Applications of SDDP in electricity markets with hydroelectricity

Andy Philpott
Electric Power Optimization Centre
University of Auckland.
www.epoc.org.nz

SESO Workshop, Paris, June 1, 2017.

Motivation

- SDDP developed in late 1980s by Mario Pereira and colleagues.
- There are now many implementations of SDDP.
- We have been using it in EPOC for 20 years.
- Talk about how we use it.

Summary

Summary

Hydrothermal scheduling

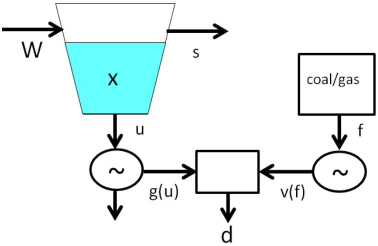

Find a water release policy over next T weeks that minimizes the expected thermal fuel cost of meeting demand over this period. Unsatisfied demand is met by shedding load at a shortage cost. A policy is not a fixed plan - it is a decision rule that says at the beginning of each week "given previous observations and actions, do the following". This means that water release actions are not predetermined but will depend on observations.

Mathematical description

$$
\begin{aligned}
\min \quad & \mathbb{E}\left[\sum_{t=1}^{T} c_t^\top v_t + V_{T+1}(x_{T+1})\right] \\
\text{s.t.} \quad & g(u_t) + v_t = d(t), \\
& x_{t+1} = x_t - A(u_t + s_t) + W_t, \\
& x_1 = \bar{x}, \\
& u_t \in \mathcal{U}, \quad v_t \in \mathcal{V}, \; s_t \geq 0, \qquad t = 1, \ldots, T.
\end{aligned}
$$

Expected marginal water value for a single reservoir

- Water release policies for a single-reservoir can be specified by an expected marginal water value measured in $/m^3 This is defined for each possible reservoir level, and is the opportunity cost of releasing one unit of water.
- The optimal policy generates from all thermal plant that are no more expensive (in $/MWh) than the expected marginal water value multiplied by a conversion factor (cubic metres per MWh).
- The water release policy solves a stage optimization problem: minimize the current cost of thermal generation plus the expected future cost of meeting demand. The slope of this future cost function at level x is (minus) the marginal water value at this level.

SDDP and electricity markets

- SDDP policies for national electricity systems seek a minimum-cost solution for a social planner.
- New Zealand has operated a deregulated electricity market since 1996.
- SDDP policies for the country are not used to value water.
- EPOC primarliy uses SDDP policies as a counterfactual to estimate how competitive the market was in past years.
- This is commonly called backtesting.

Optimal solutions and competitive market outcomes
(Philpott, Ferris and Wets, 2016)

- Water release policies from system optimization give a social planning solution.
- Social planning models correspond to perfectly competitive partial equilibrium when:
  - market participants act as risk-neutral price takers;
  - markets are complete;
  - market participants problems are convex;
  - market participants agree on probability distributions of random outcomes.
- When market participants are risk averse then social planning models correspond to perfectly competitive partial equilibrium when additionally:
  - risk measures are coherent and comparable;
  - markets for risk are complete.

Criticisms from backtesting

- The SDDP solution is only an approximation of the market, so counterfactual prices are wrong.
- How does one know that the SDDP software being used has no bugs?
- Can one be sure that SDDP gives an optimal operating policy for a whole year?
- Can SDDP policies identify the cause of any difference in outcomes (and offer a remedy)? Are counterfactual prices different because of lack of competition, or from attitudes to risk, or from something else?

Summary

Stochastic Dual Dynamic Programming (SDDP)

Computes a water release policy when there are a number of different reservoirs with (possibly) different locations, different inflow processes and different river systems. This can give different results from aggregated reservoirs. Optimal release is no longer determined by a single marginal water value. The optimal releases u given storages x and next period's inflow W, is computed by the recursive stage problem:

$$\begin{aligned} V_t(x) = \mathbb{E}[\min \quad & c_t^\top v + V_{t+1}(x - A(u+s) + W)] \\ \text{s.t.} \quad & g(u) + v = d(t), \\ & u \in \mathcal{U}, \quad v \in \mathcal{V}, \; s \geq 0. \end{aligned}$$

