Revised manuscript with changes highlighted (With Author Details)

Click here to view linked References

Capacity planning of renewable energy systems using stochastic dual dynamic programming

J. Hole[a,b,c,*], A.B. Philpott[a], O. Dowson[a]

[a]*Electric Power Optimization Centre, University of Auckland, New Zealand*
[b]*Norwegian University of Science and Technology, Norway*
[c]*Norwegian Water Resources and Energy Directorate, Norway*

Abstract

We present a capacity expansion model for deciding the new electricity generation and transmission capacity to complement an existing hydroelectric reservoir system. The objective is to meet a forecast demand at least expected cost, namely the capital cost of the investment plus the expected discounted operating cost of the system. The optimal operating policy for any level of capacity investment can be computed using stochastic dual dynamic programming. We show how to combine a multistage stochastic operational model of the hydro system with a capacity expansion model to create a single model that can be solved by existing open-source solvers for multistage stochastic programs without the need for customized decomposition algorithms. We illustrate our method by applying it to a model of the New Zealand electricity system and comparing the solutions obtained with those found in a previous study.

Keywords: OR in energy; stochastic dual dynamic programming; capacity expansion; multistage stochastic programming; policy graph

1. Introduction

Electricity systems around the world are transitioning to technologies with zero or near-zero carbon emissions (Clarke et al., 2022). The most popular sources of renewable energy (wind and solar) are either intermittent or diurnal and so some form of flexibility is required to match supply and demand when the wind is not blowing or the sun is not shining. One option is thermal peaker plants, which can quickly ramp up and down to balance supply and demand. Another option is

*Corresponding author
Email address: jarand.hole@ntnu.no (J. Hole)

to store and discharge energy from some form of storage, which may be medium-term storage in the form of hydro-reservoirs, or short-term storage in the form of batteries. Since future wind and solar generation are uncertain, the optimal operation of storage is the solution to a stochastic control problem of some complexity. A third option is to invest in interconnectors, through which neighbouring power systems can share excess electricity and flexibility.

In addition to the intermittency challenges, isolated hydro-dominated energy systems like Iceland, Brazil, and New Zealand, face a "dry-year" risk, in which a sustained period of low inflows leaves insufficient water in the reservoirs to meet demand. To compensate, the system must either be willing to periodically shed demand at high cost (and social interruption), build and operate thermal peakers as a back-up source of energy, or over-build intermittent renewable capacity to ensure sufficient supply. In each case, extra transmission may be needed to move power from the newly built generation to demand loads. The motivation for this paper is to further develop models which can be used to correctly size the investment and analyse the operation of each option.

A key challenge when analysing the operation of a hydro-dominated system is the trade-off between the value of using the water stored in a reservoir in the current period, compared with the value of keeping the water for a future period. Most commonly, the operation of the system is modelled as a multistage stochastic program and solved using stochastic dual dynamic programming (SDDP) (Pereira & Pinto, 1991). One such example is the Brazilian energy system, which has used SDDP to plan its operation for over 25 years (Maceira et al., 2008, 2018). Because of the use of both hydro and thermal generation, the problem is often called the *hydro-thermal scheduling problem*.

A number of authors have augmented multistage stochastic operational models with capacity expansion decisions. As a brief sampling of the literature, we point to the following works (Newham, 2008; Rebennack, 2014; Wu et al., 2016; PSR-Inc., 2023; Bruno et al., 2016; Thomé et al., 2019; Lara et al., 2020; Bødal et al., 2022; Jacobson et al., 2023). The literature can be categorized into two broad streams, which we refer to as *static* and *dynamic* investment models.

In *static* investment models, the capacity investment and operation are treated as a two-stage problem, in which the agent chooses a level of new capacity investment, and then evaluates the operation of the energy system with the investment complete. Static models focus on the ideal end-state of the energy system, but not the timing or sequence of decisions that are required to

get there. In its simplest form, a static model enumerates all potential investments and evaluates their capital and operating costs before choosing the best. More sophisticated static models treat the operation of the energy system as the subproblem of a two-stage model that has first-stage capacity investment decisions (modelled by binary or continuous variables) optimized using a form of decomposition. The second-stage is solved, often using SDDP, to guide the solution of the first stage. A unifying feature of static models is that they decompose the investment and operational decisions, and that their operational problems model a finite time horizon. The use of a finite-horizon model to evaluate the steady-state operational cost of the system forces models to use excessively long horizons to mitigate end-of-horizon effects. Moreover, because these models use decomposition, the authors each code a customized algorithm to solve their particular problem.

