independent. These are sampled from the historical weekly inflow series available on the EMI site [16]⁶. Full details of the DOASA model for this study can be found in the online companion [12] to this paper.

The solution to $P_1(x_1, \omega(1))$ 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 for any day d at this stage by setting $\bar{x} = x(t)$ and solving $\text{HydrovSPD}(\bar{x}, d)$ (i.e. HydrovSPD using initial reservoir storage \bar{x} and data from day d) in which the constraints

$$\alpha^k + \sum_{r \in \mathcal{R}} \beta_r^k x_r(49) \leq \theta, \quad k \in \mathcal{K},$$

at the end of the last trading period in day d are determined by the cuts defining $\mathbb{E}[C_{t+1}(x(t+1), \omega(t+1))]$. In our experimental setup we use the cuts defining $\mathbb{E}[C_2(x(2), \omega(2))]$ to define the final cost function on each day d that we run $\text{HydrovSPD}(\bar{x}_r, d)$ in the first week of the fortnight. For the second set of seven runs of $\text{HydrovSPD}(\bar{x}_r, d)$, we use the cuts defining $\mathbb{E}[C_3(x(3), \omega(3))]$. Then DOASA is re-solved using the storage at the end of the fortnight. Each DOASA solve computes 1000 cuts. In summary the experimental procedure is as follows:

Given reservoir levels $x(1)$ solve a 52-week hydrothermal scheduling problem using DOASA.

1. Set $t = 1$.
2. Solve a hydrothermal scheduling problem over weeks $\{t, \dots, t + 51\}$ using DOASA.

⁵Because of complicated environmental restrictions, we treat Lake Manapouri differently from other storage lakes. See section 4.

⁶In any year y we select inflows for each catchment in the years from $y - 35$ to $y - 1$ as equally likely random outcomes in each week. Thus for any year we have 35 (vector) outcomes per stage giving a stagewise independent scenario tree for DOASA of 35^{51} scenarios.

3. Set $\bar{x} = x(t)$.
4. For $d = 1$ to 7,
 - (a) Select α^k and β^k , $k \in \mathcal{K}$, from the cut intercepts and slopes approximating $\mathbb{E}_\omega[C_{t+1}(x(t+1), \omega(t+1))]$;
 - (b) Solve HydrovSPD(\bar{x}, d);
 - (c) Set $\bar{x} = x_r(49)$.
5. For $d = 8$ to 14,
 - (a) Select α^k and β^k , $k \in \mathcal{K}$, from the cut intercepts and slopes approximating $\mathbb{E}[C_{t+2}(x(t+2), \omega(t+2))]$;
 - (b) Solve HydrovSPD(\bar{x}, d);
 - (c) Set $\bar{x} = x_r(49)$.
6. Set reservoir levels to \bar{x} , set $t = t + 2$, and go to step 2.

DOASA also allows us to compute risk-averse policies with varying levels of risk aversion. Risk is modelled using a nested dynamic risk measure (see [22]), in which the one-step risk measure is a convex combination of the expectation and worst-case outcome of future fuel and shortage cost. In other words we use the one-step risk measure

$$\rho(Z) = (1 - \lambda)\mathbb{E}[Z] + \lambda\mathbb{W}[Z]$$

where $\lambda \in [0, 1)$, Z represents the random future cost, and

$$\mathbb{W}[Z] = \max\{Z(\omega)\}.$$

The dynamic version uses a nested form of ρ , where the risk averse certainty equivalent of a random stream of costs, say Z_1, Z_2, Z_3 , is computed using a nested formulation, which would be $\rho(Z_1 + \rho(Z_2 + \rho(Z_3)))$ in this example. A straightforward procedure for implementing this within SDDP algorithms is described in [23]. If there are M scenarios, the measure $\rho(Z)$ is equivalent to weighting all scenarios with equal probability $\frac{(1-\lambda)}{M}$ except for the most expensive scenario which receives weight $\frac{1}{M}(M\lambda - \lambda + 1)$. In our experiments we compare a risk-neutral policy ($\lambda = 0$) with a mildly risk averse policy ($\lambda = 0.1$) and a highly risk-averse policy ($\lambda = 0.3$). A value of $\lambda = 0.1$ implies that the decision maker each week believes with probability 0.1 that the worst inflow observed in this week in the last 35 years will occur, which is about four times more likely than the model

with $\lambda = 0.0$. A value of $\lambda = 0.3$ makes this worst-case inflow about 11 times more likely than the risk-neutral probability.

The DOASA model also assumes that inflows are stagewise independent, but modified by an adjustment to account for stagewise correlation. The weekly variance in the empirical inflow distribution is inflated, so that the sum of variances over two (assumed independent) consecutive weeks is the same as the variance of this sum estimated from historical data. This is called *inflow spreading* by some modellers. We call it Dependent Inflow Adjustment (DIA) and describe it more fully in [26]. On its own, DIA may not be sufficient to successfully represent stagewise dependence, for decisions are not conditioned on previous inflow observations as they would be under stagewise dependence assumptions. Risk aversion can serve as a proxy for this effect, as sequences of low inflows are treated as being more likely than what one would expect from stagewise independence. Modelling risk aversion using nested coherent risk measures can also be interpreted from the perspective of a decision maker who is unsure of the true probability distribution of future inflows, and adopts a distributionally robust model as described in [24].

4 Market comparison

We now describe a set of experiments that were carried out using data from 2017. Given costs per MWh of gas, diesel, and coal generation it is possible to compute the cost of fuel used by the large thermal generators (Contact and Genesis) to generate the electricity dispatched by the wholesale market in each historical 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, Stratford, and Whirinaki). These are the only generators that we allow to offer energy within our model.

As mentioned above, we treat Manapouri differently from other storage lakes. This is because Manapouri has complicated nonconvex environmental constraints on its operation. Ignoring these constraints leads to water releases that might

violate the legal requirements placed on the system, thereby underestimating the true cost of the dispatch. To deal with this we constrain Manapouri generation in DOASA (by using a simple decision rule to determine its release) but impose its historical generation levels in HydrovSPD. Observe that a water value for Manapouri is still computed in DOASA even though its release is fixed in HydrovSPD to historical levels. Even though we have fixed small generation to its historical level, some caution must be exercised in fixing too much hydro generation at historical levels, since this affects the value of stored water⁷.

In reporting costs, our measure will be the cost of fuel burned by the five thermal plants. The fuel used in a thermal power station is coal, natural gas or diesel, as shown in Table 1. Coal is supplied from stockpiles that are restocked under long-term contracts. Coal costs are assumed to be constant at \$4/GJ. Natural gas is supplied by take-or-pay contracts. It is assumed that in social planning the supply can be secured and the costs are wholesale prices. The quarterly average prices of natural gas for wholesale use are available from the Ministry of Business Innovation and Employment (MBIE) [17]. The quarterly average prices of diesel for commercial use in [17] are used as the costs of diesel. The quarterly average prices are converted into real dollars in December 2015 and the costs of CO₂ emissions (based on the current CO₂ price expressed in 2015 terms⁸) are added. This gives the fuel and carbon cost of coal, diesel and gas as shown in Table 2. The short-run marginal cost (SRMC) for any plant can be obtained by multiplying the heat rate (see Table 1) by the fuel and carbon cost from Table 2, and adding a variable operations and maintenance cost. These SRMC values are similar to those assumed by other authors (e.g. [7, page 6, Table 2]).

Power station	Heat rate (GJ/MWh)	Fuel
Huntly main 1-4	10.3	coal
Huntly e3p	7.2	natural gas
Huntly peaker	9.8	natural gas
Otahuhu B	7.45	natural gas
Stratford peakers	9.5	natural gas
Taranaki Combined Cycle	7.6	natural gas
Whirinaki	11	diesel

Table 1: Thermal power stations and heat rates

⁷The counterfactual model in the Commerce Commission report authored by Wolak [35] fixed all hydro generation to historical levels which bounded marginal water values at thermal fuel costs.

⁸The costs of CO₂ permits are automatically adjusted in our model for regulatory relaxations (e.g. 1 for 2 schemes) that were applied in some months.

Real fuel costs		Coal (EPOC)	Gas(MBIE)	Diesel	Coal(FNZC)	Gas(FNZC)
Year	Quarter	(2015\$/GJ)	(2015\$/GJ)	(2015\$/GJ)	(2015\$/GJ)	(2015\$/GJ)
2017	1	5.53	6.6	19.58	7.56	8.36
2017	2	5.48	5.48	20.14	7.5	6.91
2017	3	5.54	6.32	19.95	7.57	7.99
2017	4	5.66	6.48	23.78	7.69	8.18

Table 2: Fuel and CO2 costs for thermal generation in December 2015 NZ dollars (Source [17], [11]).

To enable a fair comparison with market outcomes, we have de-rated stations at which plant have been removed for planned maintenance. The weekly de-rating of a generator is taken by default to be the outage amount given in the POCP database [30]. If no data are provided for an offering generator in [30], we de-rate its plant capacity in a given week of the year by the difference between its nominal capacity and the average total offer quantity made by the plant in the same week in previous years. The schedule in POCP defines the starting and end time of scheduled maintenance for generators, which includes the offering generators and all small and run-of-river generators that we consider as fixed (e.g. Tokaanu, Rangipo and Waikaremoana). The HVDC line capacity is treated as fixed in DOASA, but will be assigned its vSPD value for HydrovSPD. All de-rating is deterministic, so random outages can be anticipated by the social planner.

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 as shown in Table 3.

(\$/MWh)	Up to 5%	Up to 10%	VOLL	North Island	South Island
Industrial	\$ 1,000	\$ 2,000	\$ 10,000	0.36	0.58
Commercial	\$ 2,000	\$ 4,000	\$ 10,000	0.27	0.17
Residential	\$ 2,000	\$ 4,000	\$ 10,000	0.37	0.24

Table 3: Load reduction costs (\$/MWh) and proportions of each load that is industrial, commercial, and residential load.

The last two columns of Table 3 show the proportion of load of each type in each island. This shows that (rounded to the nearest percentage) 58% of South Island load is industrial, 17% is commercial, and 24% is residential. The costs (in NZ\$/MWh) of shedding load are also shown in Table 3. 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). Therefore if, for example, load in the South Island was 1000 MW, we could shed up to 5% of 580 MW at \$1000/MWh and at \$2000/MWh, we could shed 5% of 410MW (170MW commercial and 240 MW residential) plus a further 5% of 580MW industrial load.

5 Experiments

We now present the results of applying DOASA and HydrovSPD to data from the calendar year 2017. This backtest was conducted for both risk-neutral ($\lambda = 0$) and risk-averse ($\lambda = 0.1$ and $\lambda = 0.3$) settings, and for two sets of fuel-price data (provided by MBIE [17] and FNZC[11] respectively). We first discuss results for MBIE fuel costs and then present results using FNZC fuel costs.

5.1 MBIE fuel costs

The results from DOASA and HydrovSPD assuming MBIE fuel prices are summarized in Table 4 shown below. The rest of this section will compare the solutions in detail.

5.1.1 Short-run costs

The counterfactual models release more water than the historical dispatch, and generate more hydroelectricity, resulting in a decrease in thermal cost (from fuel and CO₂ costs). The differences in final reservoir levels can be costed using the risk neutral cuts from DOASA. These values are shown as cost-to-go figures in the row labelled “Final storage cost”. When these are taken into account the cost savings from the counterfactual models are about \$36m, \$36m, and \$22m. In both risk averse and risk neutral counterfactuals there was a small amount of load shedding, needed to avoid HydrovSPD returning an infeasible solution. In real-time dispatch, any infeasibility in a solve of SDP is typically resolved by relaxing reserve or security constraints. We did not allow this in the counterfactual. Although the vSPD dispatch solutions in the pricing runs archived on the EMI site by the Electricity Authority are all feasible, they do not match actual historical generation, partly because of “constrained-on” dispatch and partly because river chains are block dispatched in real time. Thus vSPD dispatch solutions can turn out to be infeasible for HydrovSPD when river-chain flows are accounted for.

Total thermal generation in Table 4 increases as social risk aversion increases. (The High-risk-aversion figure of 5736 GWh matches the Historical figure by coincidence; the mix of thermal generation is different as can be seen from the different thermal cost figures of \$311m and \$309m.) As risk aversion increases, the amount of hydro generation decreases and the amount of water released also decreases, but hydro generation in all counterfactual solutions is higher than historical hydro generation. The final storage cost is computed using risk-neutral water values corresponding to that week in the final run of DOASA. As residual storage increases with risk aversion the final storage cost (i.e. cost-to-go) decreases. The

total generation figures are different across the four solutions because of different transmission losses which are higher in the counterfactual solutions than in the historical solution. This appears to come from higher levels of South Island generation in the counterfactual solutions leading to more South-North transmission flows. Counterfactual rents on the HVDC link are higher than historical rents.

2017 MBIE fuel cost		Historical	Risk neutral	Mild risk aversion	High risk aversion
Thermal cost	(\$m)	309.15	239.04	260.03	311.02
Final storage cost	(\$m)	340.08	374.54	353.34	316.02
Total cost	(\$m)	649.23	613.58	613.37	627.04
Demand violation	(MWh)	0.00	11.88	11.88	35.02
Hydro generation	(GWh)	19999.25	21186.64	20834.80	20066.48
Thermal generation	(GWh)	5736.26	4610.67	4952.09	5736.35
Total generation	(GWh)	25735.50	25797.31	25786.89	25802.83
Transmission losses	(GWh)	1267.68	1335.14	1322.50	1341.75
OTA average price	(\$/MWh)	78.23	64.97	67.62	77.82
HAY average price	(\$/MWh)	77.13	62.36	65.22	76.34
BEN average price	(\$/MWh)	76.55	50.94	55.16	62.62
Revenue	(\$m)	2857.51	2116.90	2242.81	2596.82
Cost	(\$m)	309.15	239.04	260.03	311.02
Rent	(\$m)	2548.36	1877.86	1982.77	2285.80
HVDC rent (S-N)	(\$m)	3.57	19.94	20.72	27.08

Table 4: Summary of outcomes from EMBER backtest on 2017 assuming MBIE gas costs. Generation revenue, costs, and rents are computed for large generation plant only.

5.1.2 Revenue and rent

The results of total revenue and Ricardian rent (in 2015 dollars) are shown in Table 4. The rents and costs here are calculated using generation quantities and nodal prices only for those stations that are owned by the five largest electricity companies. Moreover they account only for costs incurred during the calendar year 2017; future costs from final reservoir levels (“Final storage costs”) are not accounted for in rent calculations. Since historical wholesale prices are higher than counterfactual prices, revenues and rents are higher than counterfactual values. The differences in Ricardian rent from historical values for 2017 are \$671m, \$566m and \$263m for the three counterfactual solutions respectively. Even in the high risk aversion case, when counterfactual final reservoir levels exceed historical ones, we see that there is \$263m difference. It should be stressed that these wholesale rents do not represent total electricity company earnings as they are not adjusted for retail costs, contracts, or fixed costs. And, as discussed in the Introduction, to correspond to a competitive equilibrium, the distribution of rents between companies must be adjusted to account for risk trading between them.