Cuts

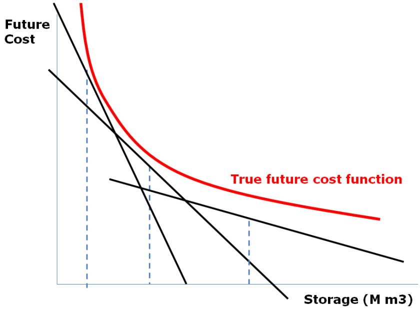

We need to define the function V_{t+1}. This is never known exactly but is approximated by the pointwise maximum of linear functions of x that are called cuts. SDDP algorithms like DOASA update these cuts using sampling procedures. For the New Zealand system about 5000-10000 cuts are needed in each week to define V_t accurately enough. This can take a lot of computation.

Bounds

- At each iteration of the algorithm SDDP gives a lower bound on the expected cost of an optimal policy. This increases as the algorithm proceeds.
- At each iteration of the algorithm SDDP gives a candidate policy (defined by the cuts). This can be simulated to estimate its expected cost and see how close it is to the current lower bound.
- Solutions are judged to be optimal (enough) when a confidence interval of their expected cost is close to the lower bound.

Inflow Models

- Inflow models preserve spatial dependence between catchments.
- Stagewise independent inflows sampled from historical record (the default for EMI-DOASA).
- Inflows sampled independently from historical values with variance increased to model correlation (DIA model in EMI-DOASA).
- PARMA models: simplest is AR1 model (PSR's SDDP and CEPEL's NEWAVE).
- PARMA models for logarithm of inflows (AR model avoiding negative inflows).
- Markov Chain model of climate state and inflows (Philpott & de Matos, 2012).

Summary

Approximating the New Zealand electricity system

- NZ has a nodal market with 250 nodes.
- Wholesale market is dispatched every 30 minutes using software called SPD. This yields generation levels for every generator and 250 prices every half hour.
- SPD inputs (offers of energy that are made by generators, network constraints and demand) are made public two weeks after the day of dispatch.
- NZ electricity regulator has developed a GAMS/CPLEX model that replicates SPD output. It is called vSPD.
- So one can run counterfactual tests on market outcomes from two weeks ago.

The SPD network

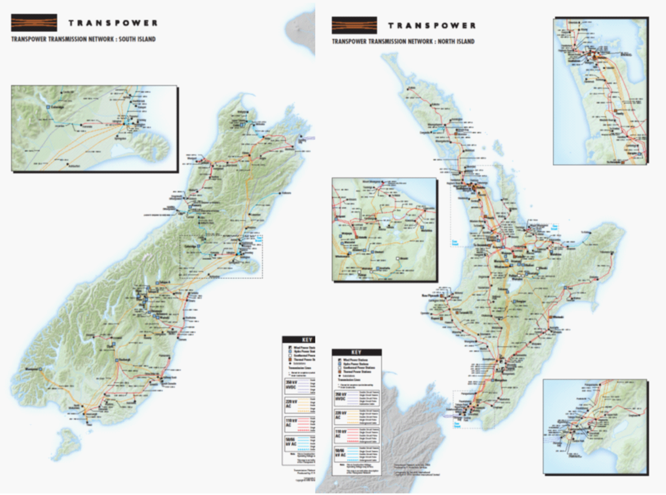

Network diagram of New Zealand transmission grid.

What about hydro?

- In a counterfactual, thermal agents offer energy at short run marginal cost (SRMC) which is cost from fuel, operations and maintenance and carbon.
- Run-of-river hydro agents, wind agents and geothermal agents offer energy at SRMC=0.
- Hydro agents with storage offer energy at SRMC=the marginal value of stored water. In the true market this is different from what one gets from SDDP. In a competitive counterfactual market SRMC can be set to the marginal value of water from SDDP