Static investment models are a typical approach to planning renewable energy systems, e.g., to have net-zero carbon emissions by 2050. The optimal capacity decisions that have to be made up to 2050 to reach this goal will depend on how the resulting system is operated in 2050, and the resulting average annual cost. However, since planners cannot specify now what values the state variables (such as reservoir storage) will have in 2050, it is difficult to justify a finite-horizon second-stage model with a given ending state. Instead it makes sense to evaluate the optimal expected discounted operating cost as if the system is operated over an infinite horizon. This yields a static investment problem with an infinite-horizon operational problem.

Dynamic investment models co-optimize the capacity investment and operation of the energy system over a finite time horizon. In contrast to static models, changes in capacity can occur more than once over the time horizon. Since capacity investments in electricity infrastructure have long lifetimes, the time horizons of dynamic models are typically several decades. Dynamic models are usually formulated as monolithic multistage optimization problems, and typically represent investment decisions using binary variables in each stage. Stochastic versions of these monolithic models represent uncertainty using a scenario tree. Examples of such models are Dominguez et al. (2016), Liu et al. (2017), and Backe et al. (2022). Because of the difficulty in solving large-scale mixed integer programming problems, various approximations or heuristics are often used to compute good solutions.

Static investment models are also useful in a dynamic context. A popular approach to modelling uncertainty in dynamic capacity planning models is to create a *multi-horizon* scenario tree

Kaut et al. (2014). These discriminate between the operational uncertainty that occurs over short time horizons (giving a fine-grain scenario tree) from the longer-run uncertainty that is realized over decades (creating a coarse-grain scenario tree). Such structures are amenable to solution by decomposition (see, e.g., Downward et al. (2020)), in which a subproblem (with a fine-grain scenario tree) is solved in each node of the coarse-grain scenario tree.

In the multi-horizon setting, the subproblems have initial state variables that are not inherited from previous coarse-grained nodes. This approach works well when the fine-grain problem involves only short-term storage (such as batteries that smooth wind and solar generation). Here a model using representative days will give a reasonable estimate of steady-state operating costs. If storage, e.g., in gas storage or hydroelectric reservoirs, is over longer periods to deal with possible energy shortages, then an estimate of steady-state operating costs is more difficult to obtain, since it will depend on the initial and terminal storage levels in any representative year. In this case, the fine-grained subproblems should be optimized as static investment problems with infinite-horizon operational problems that do not assume exogeneous initial and final storage levels.

In this paper, we propose a static capacity investment model that co-optimizes investment in new capacity and the discounted operational cost of the system in the face of uncertain inflows and variable renewable energy. Our model is based on the policy graph modelling framework of Dowson (2020) and can be solved using existing open-source multistage stochastic programming solvers such as SDDP.jl (Dowson & Kapelevich, 2021) without the need for the practitioner to code their own specialized decomposition algorithm. The single, co-optimized, infinite-horizon model formulation is the main novel contribution of this paper. We apply our model to a case study that investigates the capacity expansion of variable renewable energy in the New Zealand electricity system. We consider different investment scenarios, and to verify that the method works, we compare our optimal solutions with those computed by enumeration by Philpott & Downward (2023). The ability to solve our model without the need to implement custom decomposition algorithms is particularly important, because it means that our approach can easily be applied by other researchers to other energy systems and case studies. Thus, this paper also serves as a demonstration of the flexibility of the policy graph modelling approach and of SDDP.jl.

The paper is laid out as follows. In the next section we describe the proposed method for

the integration of capacity investments in SDDP.jl. In Section 3 we describe how we model the variable wind generation from the endogenous wind investment in SDDP.jl. In Section 4 we give a short description of the New Zealand case study, and compare the results from our investment model and previous work in Section 5, before we conclude in Section 7. Appendix A gives the full formulation of the SDDP stage problem and Appendix B provides parameters for the generation plants and load shedding modelled in the case study. More detailed results from the case study are presented in Appendix C and Appendix D.

2. Modelling investments using a policy graph

In this section we introduce the hydro-thermal scheduling problem, explain how to model the finite- and infinite-horizon variants as a policy graph, and expand the policy graph to include investment decisions.

2.1. *Hydro-thermal scheduling*

The classical version of the hydro-thermal scheduling model is a finite-horizon discrete-time stochastic optimal control problem. The *stages* of this problem are indexed $1, 2, \ldots, T$. The *states* of this problem are reservoir levels measured at the end of each stage (denoted X) and *actions* (denoted U) are water releases through generating turbines, flows through spill and river arcs, and dispatchable non-hydro generation. In each stage, each reservoir r experiences a random inflow $\omega_r(t)$. The upper-case notation used for X and U indicate that these are also random variables.