5.1.3 Generation

We now turn our attention from the summary in Table 4, to discuss the details of differences between the historical solution and the counterfactual solutions. The differences in generation between historical levels and the counterfactual models are plotted in Figures 4, 5, and 6. In all the plots in this paper we represent historical values using purple lines, risk-neutral values using blue lines, mild risk-averse values ($\lambda = 0.1$) using red lines, and high risk-averse values ($\lambda = 0.3$) using green lines. Hydro generation in the counterfactuals is increased in January and in the second half of 2017. As shown in Figure 5, thermal generation is decreased in these periods. Total thermal generation increases with risk aversion to low inflows. The differences in weekly thermal cost are shown in Figure 8. Weekly total generation (Figure 6) is the same in all three solutions except in January where increased hydro generation leads to more transmission losses (see Figure 7). For all solutions, the weekly total generation minus transmission losses gives the same sequence of numbers (corresponding to weekly historical demand minus fixed generation).

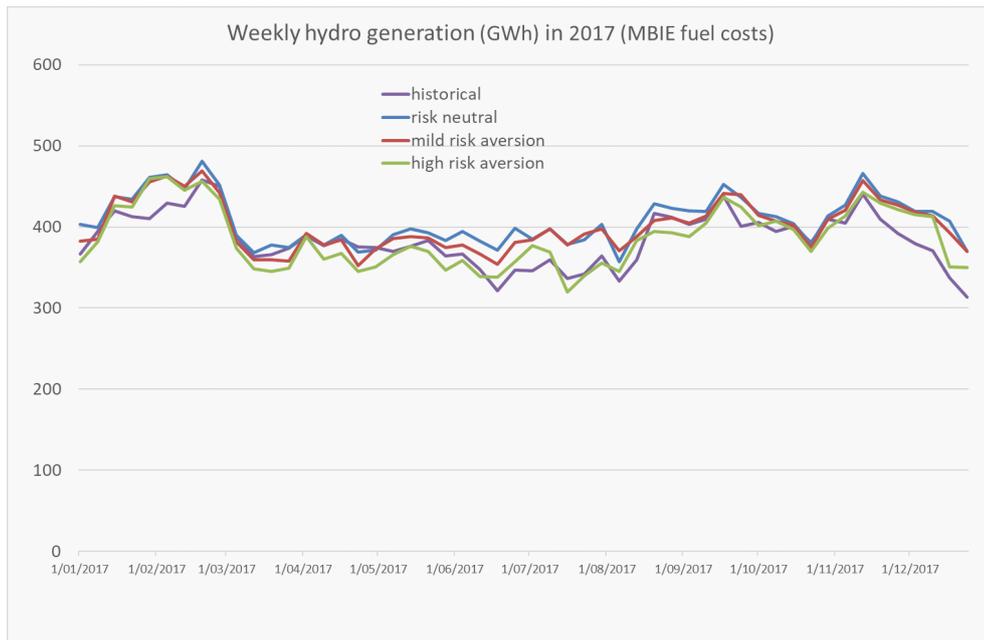


Figure 4: Weekly hydro generation (GWh) in 2017 for MBIE cost model.

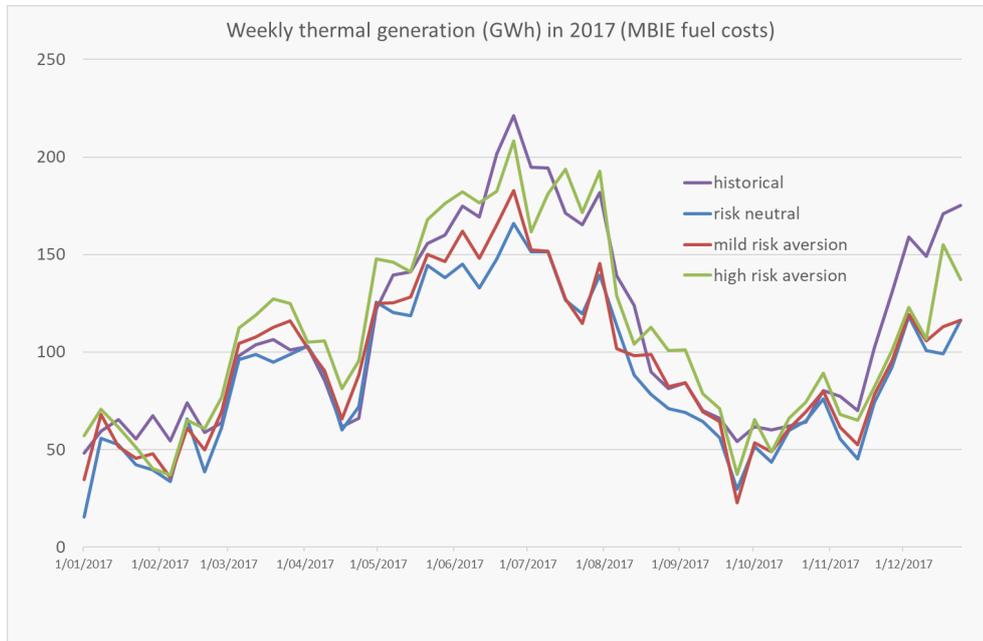


Figure 5: Weekly thermal generation (GWh) in 2017 for MBIE cost model.

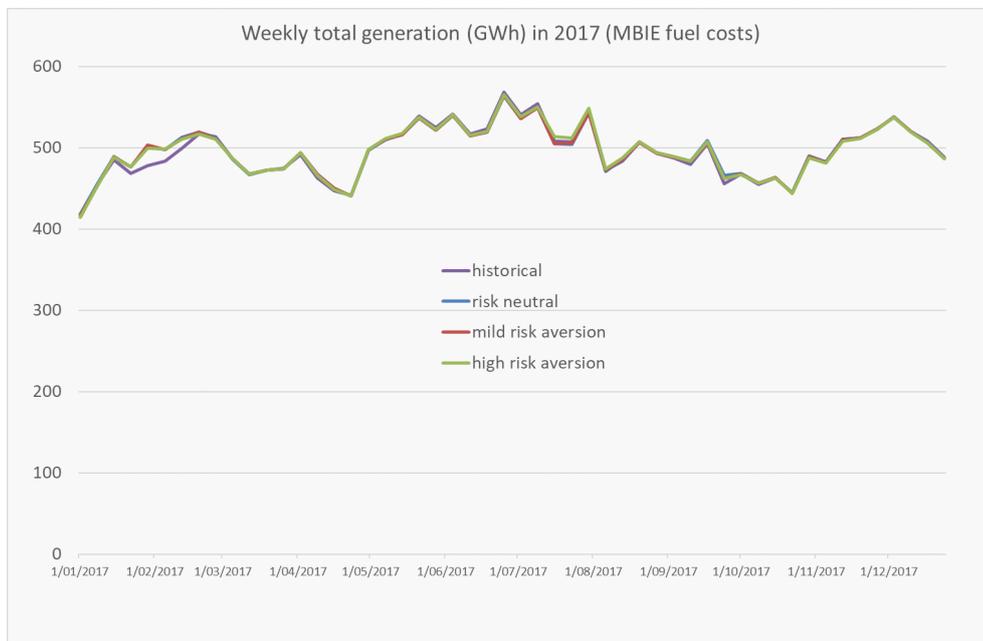


Figure 6: Weekly total generation (GWh) in 2017 for MBIE cost model.

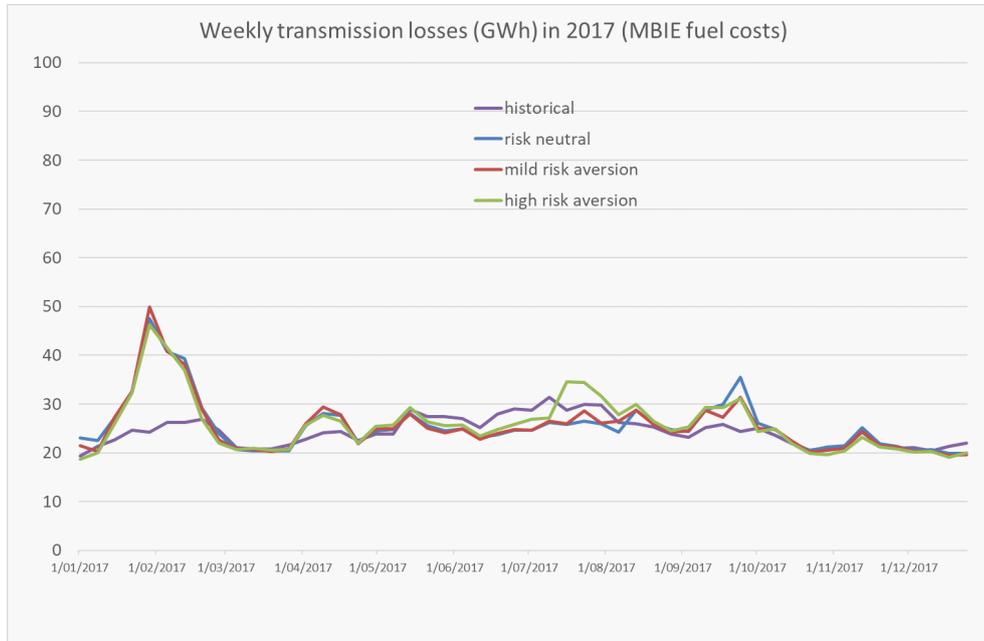


Figure 7: Weekly transmission losses (GWh) in 2017 for MBIE cost model.

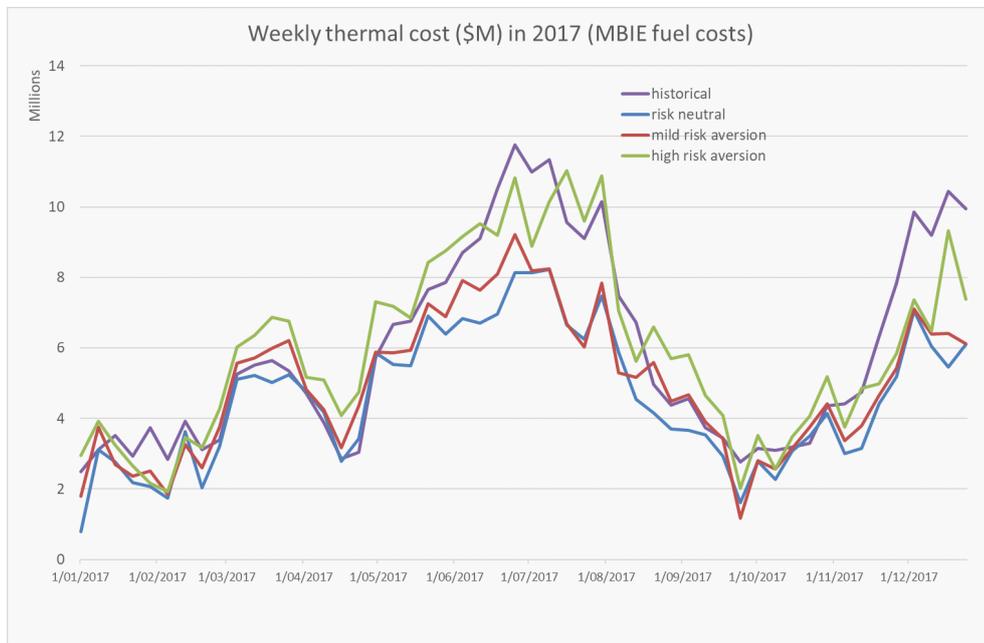


Figure 8: Weekly costs of thermal fuel and carbon emissions (\$M) in 2017 for MBIE cost model.

5.1.4 Reservoir storage

The differences in reservoir storage between historical levels and the counterfactual models are plotted for North and South Islands in Figure 9. These are broken down into the storage in individual lakes in Figures 10, 11, 12, and 13. The counterfactual reservoir levels increase with risk aversion, corresponding to an increase in thermal generation and a decrease in hydro generation (Table 4) as risk levels increase.

It is interesting to compare counterfactual reservoir levels against historical ones. The Hawea levels shown in Figure 10 are adjusted from those reported in HydrovSPD. In the historical dispatch, the Clyde and Roxburgh dams were forced to spill by virtue of a resource consent that requires this when the level of Lake Wanaka exceeds a specified threshold. This regulation is not modelled in HydrovSPD, and so we spill water from Hawea around the Clyde and Roxburgh dams to match the historical spill (Since releases from Lake Wanaka are not controllable, the level of this lake follows its historical trajectory.)

As shown in Figure 13, the level of Lake Taupo is kept higher in the counterfactual solutions than in the historical dispatch. Despite this, the counterfactual solutions generate more hydroelectricity from the Waikato river than observed historically. Table 5 shows the figures.

	Historical	Risk neutral	Mild risk aversion	High risk aversion
Generation (GWh)	5022	5413	5353	5248

Table 5: Generation from Waikato river.

There are several possible reasons for these differences. HydrovSPD is clairvoyant in each day and as reported in [32] this improves the efficiency of river chains through improved coordination with other dispatch. Tributary inflows are estimated from daily measured data at two locations using fixed proportions, and we fix intraday inflows at their average levels which may allow a more efficient dispatch. Moreover the generation output of each station on the river is computed using the fixed conversion factors listed in the generator database on the EMI site [16]. These conversion factors when compared with historical generation can sometimes overstate the average efficiency of the generating plant that must deal with varying head levels and inefficient running. HydrovSPD also does not model delays in water flows, assuming that water is instantaneously available downstream when released. High levels of spill are scheduled in some periods in the counterfactual to transfer water, when required, to downstream stations.

These modelling features of HydrovSPD allow it to extract more energy from released water than might be practically possible⁹.

The levels of Lake Pukaki and Lake Tekapo are shown in Figures 11 and 12. The generation from Tekapo A and B stations is higher in the counterfactual solutions than in the historical dispatch. As a result Lake Tekapo is drawn down more aggressively, hitting its minimum operating level in August (except in the high risk aversion case). The Ohau stations are also run harder in the counterfactual solutions than they are in the historical dispatch. As a result total Waitaki generation is higher in the counterfactual model during the winter months than in the historical dispatch. Compared with historical levels, this leads to an increase in transmission flow from South to North, and more congestion on the HVDC line from Benmore to Haywards.

In practice the operators of hydro stations using Lake Tekapo water (Genesis Power) would be reluctant to allow Lake Tekapo storage levels to drop as low as the counterfactual (blue, red and green) trajectories in Figure 12. As a company trading in energy, such low lake levels would be risky. The counterfactual solutions, however, operate the whole Waitaki river chain (indeed the whole electricity system) as a single entity. In this social planning model, the risks of low levels of Lake Tekapo are balanced by the advantages of extra generation from stations on the lower part of the Waitaki river system. These stations are operated by Meridian Energy, a different company who are in practice exposed to their own set of risks. The lack of coordination between Genesis and Meridian in the historical dispatch contributes to a solution with a higher overall cost. Ideally some of this lost coordination benefit could be captured and shared between the firms by a suitable payment for water transfer between the parties as outlined in [9].