The network with hydro reservoirs shown

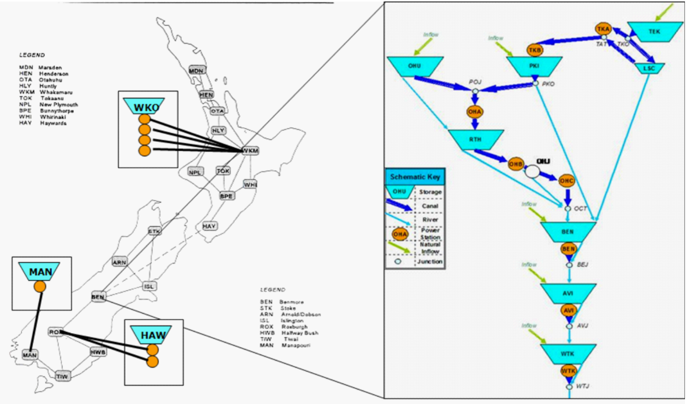

A three-node SDDP network with hydro

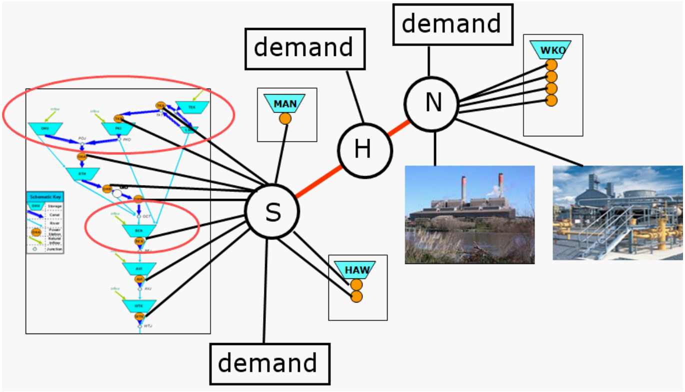

Computing a counterfactual for a calendar year

- Obtain historical reservoir levels for January 1, previous years inflow data and system data for year.
- For month $m = 1$ to 12
  - Solve SDDP model starting at beginning of month m to give an optimal policy
  - For each day d in month m
    - Solve a vSPD model for 48 trading periods linked by water conservation constraints and paying marginal value from SDDP cuts for all stored water used over the day;
    - Output quantities of interest for each period (prices, dispatch, reservoir levels);
    - Record reservoir levels for the start of the next day;
  - Next day
  - Record reservoir levels for the start of the next month;
- Next month

Summary

DOASA

DOASA stands for

Dynamic Outer Approximation Sampling Algorithm.

DOASA is an implementation of SDDP in C++ developed by Geoff Pritchard and Andy Philpott.

EMI-DOASA is the version of DOASA available on the NZ Electricity Authority EMI site for free use in New Zealand. There are CLP and Gurobi versions.

JADE
(Based on joint work with Lea Kapelevich and Oscar Dowson)

JADE stands for

Just Another DOASA Environment

JADE is an implementation of DOASA developed by Lea Kapelevich in the scripting language Julia. It uses the Julia packages StochDualDynamicProgram developed by Oscar Dowson, JuMP developed by Iain Dunning and Miles Lubin at MIT, and any LP solver (CLP, Gurobi, CPLEX, XPRESS).

JADE was developed completely independently from DOASA, but adopts the same modelling assumptions and reads the same input files. Because of its use of the JuMP modelling language, JADE can be easily configured to verify DOASA code changes and detect bugs.

JADE and DOASA lower bounds

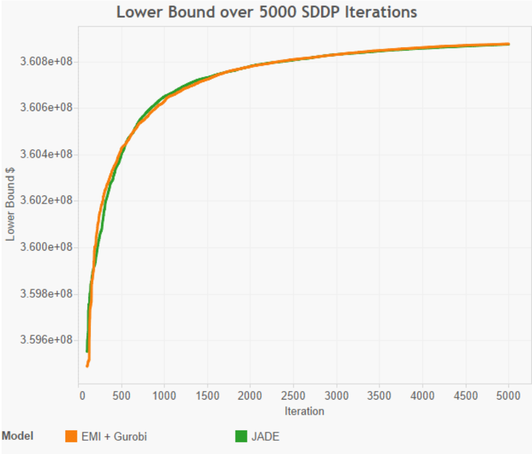

JADE and DOASA follow the same trajectories in computing lower bounds for the expected cost of the optimal policy.