⁹We plan future implementations of HydrovSPD that model head effects, reach delays and higher resolution inflows. However calibration of these models is very difficult without full access to generator data (which is proprietary).

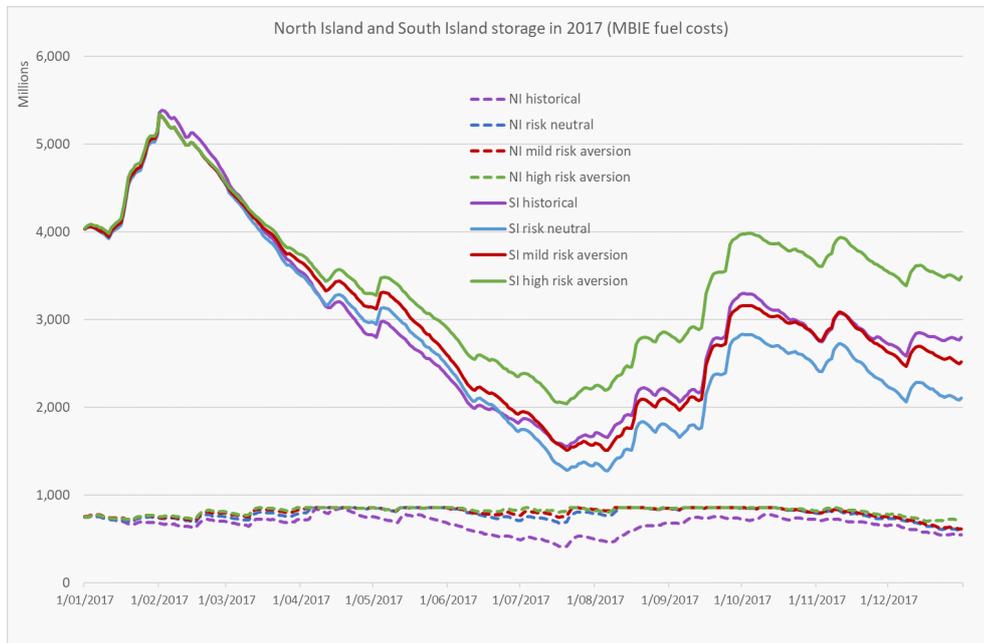


Figure 9: North Island and South Island reservoir storage (Mm³) in 2017 for MBIE cost model.

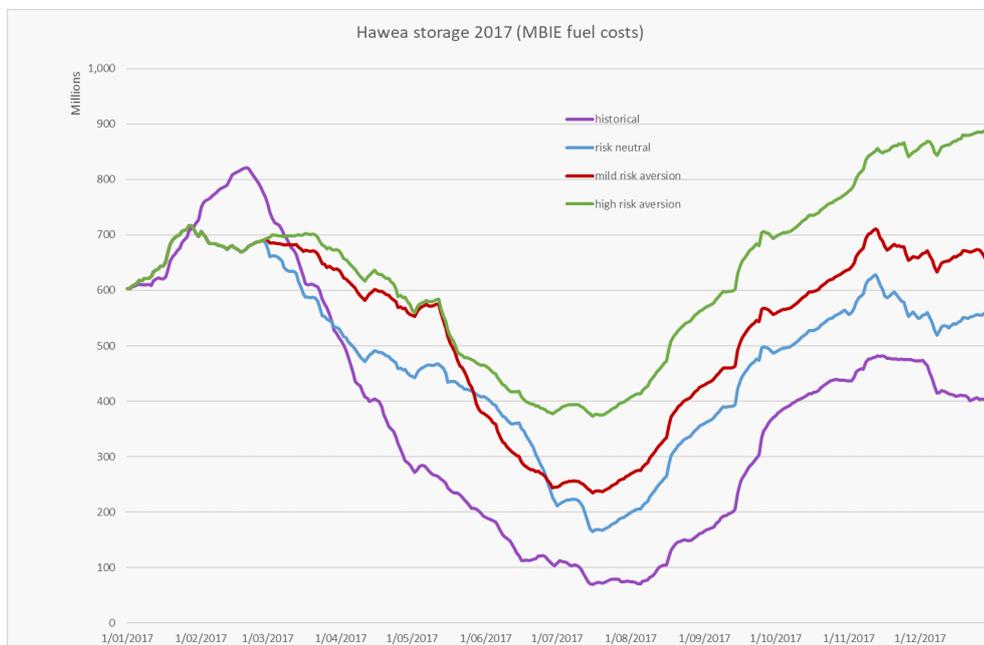


Figure 10: Lake Hawea storage (Mm³) in 2017 for MBIE cost model.

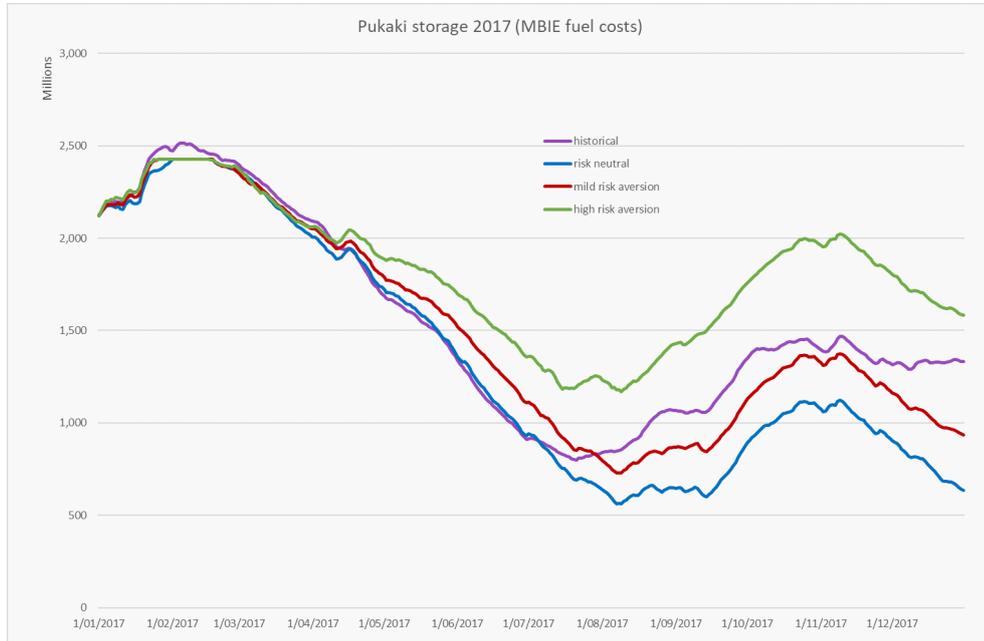


Figure 11: Lake Pukaki storage (Mm³) in 2017 for MBIE cost model.

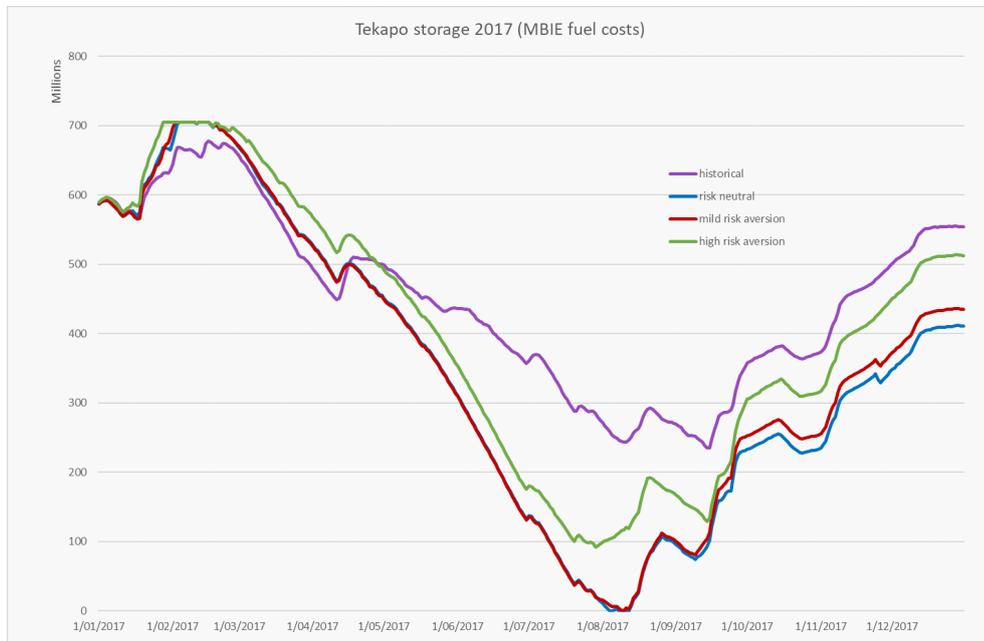


Figure 12: Lake Tekapo storage (Mm³) in 2017 for MBIE cost model.

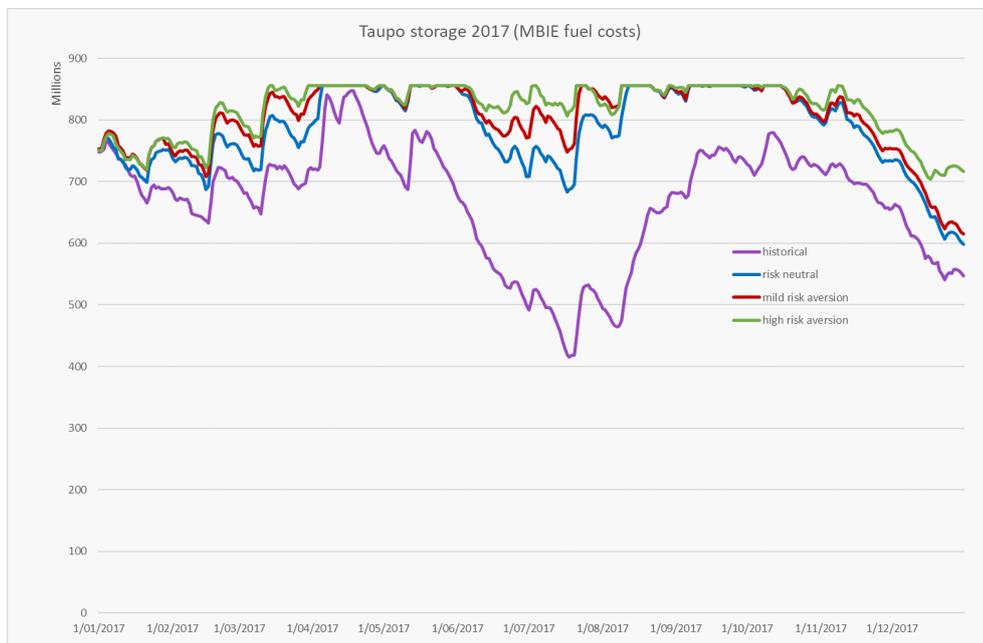


Figure 13: Lake Taupo storage (Mm³) in 2017 for MBIE cost model.

5.1.5 Nodal prices

Plots of prices that are generated by the counterfactual policies are shown below. As mentioned above HydrovSPD can sometimes return an infeasible solution, resulting in very high prices (\$500,000/MWh) in the counterfactual model to signal the infeasibility. In the market dispatch model SPD, such an outcome would be resolved by the system operator relaxing a reserve requirement, until the dispatch was feasible. We model this in the counterfactual as follows. For any trading periods that have some nodal price of at least \$10,000/MWh (our choice of VOLL) we check if there is any demand violation. If so then any nodal prices above \$10,000/MWh are set to that level. If there is no demand violation in the trading period then all nodal prices in that period are capped at \$1,000/MWh. We chose this value after observing that 99.8% of historical trading periods without demand violation have prices below \$1,000/MWh.

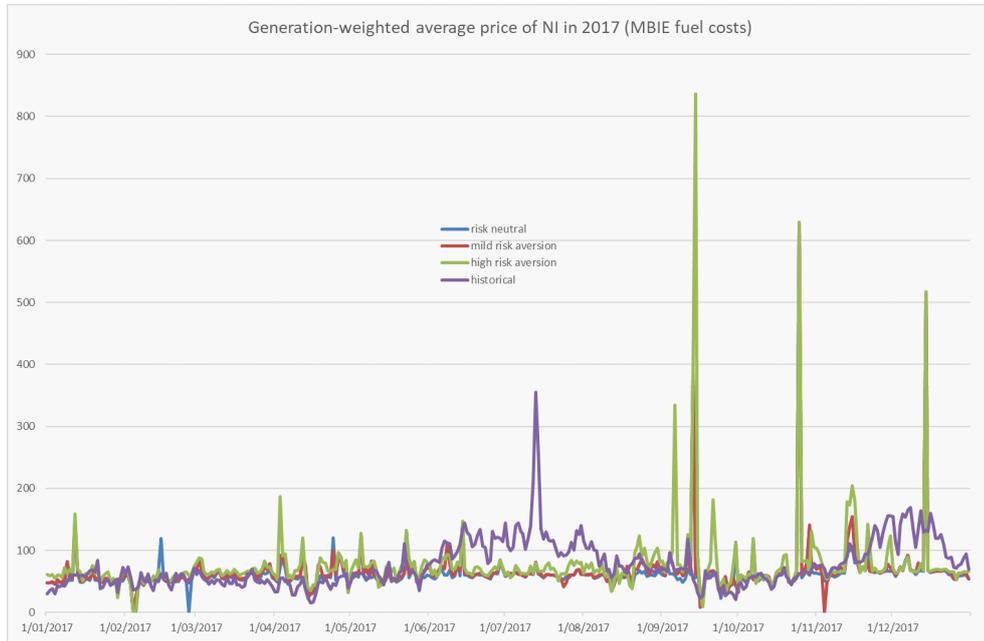


Figure 14: Generation-weighted daily average prices (\$/MWh) over North Island grid injection points (for large generators) in 2017 for MBIE cost model.