Simulated value of policies

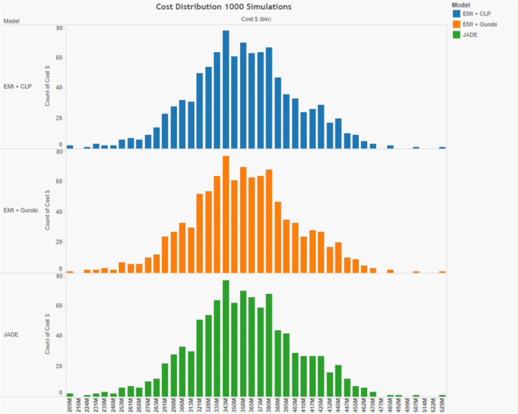

Policies give essentially the same distributions when simulated with Monte Carlo inflows.

Historical simulations

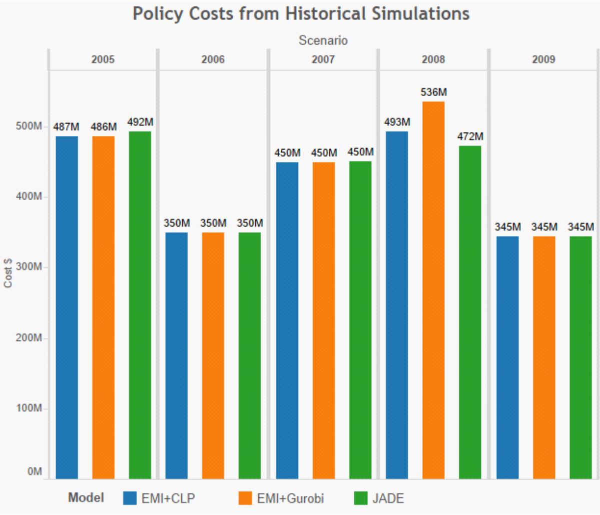

Solution times

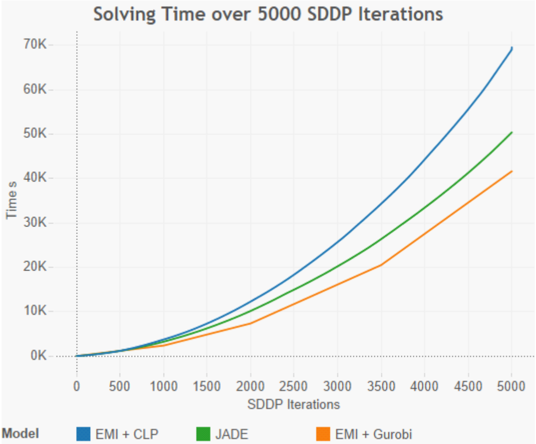

JADE using the Gurobi solver is faster than EMI-DOASA using CLP, but a bit slower than EMI-DOASA using Gurobi.

Summary

Test JADE with three and six load blocks

	3 blocks	6 blocks
Lower bound, $	320186241.7	320245280.6
Hours to solve	8.8	15.0

Solve time increases with load blocks.

Year		
2000	201919333.9	203885338.3
2001	373739554.9	373822531.2
2002	251224827.4	251220099.3
2003	298505773.3	298631336.6
2004	250365498.3	250382166.8
2005	439310054.1	439209067.4
2006	261982882.5	262007034.6
2007	353312988.9	353283433.9
2008	300137490	300124921.7
2009	264124711.9	264222247.9

Simulated expected cost of policies computed using 3 and 6 blocks in a model using 6 blocks.

Test JADE with different inflow models

Inflow model	Scenario cost (NZ)$M				
	2005	2006	2007	2008	2009
Independent	491.90	349.62	450.57	553.69	345.51
AR1	486.22	351.12	452.66	454.45	348.43
Log AR1	489.97	349.71	450.87	600.79	343.64
DIA 2	485.40	349.77	450.54	453.41	346.14
DIA 5	484.00	350.00	451.00	516.00	349.00
Risk Averse	484.82	349.68	451.47	447.09	345.20

JADE policy computed with 5000 cuts and simulated in historical years.

Test JADE with cut selection
(De Matos, Philpott and Finardi, 2015)

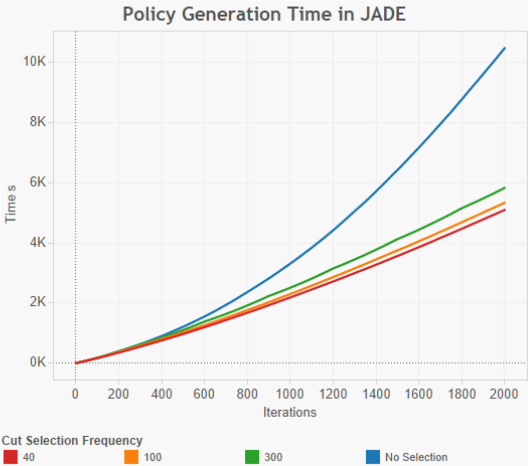

In each stage optimization, Level 1 dominance cut selection ignores dominated cuts at visited states (but restores them at the end).

Cut selection policy evaluation

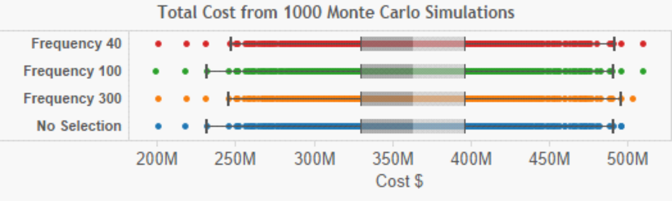

Effect of cut selection on simulated policy value.

Historical simulations with cut selection

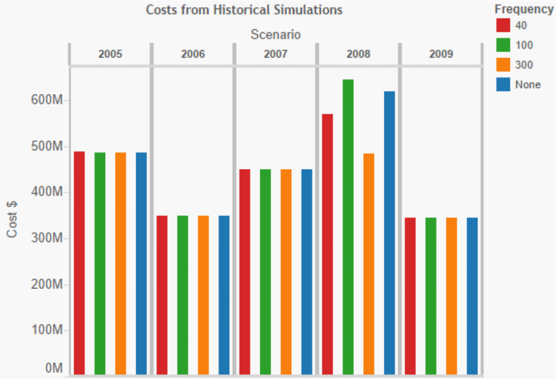

Cut selection affects policy in dry year.

SDDP policies are random

Year	Average	Standard Dev.
2000	286,225,815	82,305
2001	475,933,285	88,075
2002	342,189,502	100,076
2003	387,975,202	5,158,708
2004	256,211,965	348,440
2005	487,581,431	529,447
2006	349,597,459	64,129
2007	449,998,785	8,412
2008	501,068,672	42,335,412
2009	345,045,595	19,058

Table: Sample average and standard deviation ($) from simulating 10 SDDP optimal policies.

Summary

Plans and policies

- Multistage stochastic programming formalizes the planning process, to give a mathematical description of a plan with contingencies, assuming a given scenario tree for uncertain parameters throughout.
- To enact the plan, we take the first-stage action then recompute a new first-stage action after time has passed and new information (i.e. new data) becomes available. This typically involves a new scenario tree that can be quite different.
- If viewed as a policy, the solution to a multistage stochastic program is really just a heuristic.

Convergence of SDDP

- Suppose we solve SDDP from an initial state $x(0)$ to give a set of cuts for each stage. The upper bound and lower bound appear to have converged.
- Now starting from $x(0)$ simulate the policy many times with new random seeds. One finds the policy performs worse than expected. Is this perhaps post-decision disappointment?
- It's not, as this still happens with a deterministic two-stage problem.

Two-stage example
(Dowson, 2016)

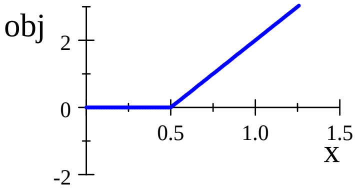

Plot of $f(x) = x + \max\{-x, 3x - 2\}$

Minimizing the function $f(x)$ over $[0, 1]$ is equivalent to

$$\begin{aligned}
\text{P: min} \quad & x + y_2 \\
\text{s.t.} \quad & x - y_1 = 0 \\
& y_1 + y_2 \geq 0 \\
& -3y_1 + y_2 \geq -2 \\
& x \in [0, 1]
\end{aligned}$$

This has optimal solution set $[0, \frac{1}{2}]$.