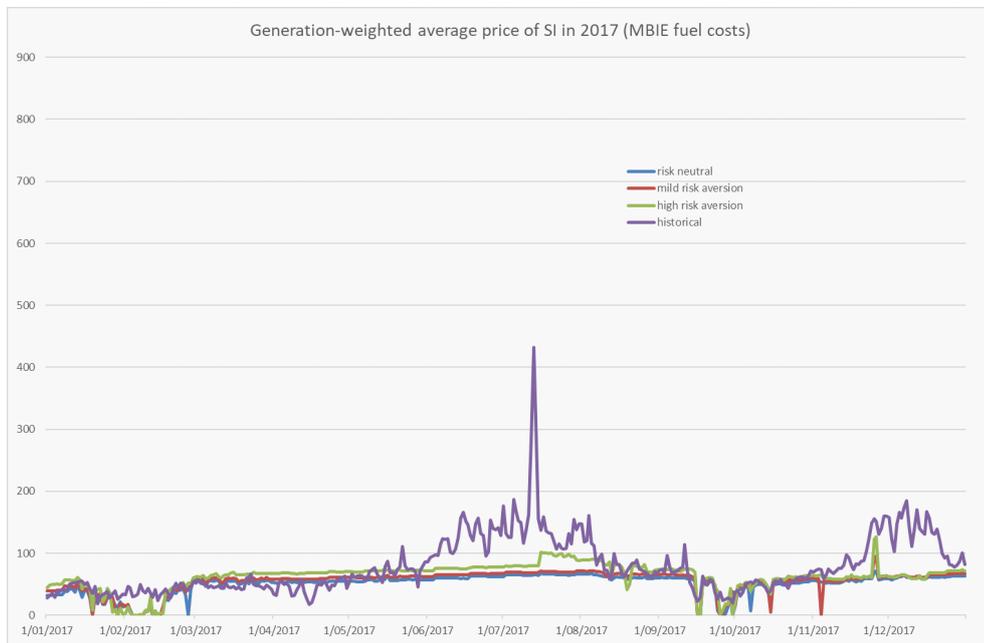


Figure 15: Generation-weighted daily average prices (\$/MWh) over South Island grid injection points (for large generators) in 2017 for MBIE cost model.

After capping prices, as above, we compute generation-weighted average prices (GWAPs) over 2017 as shown in Figure 14 (North Island) and Figure 15 (South Island) for both risk-neutral and risk-averse solutions. The most striking similarity between these plots is the degree to which historical South Island GWAPs follow those in the North Island.

This is confirmed by the time-weighted average prices at Benmore, Haywards and Otahuhu as shown in Table 4. These estimates are approximately equal in the historical dispatch, but the average Benmore price is \$11 lower, \$10 lower, and \$14 lower in the three counterfactual models. The difference implies that South Island electricity prices in 2017 were higher than perfectly competitive levels in 2017. Daily time-weighted average prices for historical and counterfactual models are plotted for Otahuhu, Haywards and Benmore nodes in Figures 16, 17, and 18 respectively. They show a close match in historical daily prices at Haywards and Benmore that is not observed in the counterfactual prices.

The price differences between the Benmore and Haywards nodes are higher on average in the counterfactual models. This results in higher South-North transmission rentals collected from periods when HVDC congestion occurs. These are of the order of \$20 million. Figure 19 shows the accumulation of these transmission rents over the year for all four solutions. During winter when South-North flows are small the rental accrual is flat. At the start and end of the year, there are substantial differences in accrual of transmission rents when Benmore prices are lower than Haywards as shown in Figures 17 and 18.

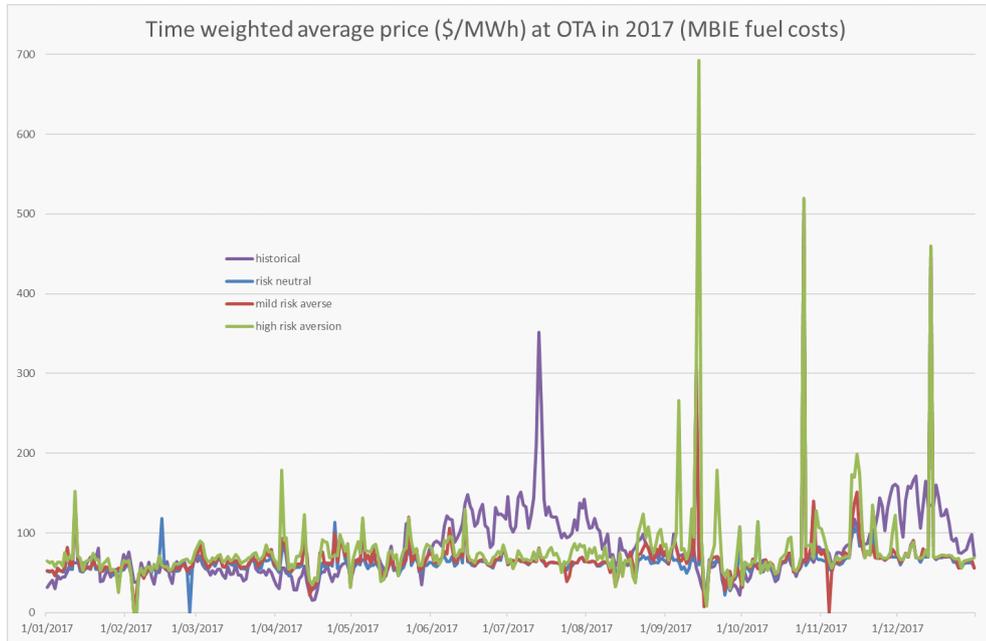


Figure 16: Time-weighted daily average prices (\$/MWh) at OTA in 2017 for MBIE cost model.

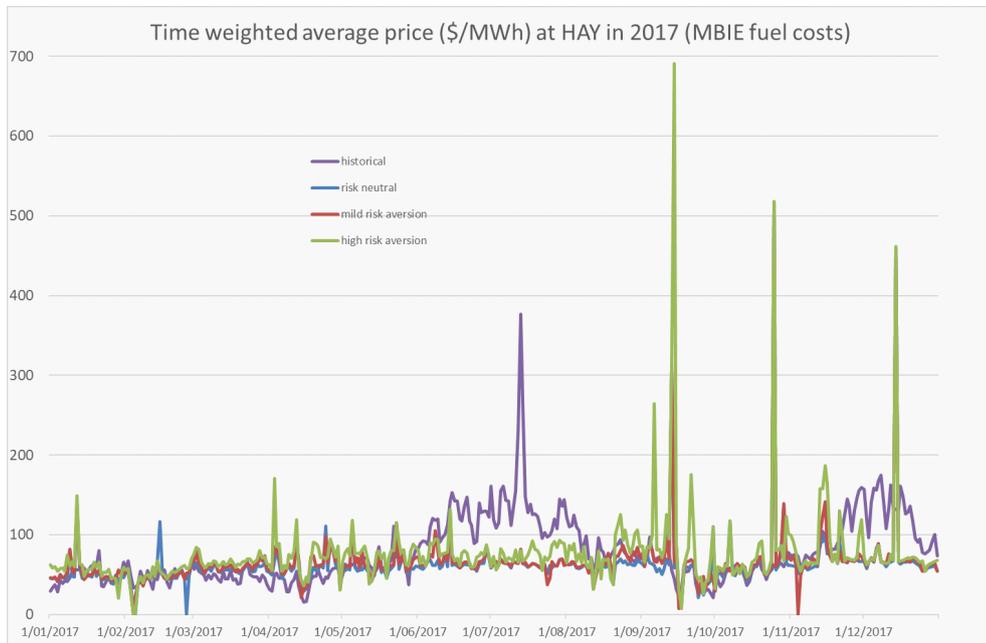


Figure 17: Time-weighted daily average prices (\$/MWh) at HAY in 2017 for MBIE cost model.

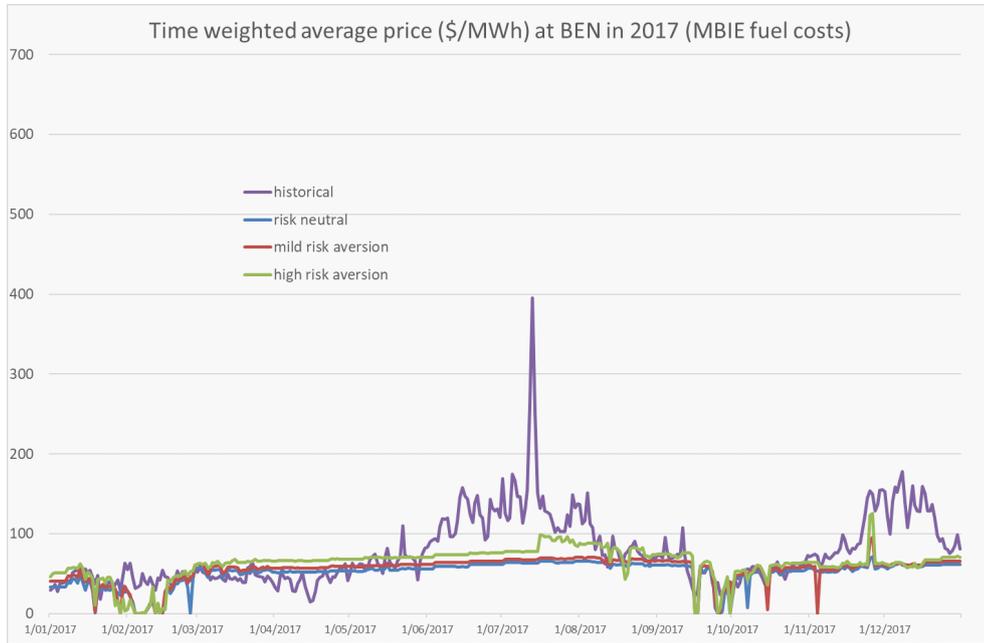


Figure 18: Time-weighted daily average prices (\$/MWh) at BEN in 2017 for MBIE cost model.

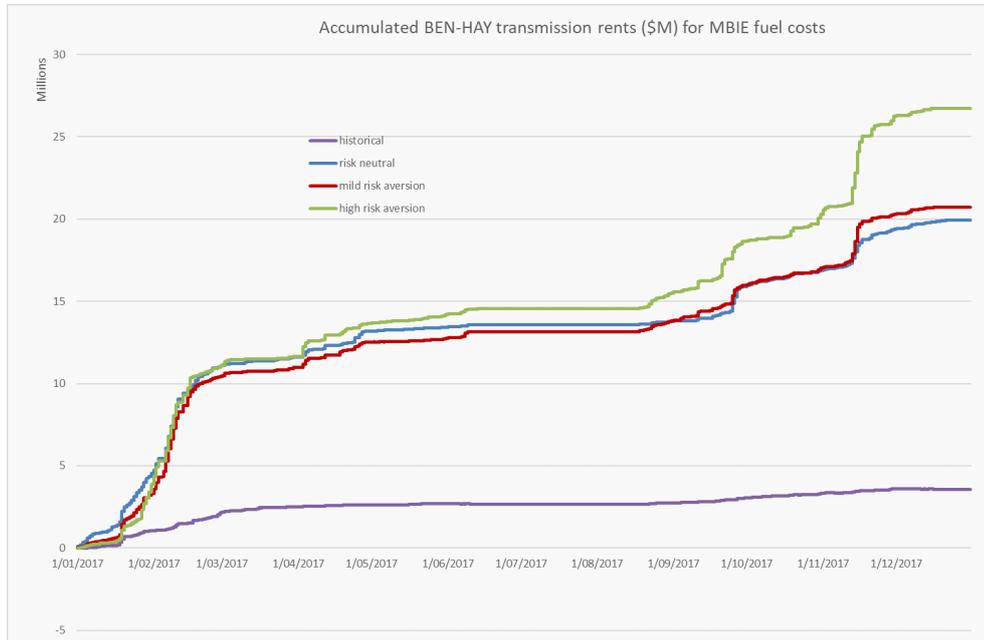


Figure 19: Cumulative HVDC rentals over 2017 for MBIE cost model.

5.2 FNZC fuel costs

A second set of experiments were carried out with gas costs provided by First NZ Capital Securities Ltd (FNZC)¹⁰ that result in higher variable costs as shown in the last two columns of Table 2. The use of gas and coal by thermal plants is complicated by take-or-pay contracts. If a thermal generator holds a gas contract for more than they need then they might offer below a nominal fuel cost in order to burn the excess at an apparent loss. On the other hand a generator who is short of gas might regard the opportunity cost of gas to be higher than what was paid in a take-or-pay contract. Gas cost is further complicated by ownership. In 2017 the Genesis group had a 46% stake in the Kupe field. This makes reported payments for gas for the group significantly lower than they would be otherwise, where the cost of gas as an operating expense for Huntly power station is interpreted as an opportunity cost, i.e. the foregone value of not selling it elsewhere. The results for FNZC fuel prices are summarized in Table 6 shown below. The rest of this section will compare the solutions in detail.

¹⁰We are grateful to Nevill Gluyas for providing us with these estimates of fuel prices.

2017 FNZC fuel cost		Historical	Risk neutral	Mild risk aversion	High risk aversion
Thermal cost	(\$m)	385.88	261.93	285.76	334.55
Final storage cost	(\$m)	458.41	472.25	450.09	413.79
Total cost	(\$m)	844.30	734.18	735.85	748.34
Demand violation	(MWh)	0.00	7.47	30.62	31.99
Hydro generation	(GWh)	19999.25	20829.20	20458.27	19756.54
Thermal generation	(GWh)	5736.26	4979.95	5351.37	6060.65
Total generation	(GWh)	25735.50	25809.15	25809.65	25817.19
Transmission losses	(GWh)	1267.68	1344.27	1345.67	1359.76
OTA average price	(\$/MWh)	78.23	80.15	86.09	96.43
HAY average price	(\$/MWh)	77.13	76.60	83.08	94.59
BEN average price	(\$/MWh)	76.55	67.14	72.96	82.18
Revenue	(\$m)	2857.51	2663.39	2902.55	3275.90
Cost	(\$m)	385.88	261.93	285.76	334.55
Rent	(\$m)	2471.62	2401.46	2616.79	2941.35
HVDC rent (S-N)	(\$m)	3.57	24.90	27.34	32.68

Table 6: Summary of simulation results with FNZC gas costs

5.2.1 Short-run costs

As before, the risk neutral and mild risk-averse counterfactual models release more water than the historical dispatch, and generate more hydroelectricity, resulting in a decrease in thermal cost (from fuel and CO₂ costs). Including the differences in final reservoir levels costed using the risk neutral cuts from DOASA, we obtain total cost savings from the counterfactual models of about \$110m, \$108m, and \$96m. In both risk-averse and risk-neutral counterfactuals there was a small amount of load shedding, needed to avoid HydrovSPD returning an infeasible solution.

Total thermal generation in Table 6 increases as social risk aversion increases. As risk aversion increases the amount of hydro generation decreases and the amount of water released also decreases. Observe that the high-risk-aversion counterfactual solution uses less hydro than the historical solution, and substantially more thermal generation (334 GWh). Prices, revenues and rents for this counterfactual all exceed the historical values by substantial amounts. In all cases, counterfactual rents on the HVDC link are substantially higher than historical rents.

5.2.2 Revenue and rent

The results of revenue, cost and Ricardian rent (in 2015 dollars) are shown in Table 6. It is remarkable that under the FNZC fuel cost assumptions, historical rents are *lower* than those that would be earned under perfect competition, except if agents are risk-neutral when they are \$70m more. In the risk neutral case the counterfactual revenue is nearly \$200m lower than historical but the fuel and CO₂ bill for generators is \$124m smaller, which gives them Ricardian rents that are

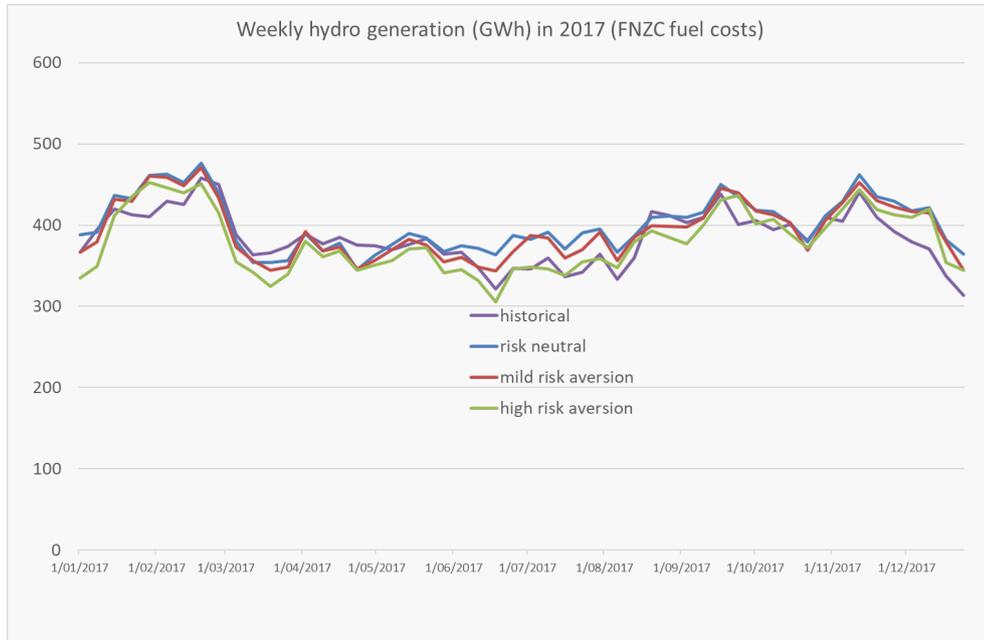


Figure 20: Weekly hydro generation (GWh) in 2017 for FNZC cost model.

close to historical levels.

5.2.3 Generation and storage

In the interests of completeness we now present plots for the counterfactual solutions obtained with FNZC costs that match the plots from the previous section.

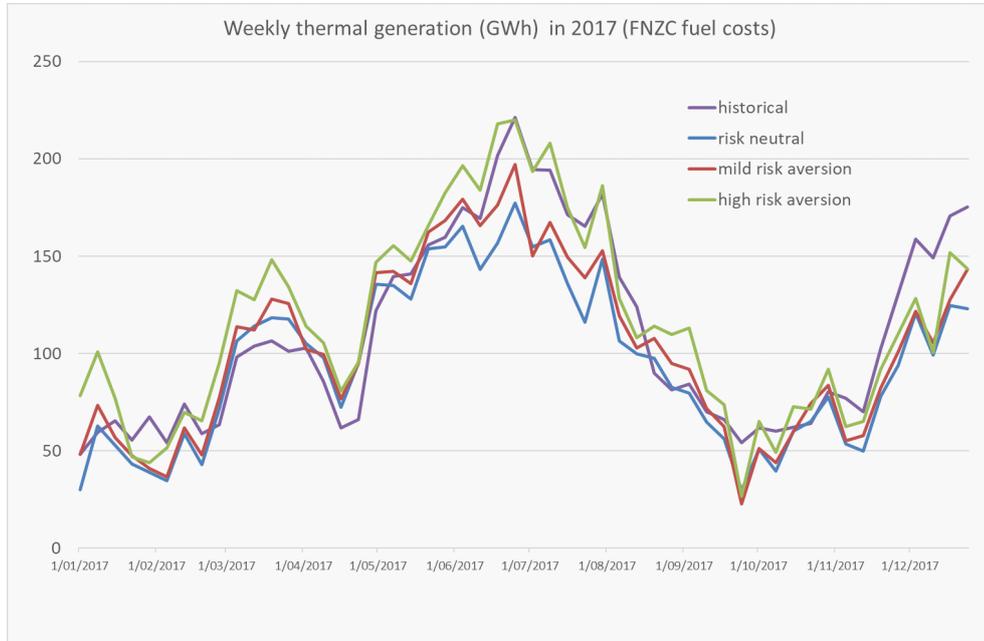


Figure 21: Weekly thermal generation (GWh) in 2017 for FNZC cost model.

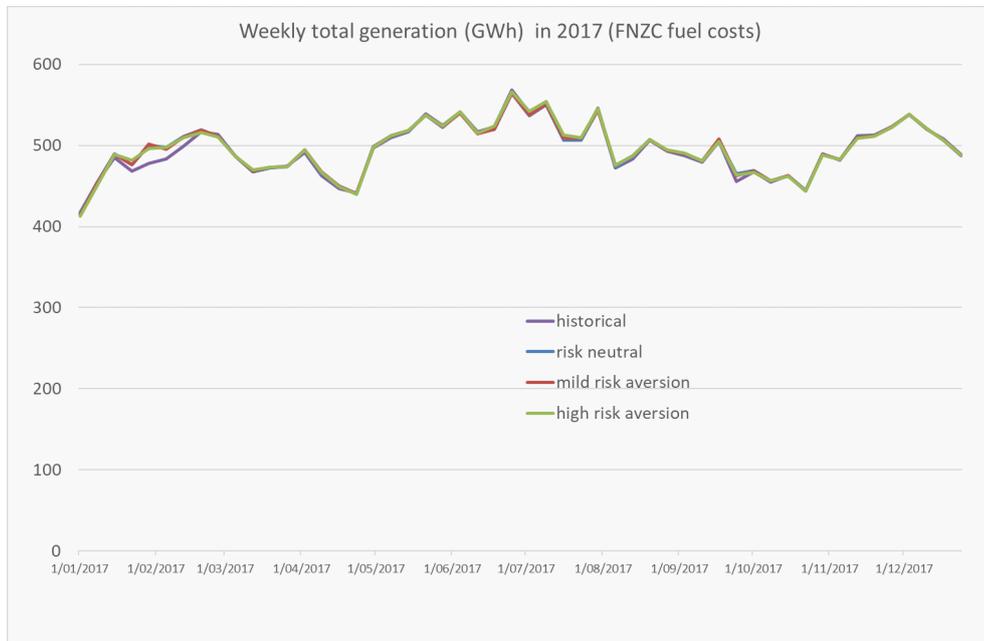


Figure 22: Weekly total generation (GWh) in 2017 for FNZC cost model.

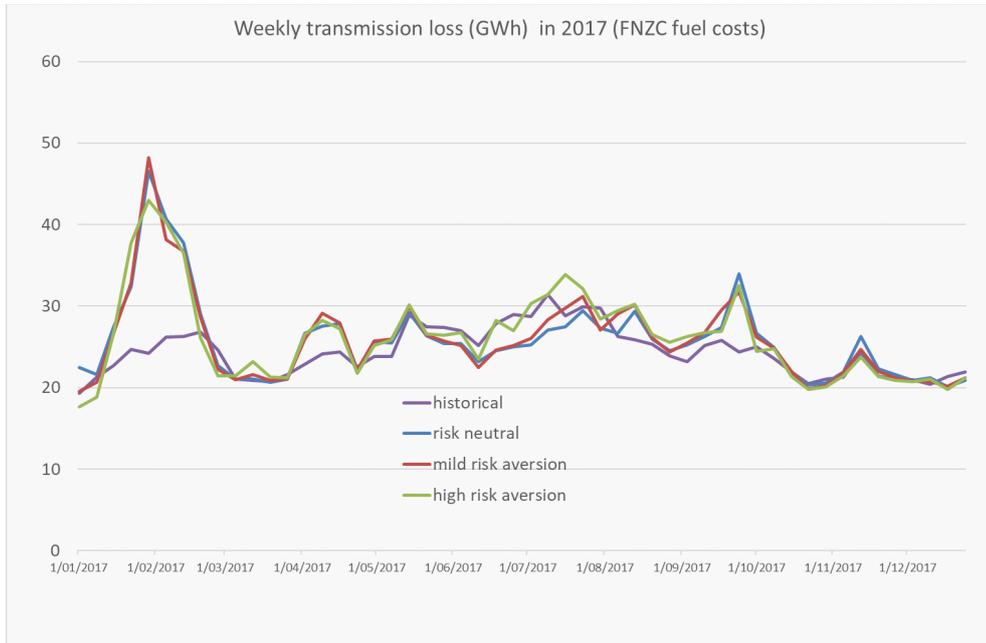


Figure 23: Weekly transmission losses (GWh) in 2017 for FNZC cost model.

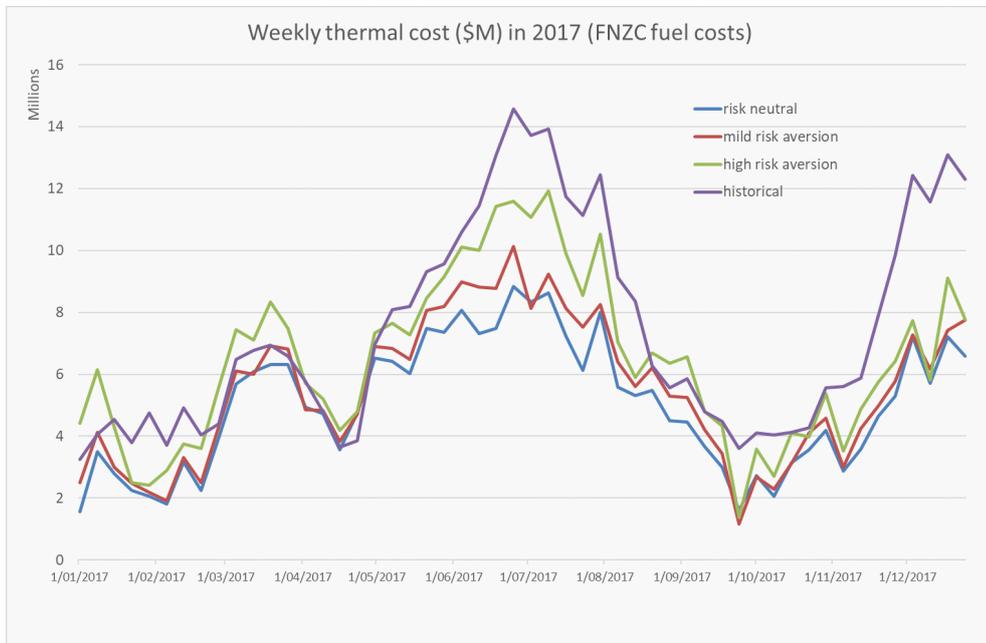


Figure 24: Weekly costs of thermal fuel and carbon emissions (\$m) in 2017 for FNZC cost model.

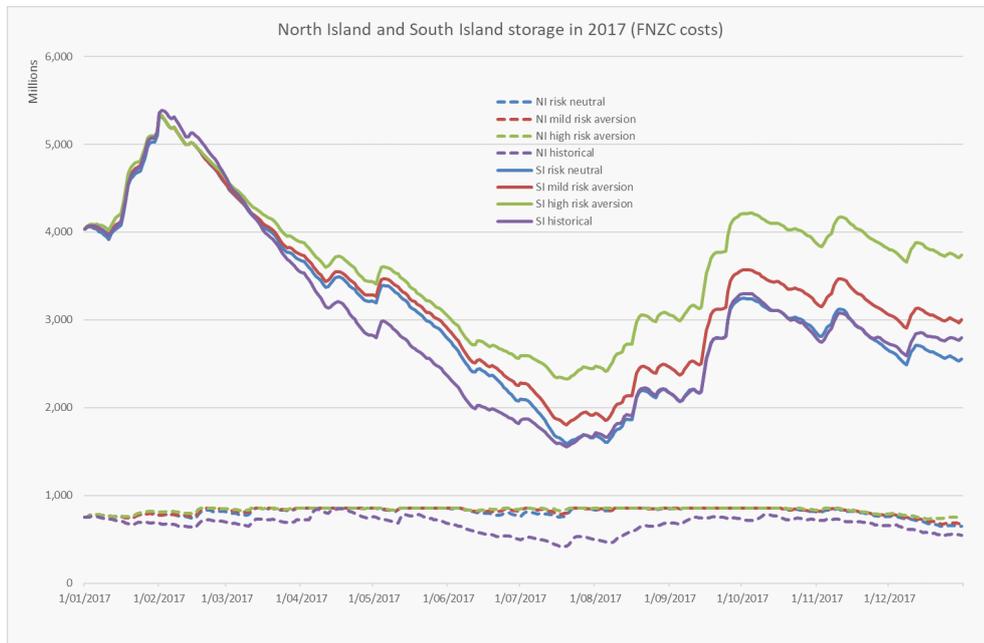


Figure 25: North Island and South Island reservoir storage (Mm^3) in 2017 for FNZC cost model.

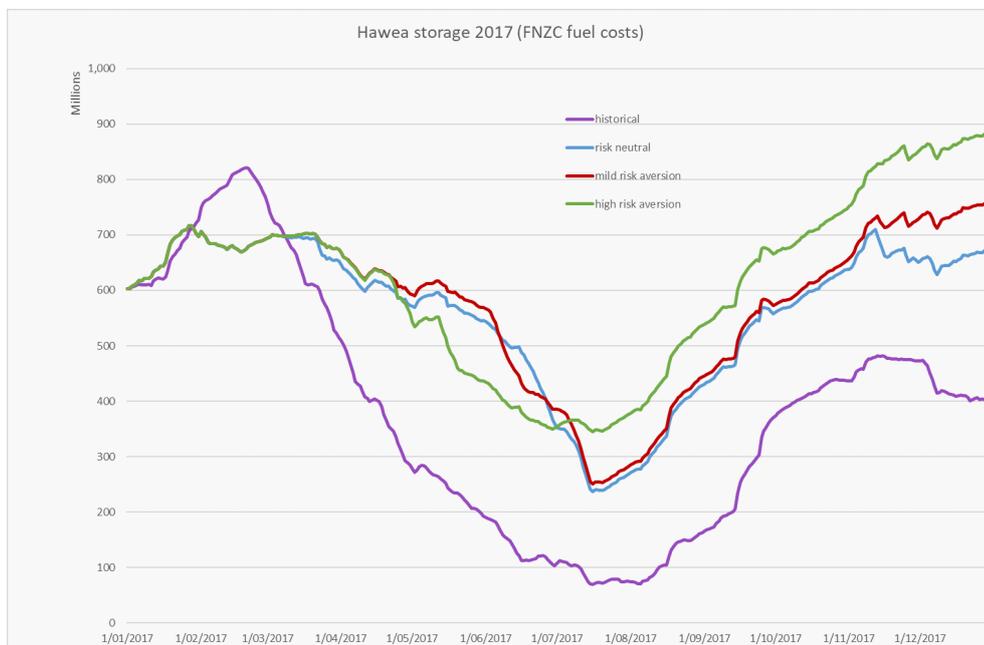


Figure 26: Lake Hawea reservoir storage (Mm^3) in 2017 for FNZC cost model.