Benders decomposition

Start at solution $(0, -1)$ to

$$\begin{aligned} \text{MP: } \min \quad & x + \theta \\ \text{s.t.} \quad & x \leq 1 \\ & \theta \geq -1 \\ & x \geq 0 \end{aligned}$$

At $x = 0$, we get

$$\begin{aligned} \text{SP: } \min \quad & y_2 & & \\ \text{s.t.} \quad & y_1 & = x & \quad [\pi] \\ & y_1 + y_2 & \geq 0 & \\ & -3y_1 + y_2 & \geq -2 & \end{aligned}$$

has solution $y_1 = y_2 = 0$, $\pi = -1$, and cut $\theta \geq 0 + -1(x - 0)$.

Benders decomposition

$$\begin{aligned} \text{MP: } \min \quad & x + \theta \\ \text{s.t.} \quad & x \leq 1 \\ & \theta \geq -1 \\ & \theta \geq -x \\ & x \geq 0 \end{aligned}$$

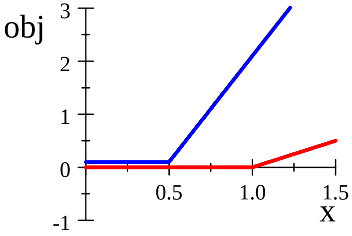

True objective function is blue. Approximation of objective function (red) gives an optimal solution $x = 0$, $\theta = 0$.

Suppose we now simulate the policy

The policy is defined by

$$\begin{array}{rl} \text{MP: min} & x+\theta \\ \text{s.t.} & x \leq 1 \\ & \theta \geq -1 \\ & \theta \geq -x \\ & x \geq 0. \end{array}$$

But MP also has solution $(1,-1)$. At $x=1$, we get

$$\begin{array}{rrll} \text{SP: min} & y_2 & & \\ \text{s.t.} & y_1 & = x & [\pi] \\ & y_1+y_2 & \geq 0 & \\ & -3y_1+y_2 & \geq -2 & \end{array}$$

has solution $y_1 = 1, y_2 = 1$,and optimal value 1. So the cuts at termination do not define an optimal policy if stage problems have alternative solutions.

What is the missing cut?

We require enough cuts possibly at every optimal solution to a stage problem.

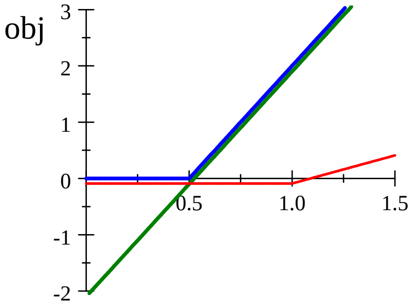

If we add a cut $\theta \geq 3x - 2$ at $x = 1$ then the overall objective of MP is the maximum of red and green functions = the true objective of the problem (blue).

The takeaway message for SDDP

- SDDP terminates when the lower bound and upper bound on optimal value are sufficiently close.
- This means that the first stage decision is the first stage decision of some policy that is close to optimal.
- SDDP appears to give an optimal policy, but only only for the states it visits. It is not an optimal policy until cuts are computed for every possible state we could be in.
- It is essential in backtesting tore-solve SDDP model with updated reservoir storages as simulation proceeds.
- SDDP policies are random objects based on forward path samples and might give variable results when applied to a given realization (e.g. a backtest).

Summary

Backtesting 2008

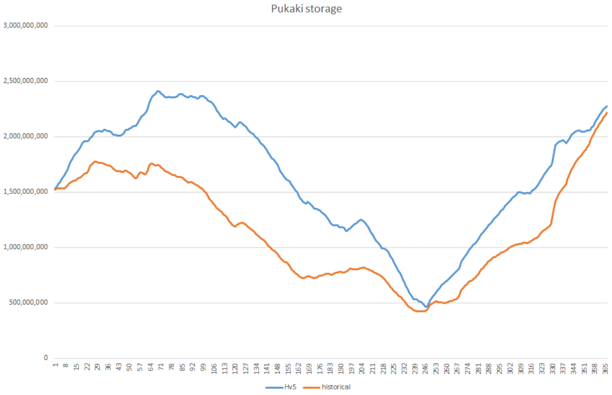

Backtesting 2008

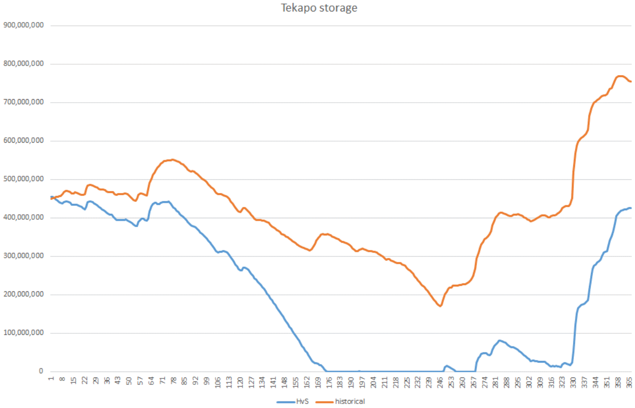

Wholesale prices in 2008

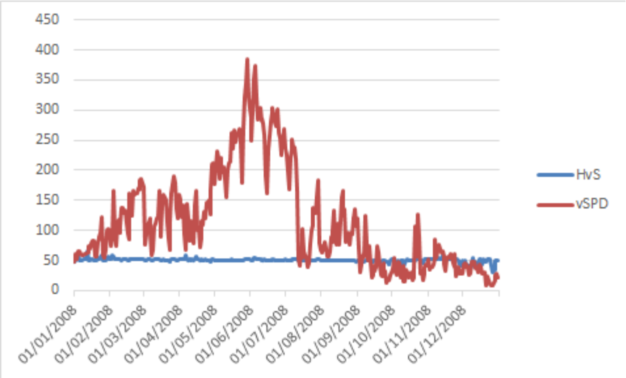

Daily average Auckland prices in 2008.

Backtesting contract prices for 2012
(Joint work with Hamish Mellor)

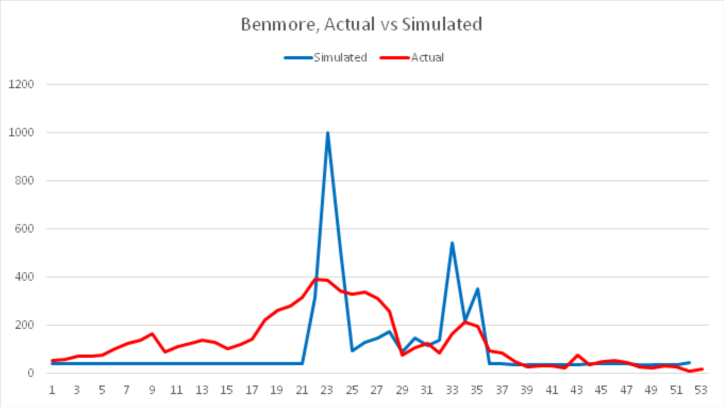

Historical weekly average prices at Benmore in 2008 compared with DOASA predictions at start of year (no rolling horizon).

Marking up thermal offers

GENERATOR	HEAT_RATE	CAPACITY	$/MWh (coal)
Huntly_main_g1_a	1	95	4
Huntly_main_g1_b	30	44	120
Huntly_main_g1_c	50	41	200
Huntly_main_g1_d	75	40	300
Huntly_main_g1_e	100	40	400

Each large thermal unit can be represented in DOASA as several with increasing heat rates.

2008 price comparison

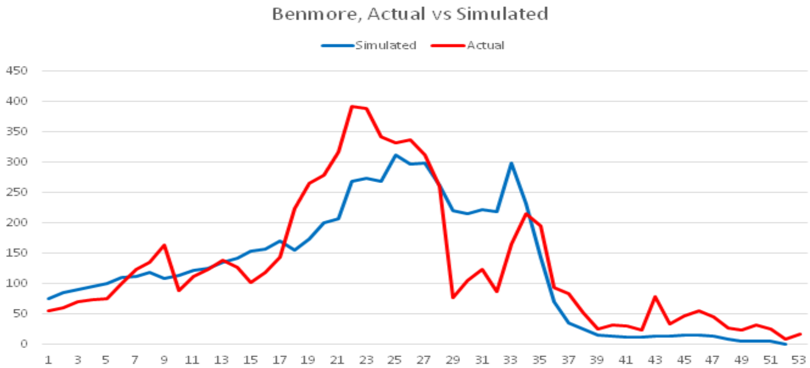

Actual and counterfactual prices for 2008 with marked up fuel costs, no rolling horizon.