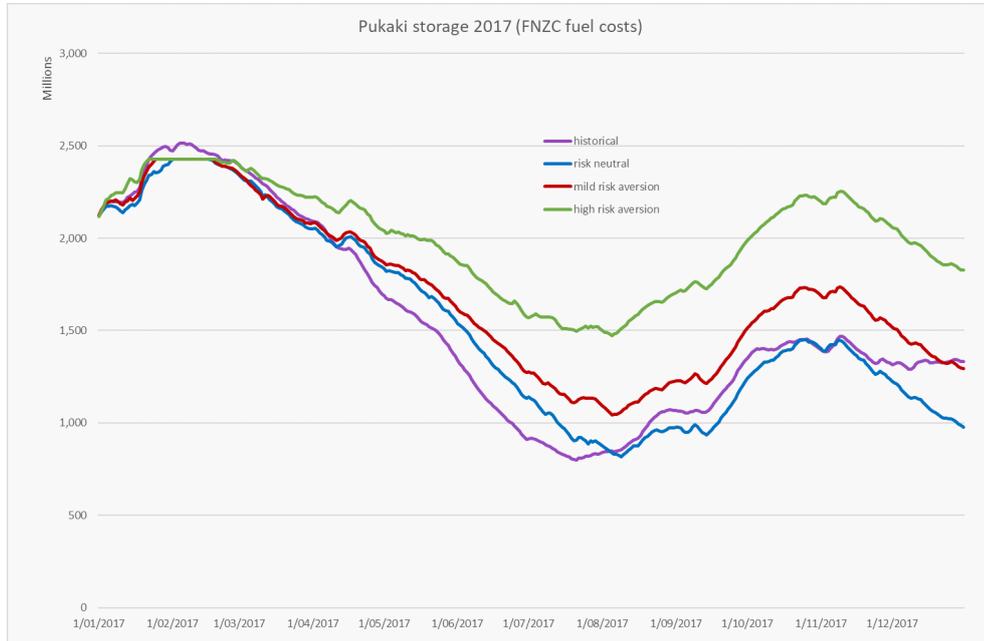


Figure 27: Lake Pukaki reservoir storage (Mm³) in 2017 for FNZC cost model.

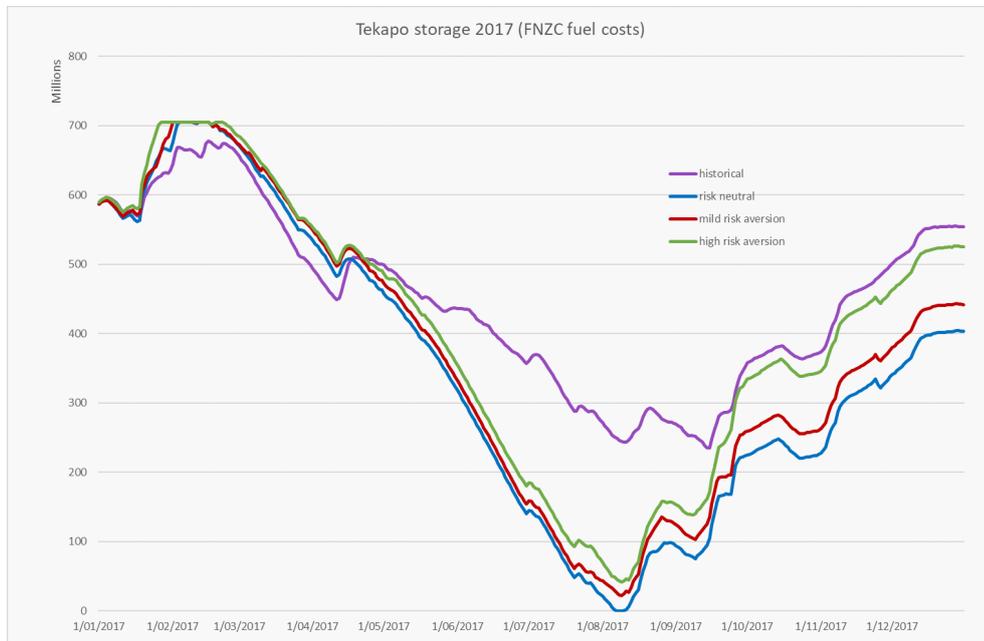


Figure 28: Lake Tekapo reservoir storage (Mm³) in 2017 for FNZC cost model.

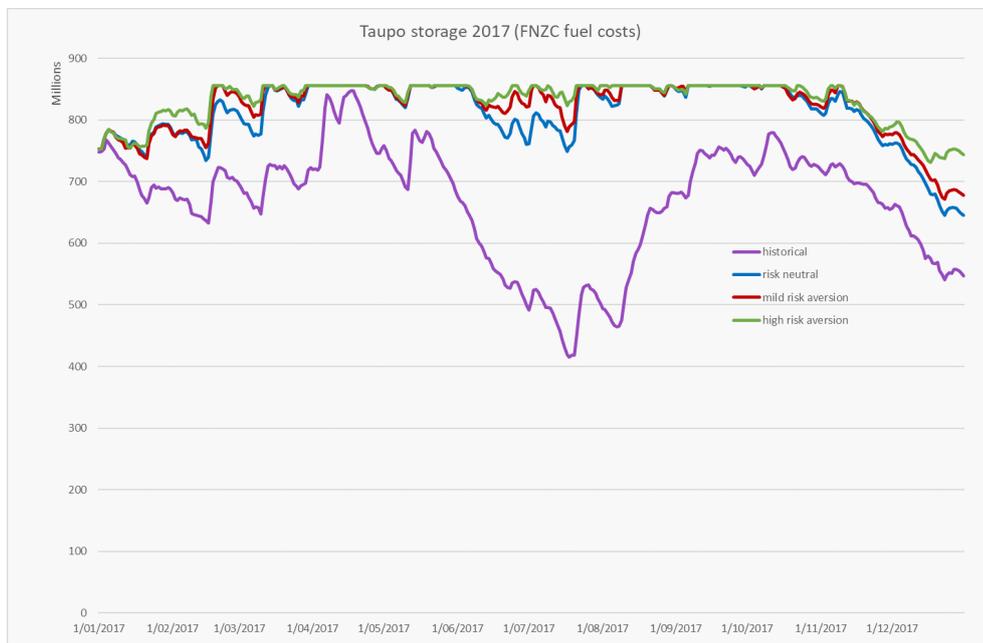


Figure 29: Lake Taupo reservoir storage (Mm³) in 2017 for FNZC cost model.

5.2.4 Nodal prices

Prices that are generated by the counterfactual policies and the rents accruing are shown below. The generation-weighted average prices (GWAPs) are shown in Figure 30 for the North Island and Figure 31 for the South Island. These show similar effects as in the previous simulations with MBIE fuel costs. Observe that all GWAPs are increased owing to the assumptions of higher fuel cost.

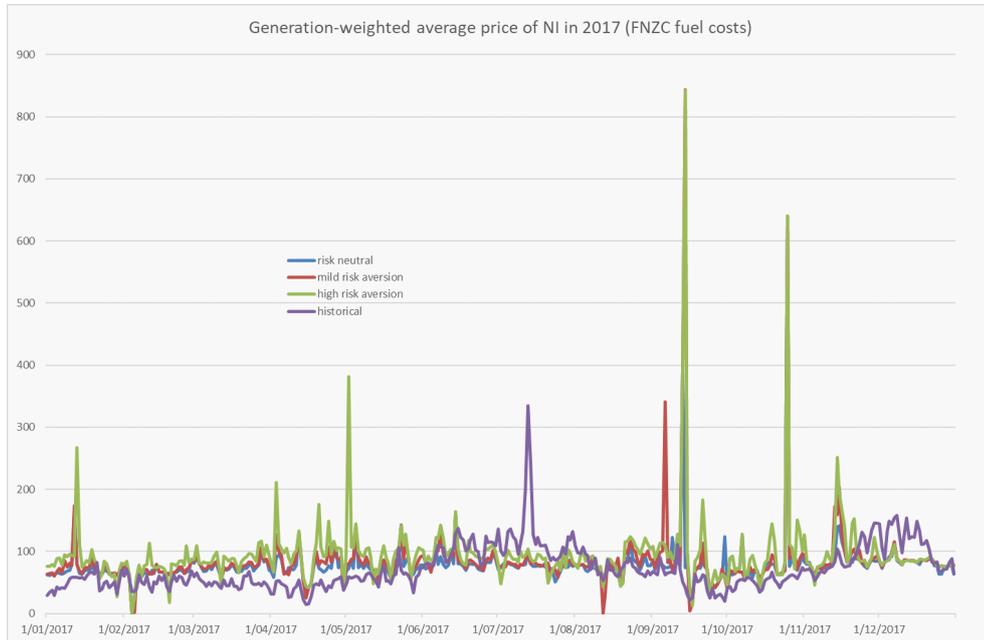


Figure 30: Generation-weighted daily average prices (\$/MWh) over North Island grid injection points (for large generators) in 2017 for FNZC cost model.

Time-weighted daily average prices are plotted in Figures 32, 33, and 34. These show similar differences in Benmore and Haywards prices from HVDC congestion in the counterfactual models that does not appear in historical dispatch. The accrued HVDC rents are shown in Figure 35. Counterfactual HVDC South-North transmission rents accrued over 2017 are significantly greater than historical values as shown in Table 6.

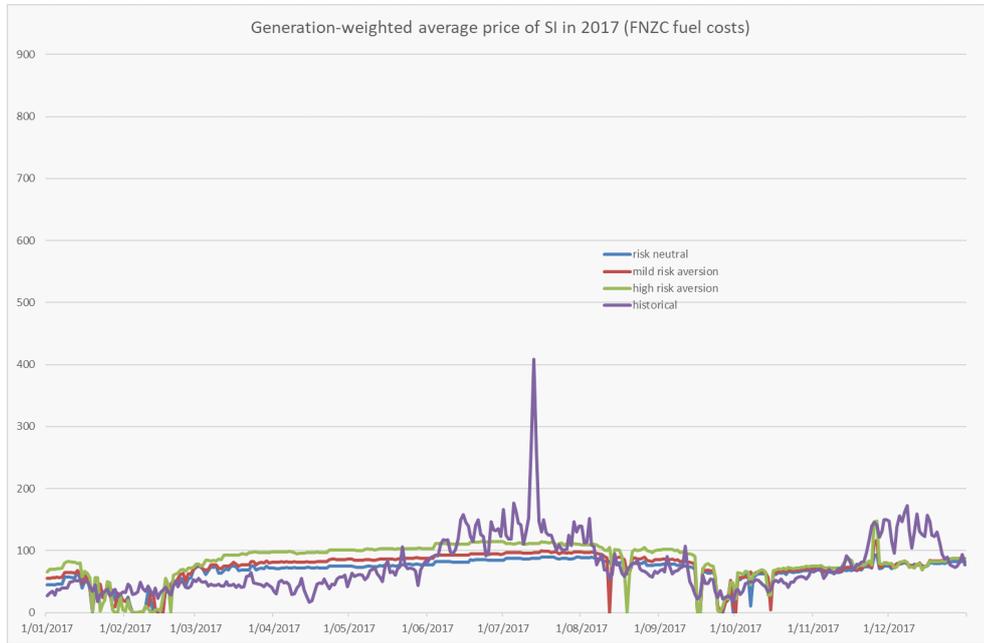


Figure 31: Generation-weighted daily average prices (\$/MWh) over South Island grid injection points (for large generators) in 2017 for FNZC cost model.

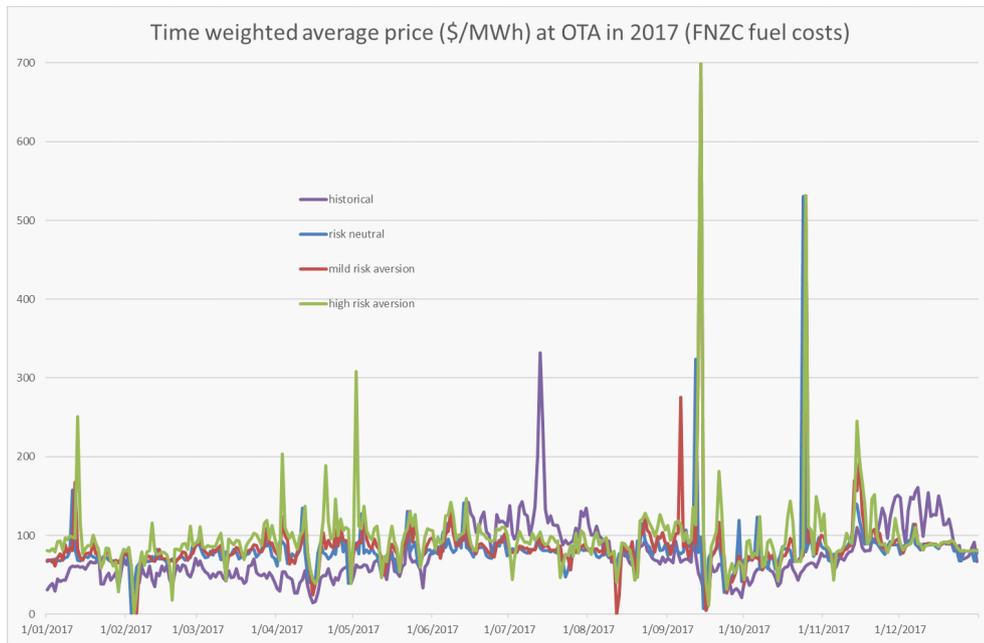


Figure 32: Time-weighted daily average prices (\$/MWh) at OTA in 2017 for FNZC cost model.

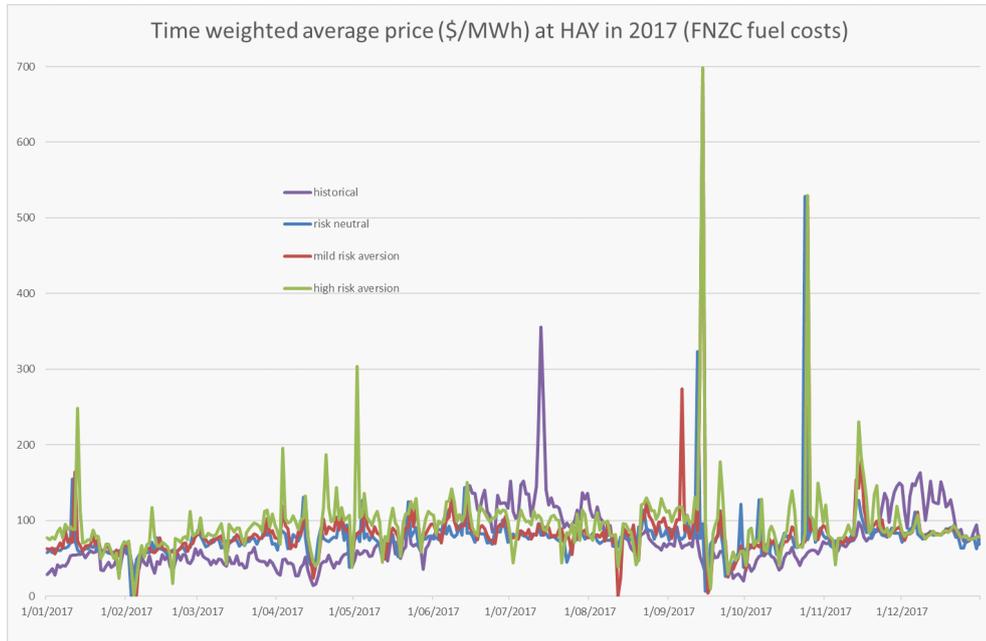


Figure 33: Time-weighted daily average prices (\$/MWh) at HAY in 2017 for FNZC cost model.

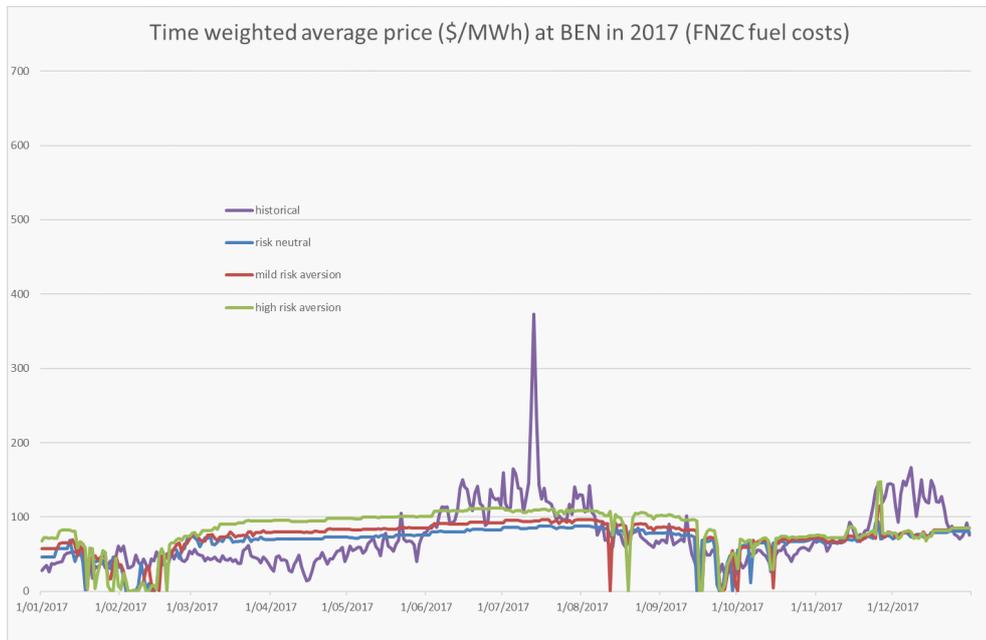


Figure 34: Time-weighted daily average prices (\$/MWh) at BEN in 2017 for FNZC cost model.

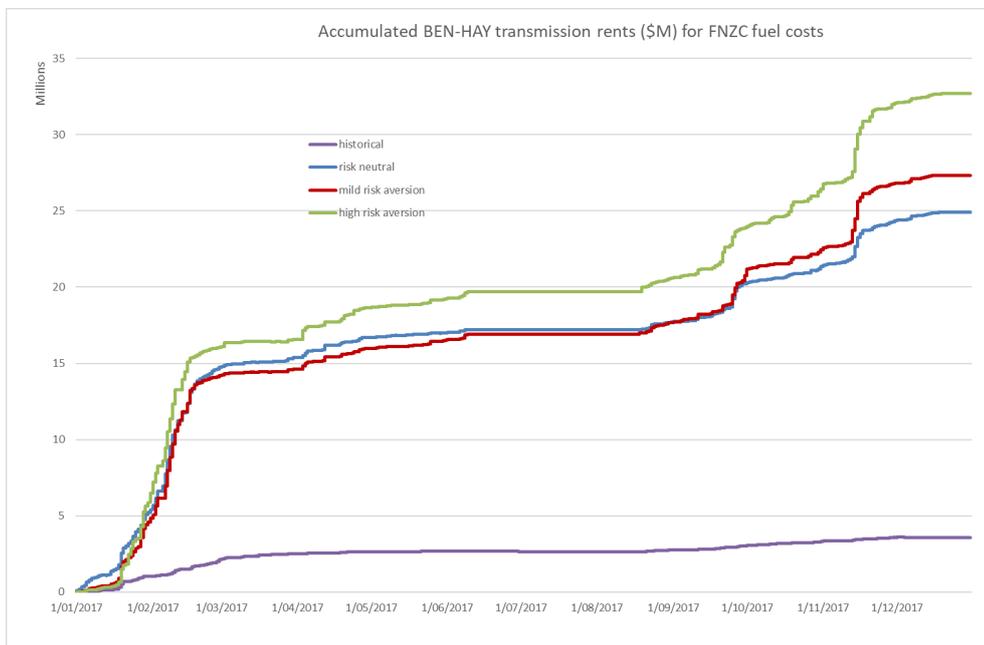


Figure 35: Cumulative HVDC rentals over 2017 for FNZC cost model.

6 Conclusions

In this paper we have described some experiments with stochastic optimization models of the New Zealand wholesale electricity market that provide counterfactual outcomes for competitive markets. As discussed in the introduction these solutions provide some insight into deviations of historical outcomes from perfectly competitive counterfactual solutions (with varying levels of risk aversion). We conclude the paper with some observations.

1. The outcomes from MBIE and FNZC cost assumptions are different. In the former case the counterfactual solutions have similar costs to the historical solution but prices, revenues, and rents are lower. In contrast the FNZC counterfactual models produce high prices and rents but short-run costs are much lower than historical values. Observe that we compute both historical costs and counterfactual costs using the assumed fuel-cost estimates. The prices from the FNZC risk-neutral counterfactual model give the closest match to historical values, while both risk-averse counterfactual solutions yield higher prices.

These outcomes point to the need for regulators to produce audited estimates of competitive short-run gas costs that can be used in benchmark studies. For benchmark studies, an opportunity cost of gas (which will drive thermal offer prices) is not very helpful. Opportunity costs will be driven by alternative use, one of which is competing electricity generation. The opportunity cost then amounts to estimating future electricity prices and spark spreads for thermal generation, which is likely to result in counterfactual solutions that reproduce those price estimates. In other words, we need to assume that gas costs come from a competitive gas market with other uses for the fuel, in order to assess levels of competition in the wholesale electricity market.

2. The inflow processes used in DOASA are assumed to be stagewise independent. This means that a sequence of dry weeks will occur in the model with lower probability than it would in reality. When reservoir levels are low, and low inflows persist, this assumption will tend to produce optimistic estimates of future costs. The marginal water values at low reservoir levels are therefore likely to be lower in our model than in a model with serial dependence. We have attempted to account for this dependence using an inflow adjustment (DIA). With these adjustments, and a suitable choice of risk aversion, the results of our counterfactual experiments can be made to match historical reservoir levels reasonably closely at least at a national level.

3. One result from this study that is common to both MBIE and FNZC cost estimates is the difference in HVDC rents accrued by the counterfactual solutions and historical dispatch, as shown in Figures 19 and 35. Historical South Island prices are closer to North Island values than they are in all counterfactual solutions. The differences depend on levels of risk aversion but are around \$10/MWh. There are several implications of this difference.
 - (a) Water values are computed by generators based on expectations of future prices. If these are \$10/MWh higher on average than competitive values then one would expect these premia to be reflected in water values. Many water valuation models for hydro generators are calibrated to historical prices. These models then determine offers of the hydro generators in the wholesale market, who can reasonably claim to be offering energy at short-run marginal cost (defined by the expected marginal value of the stored water). However this mis-estimates the perfectly competitive opportunity cost, since, according to our counterfactual models, these water values are higher than they would be under (risked) perfect competition.
 - (b) One might conjecture that lower South Island water values observed in the counterfactual solutions are a consequence of modelling HVDC transmission constraints and losses in DOASA. In other words a simpler one-node model of the electricity system would equilibrate water values in Pukaki and Taupo, and so one would expect more uniform prices across the country. This would imply that being able to model transmission in DOASA is a key point of difference that makes it an improvement over single-node water value models. We tested this in 2017, by running DOASA with no transmission constraints or losses. The results for 2017 still showed some price separation between South and North Islands in HydrovSPD, so more analysis is needed to confirm this conjecture.
 - (c) The recent Transmission Pricing Methodology Review Decision [3] has stated that the HVDC charge on South Island generators “inefficiently discourages investment in South Island generation”. The review states “Dampening investment in generation pushes electricity prices higher than they need to be. The Authority considers the new guidelines will contribute to unlocking renewable generation in the South Island and lower generation costs for the long-term benefit of New Zealand consumers.” This statement is true when South Island wholesale energy prices are competitive, but our experiments indicate they are around \$10/MWh higher in 2017. One might argue that these South

Island premia arise from generators recovering transmission charges that they feel are unfairly applied. If this is true then, under the proposed beneficiary-pays pricing regime, one might hope to see lower wholesale prices in the South Island.

4. The counterfactual outcomes we report are confined to the wholesale electricity market, and our models do not include forward contracts or retail sales. It can be argued that most electricity produced is sold at retail or contract prices, and so it is these prices that should be used in counterfactual comparisons. Unfortunately it is difficult to do this directly as historical contract quantities and prices are not in the public domain. Under the assumption of perfect competition, our counterfactual models assume that generators make contract decisions to share risk. The social planning solution that we compute emerges from the optimal trading of risk between market participants using these sorts of instruments. We do not attempt to compute these trades as part of our counterfactual solution; they will alter the distribution of Ricardian rents that accrue to each agent in our risk-averse social optimum¹¹. Some of the rents accruing to generators in periods of high prices will be exchanged through hedge contracts with large consumers for compensating payments when prices are low. On the other hand, if generators make some sales in the retail market (i.e. are gentailers) then this proportion of their revenue is hedged with their retail arm. So gentailers will retain this proportion of wholesale rents.

We also warn against drawing too many conclusions from differences in Ricardian rent observed in a single year (i.e. assuming that these hold in other years). Differences can vary over historical years. In some periods in 2009, historical rents were negative for some generating plant. It would be wrong to deduce that generators earned less in these periods, since their contract prices can be above the average value of wholesale prices. On the other hand, an average of prices computed over several years should reduce this effect, while still being biased below contract prices, since in most electricity markets (including New Zealand) contracts are traded at a premium to expected spot prices (see [5]), so counterfactual wholesale prices averaged over several years would be expected to be underestimates of contract prices.

5. As observed above, it is possible (at least in theory) for a Walrasian equilibrium to give a stochastic process of prices with respect to which every

¹¹Observe that the redistribution of rents through contracts is zero sum, so total rents will be the same.

agent optimizes its own expected benefit with the outcome of maximizing total expected welfare. Such an equilibrium might give a sample path of prices as observed for example in Figure 17. As shown by [4], the stochastic process of prices that yields an equilibrium might be very complicated with none of the stagewise independence properties that make computing optimal policies relatively easy using dynamic programming. From the storage trajectories shown in Figure 9 and matching price trajectories shown in Figure 14 and Figure 15, one can see that there are many different price sequences that will support prudent hydro reservoir releases. Consumers value electricity highly, and price has traditionally been a poor instrument to control short-term demand. Inelastic demand means offer prices early in the year in response to a dry-winter forecast may not lead to much change in consumption or even any change in dispatch. Observed price increases in these circumstances align with broad economic incentives, but this would be true for any price increase, at least up to the point where entry of new generation is prompted.

6. Do the results in this paper support the assertion that the New Zealand wholesale market is close to perfectly competitive? The short-run costs of the counterfactual solutions are reasonably close to historical values in 2017, at least in the case of MBIE fuel costs, so the wholesale market seems to lose little in productive efficiency. The counterfactual solutions, however, have a different generation mix, with more hydro generation and less thermal generation than historical levels. Counterfactual South Island hydro generation is higher in the early and later parts of the year leading to larger HVDC transfers to the North. Increased hydro generation leaves reservoir levels lower at the end of 2017; a higher future cost is exchanged for a lower thermal fuel cost. The short-run costs of the counterfactual solutions are significantly lower than historical values when FNZC cost estimates are used, as fuel costs now account for a higher fraction of overall cost. Accurate estimates of gas costs are clearly a key requirement to estimating productive efficiency.

Historical prices and revenues in the market are higher than counterfactual values, at least in the case of MBIE fuel costs. Some authors ([6], [31]) have estimated price markups using benchmark counterfactual models based on conjectured models of generator behaviour and statistically estimated water values. In contrast, our behavioural model assumes profit-maximizing price-taking agents, and is based on a theoretical alignment between perfect competition under risk and social optimization, and water values are computed using “bottom up” stochastic optimization models. A stochastic optimization avoids some of the foresight bias observed in the results of [35].

We have not eliminated bias entirely as HydrovSPD admits full clairvoyance of intra-day inflows. In addition, the benchmark model will still have some residual bias from relaxing other information constraints that will be present in a real setting.

References

- [1] T. Alvey, D. Goodwin, X. Ma, D. Streiffert, and D. Sun. A security-constrained bid-clearing system for the New Zealand wholesale electricity market. *IEEE Transactions on Power Systems*, 13(2):340 – 346, 1998.
- [2] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9:203–228, 1999.
- [3] Electricity Authority. <https://www.ea.govt.nz/dmsdocument/26851-tpm-decision-paper>, 2020.
- [4] K. Barty, P. Carpentier, and P. Girardeau. Decomposition of large-scale stochastic optimal control problems. *RAIRO-Operations Research*, 44(03):167–183, 2010.
- [5] F. Bevin-McCrimmon, I. Diaz-Rainey, M. McCarten, and G. Sise. Liquidity and risk premia in electricity futures. *Energy Economics*, 75:503–517, 2018.
- [6] O. Browne, S. Poletti, and D. Young. Simulating market power in the new zealand electricity market. *New Zealand Economic Papers*, 46(1):35–50, 2012.
- [7] T. Denne. Environmental costs of electricity generation. Technical report, COVEC, 2015.
- [8] O. Dowson, D.P. Morton, and B.K. Pagnoncelli. Multistage stochastic programs with the entropic risk measure. <http://www.optimization-online.org>, 2020.
- [9] M.C. Ferris and A.B. Philpott. Dynamic risk equilibrium. Technical report, Electric Power Optimization Centre, 2018.
- [10] P.C. Fishburn. *Nonlinear preference and utility theory*. Johns Hopkins University Press Baltimore, 1988.
- [11] N. Gluyas. FNZC estimates of fuel prices. Private communication, 2018.