2012 comparison

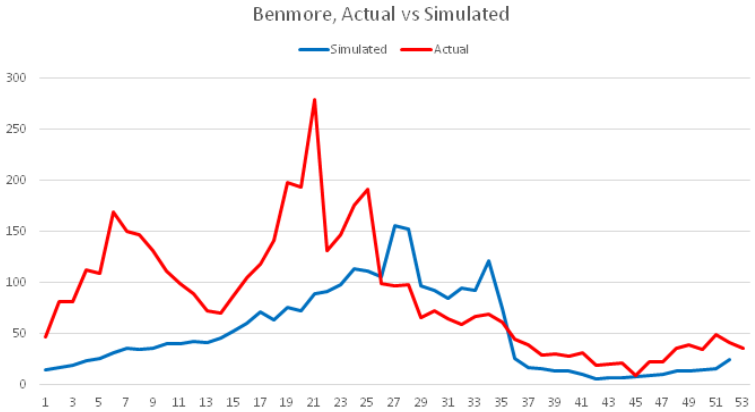

Simulated time average prices in 2012

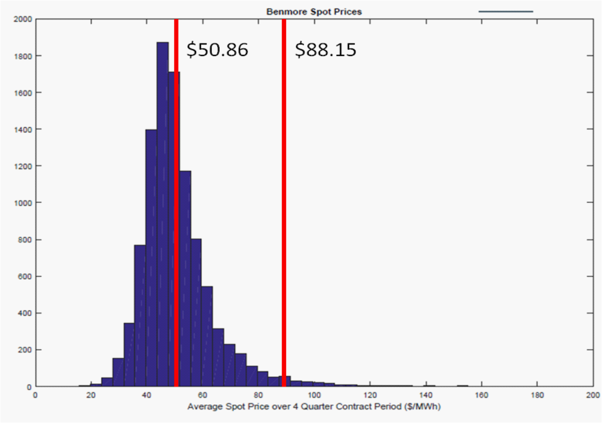

Benmore time average price simulations using DOASA policy for 2012. Highest time average price is $450/MWh, but expected price is $50.86/MWh. CFDs traded 1/1/2012 at $88.15/MWh.

Rooftop solar evaluation
(Joint work with Jerome Clavel, EPFL)

- Run DOASA with different assumptions on level of rooftop solar installations.
- Use NIWA irradiation figures and panel efficiencies to reduce demand in trading periods and times of year.
- Estimate panel costs at (NZ)$500-600 per square metre (accounting for inverters).
- This amounts to $15/m2/year over 35-40 year life.

DOASA results

Assuming 2008 demand, 1 km^2 of panels decreases expected cost by \$6 M. Current capital cost of panels and inverters = \$15M/year (\$15/m^2 / *year* assuming 35 years life).

The End

THE END

References

- Pereira M.V.F. and Pinto L.M.V.G. Multi-stage stochastic optimization applied to energy planning. *Math Prog*, 52(1):359–375, 1991.
- De Matos, V., Philpott, A.B. and Finardi, E.C., Improving the performance of Stochastic Dual Dynamic Programming, *Journal of Computational and Applied Mathematics*, 290, 196-208, 2015.
- Philpott, A. B., Matos, V. L. D., and Finardi, E. C., On Solving Multistage Stochastic Programs with Coherent Risk Measures. *Operations Research*, 61, 957-970, 2013.
- Mellor, H., Electricity Contracts and Price Premia, Project report, Engineeering Science, 2016.

References

- Kapelevich, L., DOASA in Julia, Project report, Engineeering Science, 2016.
- Clavel, J., Optimizing solar energy in New Zealand, Masters Thesis, EPFL, 2016.
- Dowson, O., Dual degeneracy in Benders decomposition and its implication for Stochastic Dual Dynamic Programming, EPOC Technical report, 2016.