- [12] Z. Guan. EMBER online companion documents. <http://www.epoc.org.nz/ember2.html>.
- [13] W.W. Hogan. Electricity scarcity pricing through operating reserves. *Economics of Energy & Environmental Policy*, 2(2):65–86, 2013.
- [14] C. Kok, A.B. Philpott, and G. Zakeri. Value of transmission capacity in electricity markets with risk averse agents. *www.epoc.org.nz*, 2018.
- [15] P. Lino, L.A.N. Barroso, M.V.F. Pereira, R. Kelman, and M.H.C. Fampa. Bid-based dispatch of hydrothermal systems in competitive markets. *Annals of Operations Research*, 120(1):81–97, 2003.
- [16] EMI-Electricity Market Information System. New Zealand Electricity Authority. <http://reports.ea.govt.nz/emi.htm>.
- [17] Ministry of Business Innovation and Employment. Energy statistics: Prices. <https://www.mbie.govt.nz/info-services/sectors-industries/energy/energy-data-modelling/statistics/prices>.
- [18] New Zealand Electricity Authority. Vectorized Schedule Price and Dispatch. <https://www.emi.ea.govt.nz/Wholesale/Tools/vSPD>.
- [19] D.M.G. Newbery and J.E. Stiglitz. Pareto inferior trade. *The Review of Economic Studies*, 51(1):1–12, 1984.
- [20] Parsons Brinckerhoff. 2011 NZ Generation Data Update. <https://www.mbie.govt.nz/info-services/sectors-industries/energy/energy-data-modelling/technical-papers/pdf-library/2011%20NZ%20Generation%20Data%20Update%20v006a.pdf>, 2012.
- [21] M. V. F. Pereira and L. M. V. G. Pinto. Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, 52:359–375, 1991.
- [22] A.B. Philpott and V.L. De Matos. Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion. *European Journal of Operational Research*, 218(2):470–483, 2012.
- [23] A.B. Philpott, V.L. de Matos, and E. Finardi. On solving multistage stochastic programs with coherent risk measures. *Operations Research*, 61(4):957–970, 2013.
- [24] A.B. Philpott, V.L. de Matos, and L. Kapelevich. Distributionally robust SDDP. *Computational Management Science*, 15(3-4):431–454, 2018.

- [25] A.B. Philpott and Z. Guan. On the convergence of stochastic dual dynamic programming and other methods. *Operations Research Letters*, 36:450–455, 2008.
- [26] A.B. Philpott and Z. Guan. Models for estimating the performance of electricity markets with hydro-electric reservoir storage. Technical report, Electric Power Optimization Centre, 2013.
- [27] A.B. Philpott and Z. Guan. Benchmarking wholesale hydroelectricity markets with risk-averse agents. Technical report, Electric Power Optimization Centre, 2018.
- [28] A.B. Philpott and Z. Guan. On the performance of the New Zealand wholesale electricity market 2008-2017 (forthcoming). Technical report, Electric Power Optimization Centre, 2021.
- [29] A.B. Philpott, Z. Guan, J. Khazaei, and G. Zakeri. Production inefficiency of electricity markets with hydro generation. *Utilities Policy*, 18(4):174 – 185, 2010.
- [30] POCP. Database of planned outages. <http://pocp.redspider.co.nz>.
- [31] S. Poletti. Market power in the NZ wholesale electricity market 2010-2016. In *Local Energy, Global Markets, 42nd IAEE International Conference, May 29-June 1, 2019*. International Association for Energy Economics, 2019.
- [32] N.R. Porter. Intra-day uncertainty and efficiency in electricity markets. Masters thesis, Electric Power Optimization Centre, University of Auckland, downloadable from www.epoc.org.nz, 2015.
- [33] S. Stoft. *Power system economics*. IEEE Press and John Wiley and Sons, 2002.
- [34] Wikipedia. National Electricity Market. https://en.wikipedia.org/wiki/National_Electricity_Market.
- [35] F.A. Wolak. An assessment of the performance of the New Zealand wholesale electricity market. Technical report, New Zealand Commerce Commission, 2009.

Price discovery can be inefficient

Andy Philpott
Electric Power Optimization Centre
University of Auckland

December 20, 2021

Abstract

New Zealand has an energy-only wholesale electricity market that operates in real time. Before real time, the system operator publishes provisional prices from a sequence of auctions that use forecast demand. Electricity generators respond to these provisional prices by updating their offers. We show how this process can lead to inefficient dispatch and electricity prices.

1 Introduction

New Zealand has an energy-only wholesale electricity market that operates in real time. Before real time, the system operator publishes provisional prices from a sequence of auctions that use forecast demand. Electricity generators respond to these provisional prices by updating their offers. We show how this process can lead to inefficient dispatch and electricity prices.

We begin by presenting a stylized two-period model with two competing generators to show how an inflated forecast of demand might lead to an inefficient dispatch even if this forecast is corrected. We then discuss a setting in which a sequence of forecasts are made that converge to the final demand. This gives a sequence of provisional dispatches and prices, called a *pre-dispatch*, that models what happens in the New Zealand wholesale market. We prove a theorem that shows under mild assumptions that this process can lead to inefficient dispatch and prices. The final section of the paper looks at an example of a particular day, May 5, 2020.

2 A two-period model

Consider a simple example in which a competitive generator (1) with strictly increasing marginal cost function $c_1(u_1)$ competes with a generator (2) with capacity b and marginal cost $c_2 < c_1(0)$. If demand is $d > b$ then the dispatch

for generator 1 and 2 respectively is $u_1 = d - b$, $u_2 = b$, the dispatch cost is

$$C = c_2 b + \int_0^{d-b} c_1(u) du,$$

and the clearing price is $\pi(d) = c_1(d - b)$. If generator 2 increases their offer price to any price $p < \pi(d)$ then the optimal dispatch and clearing price are the same. This could be interpreted as price-taking behaviour, i.e., generator 2 is not seeking to influence the price by their offer.

Now suppose that the actual load d is replaced by a forecast \hat{d} , and a provisional auction run with this forecast. If $\hat{d} > d$ then the provisional clearing price is $\pi(\hat{d}) = c_1(\hat{d} - b) > \pi(d)$. Now if generator 2 increases their offer price to $p = \pi(\hat{d})$ then the optimal dispatch and clearing price are the same as long as demand is the same as forecast. However if we dispatch plant to meet the actual load d then we obtain

$$\begin{aligned} \hat{u}_1 &= c_1^{-1}(\pi(\hat{d})) = \hat{d} - b > d - b, \\ \hat{u}_2 &= d - \hat{u}_1 < b, \end{aligned}$$

and the clearing price is $\pi(\hat{d})$. This is an inefficient dispatch since $c_2 < c_1(0) \leq c_1(u)$, $u \geq 0$ implies

$$c_2(\hat{d} - d) = \int_{d-b}^{\hat{d}-b} c_2 du < \int_{d-b}^{\hat{d}-b} c_1(u) du,$$

so the cost of (\hat{u}_1, \hat{u}_2) is

$$c_2(b + d - \hat{d}) + \int_0^{\hat{d}-b} c_1(u) du > c_2 b + \int_0^{d-b} c_1(u) du = C.$$

Observe that the price markup $\pi(\hat{d}) - \pi(d)$ persists, even if the forecast $\hat{d} > d$ is corrected in a later period to a correct forecast. Moreover the dispatch is less efficient than the perfectly competitive dispatch. This simple example shows that inaccurate demand forecasts can lead to price markups when the uncertainty in demand is resolved.

3 Predispatch auctions

In this section we are interested in more general circumstances in which inaccurate predispatch demand forecasts lead to inefficient outcomes. Consider a setting where the same two generators as in the previous section compete, but a sequence of auctions is run for a given trading period for increasingly accurate demand forecasts \hat{d} . Suppose the ISO solves a sequence of dispatch problems $P(k)$, $k = 1, 2, \dots$, with forecast demand d_k yielding clearing price $\pi(d_k)$. For convenience, we make the following assumption about forecasts.

Assumption 1 Demand forecasts satisfy $d_k > b$ and are accurate enough so that $|d_k - d_l| < b$ for every $k, l > 0$.

We require the following result.

Lemma 2 Suppose generator 2 offers b at price p and demand is $d > b$. Then the clearing price $\pi(d, p)$ is continuous and nondecreasing in d and p .

Proof. Suppose d is fixed. Then

$$\pi(d, p) = \begin{cases} c_1(d - b), & p < c_1(d - b), \\ p, & c_1(d - b) \leq p < c_1(d), \\ c_1(d), & p \geq c_1(d), \end{cases}$$

which is continuous and nondecreasing in p .

Suppose p is fixed. Then

$$\pi(d, p) = \begin{cases} c_1(d), & d \leq c_1^{-1}(p), \\ p, & c_1^{-1}(p) < d \leq c_1^{-1}(p) + b, \\ c_1(d - b), & d > c_1^{-1}(p) + b, \end{cases}$$

which is continuous and nondecreasing in d . ■

We now prove the main result of this paper.

Proposition 3 Suppose the ISO solves a sequence of dispatch problems $P(k)$, $k = 1, 2, \dots$, with forecast demand $d_k \rightarrow d$, each yielding clearing price $\pi(d_k)$. Assume the forecasts satisfy Assumption 1. If for $k = 0, 1, 2, \dots$, generator 2 offers b at price $p = \pi(d_k)$ to problem $P(k + 1)$, where $\pi(d_0) = c_2$, then $\pi(d_k) \rightarrow \pi(d)$ if and only if $d_k \leq d$ for every k .

Proof. We show that $\pi(d_k) = \max_{l \leq k} \{c_1(d_l - b)\}$, so $\pi(d_k)$ is a nondecreasing sequence. This shows that $\pi(d_k)$ does not converge to $\pi(d)$ if any demand forecasts $d_k > d$. Conversely $\max_{l \leq k} \{c_1(d_l - b)\}$ is easily shown to converge to $c_1(d - b)$ if $d_k \uparrow d$.

Suppose as an induction hypothesis that for some k , $\pi(d_k) = \max_{l \leq k} \{c_1(d_l - b)\}$. This is true for $k = 1$, since

$$\pi(d_1) = c_1(d_1 - b) > c_2 = \pi(d_0).$$

By Lemma 1, we have that the clearing price resulting from problem $P(k + 1)$ is

$$\begin{aligned} \pi(d_{k+1}) &= \begin{cases} c_1(d_{k+1} - b), & \pi(d_k) < c_1(d_{k+1} - b), \\ \pi(d_k), & c_1(d_{k+1} - b) \leq \pi(d_k) < c_1(d_{k+1}), \\ c_1(d_{k+1}), & \pi(d_k) \geq c_1(d_{k+1}), \end{cases} \\ &= \begin{cases} c_1(d_{k+1} - b), & \max_{l \leq k} \{c_1(d_l - b)\} < c_1(d_{k+1} - b), \\ \max_{l \leq k} \{c_1(d_l - b)\}, & c_1(d_{k+1} - b) \leq \max_{l \leq k} \{c_1(d_l - b)\} < c_1(d_{k+1}), \\ c_1(d_{k+1}), & \max_{l \leq k} \{c_1(d_l - b)\} \geq c_1(d_{k+1}). \end{cases} \end{aligned}$$

Now $\max_{l \leq k} \{c_1(d_l - b)\} \geq c_1(d_{k+1})$ implies that $c_1(d_l - b) \geq c_1(d_{k+1})$ for some $l \leq k$, so $d_l - b \geq d_{k+1}$ contradicting Assumption 2. Thus the third range for $\pi(d_{k+1})$ cannot occur and

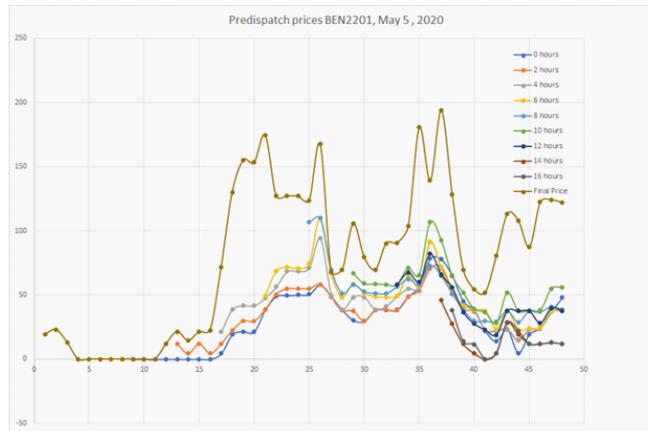
$$\begin{aligned} \pi(d_{k+1}) &= \begin{cases} c_1(d_{k+1} - b), & \max_{l \leq k} \{c_1(d_l - b)\} < c_1(d_{k+1} - b), \\ \max_{l \leq k} \{c_1(d_l - b)\}, & c_1(d_{k+1} - b) \leq \max_{l \leq k} \{c_1(d_l - b)\}, \end{cases} \\ &= \max_{l \leq k+1} \{c_1(d_l - b)\}, \end{aligned}$$

which yields the result. ■

This result shows that a sequence of solutions to dispatch auctions in which agents seek to discover the price might lead to an inefficient outcome. The setting of the theorem with two generators having specific cost structures does not represent all the nuances of the New Zealand market, but simply highlights a deleterious feature of a repeated auction mechanism. Precluding generators from updating offers in light of pre-dispatch information would converge on a perfectly competitive solution if their original offers were at marginal cost.

4 Example

How often does the phenomenon described above occur in actual pre-dispatch schedules? It would be interesting to examine forecast demands for pre-dispatch and corresponding offers and clearing prices, and see whether the behaviour described above occurs. Pre-dispatch prices are published on the EMI site of the New Zealand Electricity Authority. We extracted these for May 5, 2020, and plot them below.



Predispatch prices and final prices for Benmore on May 5, 2020. Each coloured line shows a predispatch price schedule (PRSL) computed for the next 72 trading periods, but displayed up to the end of the day.

Observe that predispatch prices are generally increasing, and converge to prices that are below the final price. It seems unlikely that demand forecasts were also increasing towards their true final values as the forecast interval gets smaller, as this would imply a systematic bias (although a cold front getting colder faster than predicted might lead to this outcome). A systematic examination of historical demand forecasts should be able to resolve this.

5 Conclusion

It is difficult to say if the imperfection in predispatch described above warrants a redesign of the New Zealand auction mechanism. It has always been taken as an article of faith in the New Zealand wholesale market that more frequent auctions lead to better outcomes, as they enable price discovery. This is certainly the case in auction settings where purchasers' valuations are private information and the seller wishes to extract (and capture) as much of this value as she can. In an electricity market that is regulated to maximize the long-term benefit of consumers, it would seem to be strange to move towards more frequent predispatch auctions that might decrease this benefit.