The general formulation of this hydro-thermal scheduling problem (HTP) is as follows:

$$
\begin{aligned}
\textbf{HTP}: \min \quad & \mathbb{E}[\textstyle\sum_{t=1}^{T} C_t(X(t), U(t))] \\
\text{s.t.} \quad & X(t) = f_t(X(t-1), U(t), \omega(t)) && t = 1, 2, \ldots, T \\
& X(0) = \bar{x} \\
& U(t) \in \mathcal{U}_t(x(t-1), \omega(t)) && t = 1, 2, \ldots, T \\
& X(t) \in \mathcal{X}_t && t = 1, 2, \ldots, T.
\end{aligned}
$$

Here, $C_t(X(t), U(t))$ is the cost of meeting demand (possibly with load shedding) in stage t from dispatchable generation. For example, to model load shedding, C_t could be a higher-fidelity

model of the following form:

$$\begin{aligned} C_t(x,u) = \min_{q} \quad & g_t(u) + \sum_{i\in\mathcal{I}} c_i \cdot q_i \\ \text{s.t.} \quad & \sum_{i\in\mathcal{I}} q_i = d_t - h_t(u) \\ & 0 \le q_i \le \bar{q}_i \qquad \forall i \in \mathcal{I} \end{aligned} \tag{1}$$

where q_i is the quantity of load shedding that is available in a set of tranches $i \in \mathcal{I}$ at cost c_i, d_t is net demand, g_t is the cost of dispatching u, and h_t is the quantity of electricity generated from decisions u.

The finite horizon ignores actions after stage T, but future costs and constraints on $X(T)$ can be modelled by suitable choices of $C_T(X(T), U(T))$ and $\mathcal{X}_T$. Note that these exogeneous modelling choices will affect the optimal solution of **HTP**. We revisit this issue below.

The transition function $f_t(X(t-1), U(t), \omega(t))$ maps the incoming state $X(t-1)$ at the beginning of stage t to a random outgoing state $X(t)$ at the end of stage t. For example, we could set:

$$f_t(X(t-1), U(t), \omega(t)) = X(t-1) + B \cdot U(t) + \omega(t),$$

where the matrix B maps the control flows $U(t)$ to the change in reservoir level. The vector of inflows $\omega(t)$ is assumed to have a finite probability distribution (drawn from the sample space $\Omega(t)$) that is independent of that in any previous stage. Observe that random variables X and U are required to be measurable with respect to the history of the random inflow process, so they must satisfy *non-anticipative* constraints that we have suppressed in the formulation.

2.2. *Policy graphs*

The model **HTP** can be defined by a *policy graph* as described in Dowson (2020). The policy graph defines the dynamic structure of the decision problem we are modelling, showing how actions affect the state variables and how realizations of the random variables are revealed over time.

A policy graph, $\mathcal{G} = (R, \mathcal{N}, \Phi)$, is composed of a root node R, a set of nodes $\mathcal{N}$, and an $|\mathcal{N}|+1$ by $|\mathcal{N}|$ matrix Φ that specifies the transition probabilities, with entries ϕ_{ij}. The *children* of node i are $i^+ = \{j \in \mathcal{N} : \phi_{ij} > 0\}$. The nodes of a policy graph form a time-homogeneous absorbing Markov chain. For each node i there is a probability, $1 - \sum_{j\in i^+} \phi_{ij} \ge 0$, of transitioning to a zero-cost absorbing state, which ends the sequence of transitions. We do not explicitly model the absorbing state.

The solution to a policy graph is a policy, $\{\pi\}_{i\in\mathcal{N}}$, comprised of a decision rule, $\pi_i(x, \omega_i)$, for each node i that maps the incoming state variable x and realization of a random variable ω_i to a feasible control variable u and an outgoing state variable x', where $(u, x') = \pi_i(x, \omega_i) \in \mathcal{X}_i(x, \omega_i)$ for a cost of $C_i(x, u, x', \omega_i)$. The random variable ω_i is independent of x and of ω_j at all other nodes, and has a finite sample space Ω_i with probability mass function $\mathbb{P}_i(\omega)$, $\omega \in \Omega_i$.

We communicate the formulation of a policy graph by giving the set of nodes $\mathcal{N}$, the transition matrix Φ, the state at the root x_R, the random variable ω_i with $\mathbb{P}_i$ for all $i \in \mathcal{N}$, and the subproblems $\mathbf{SP}_i(x, \omega_i)$:

$$\begin{aligned} \mathbf{SP}_i(x, \omega_i): \min_{u,x'} \quad & C_i(x, u, x', \omega_i) \\ \text{s.t.} \quad & (u, x') \in \mathcal{X}_i(x, \omega_i). \end{aligned}$$

The traditional cost-to-go functions are not explicitly modelled; these can be constructed implicitly from the structure of the graph.

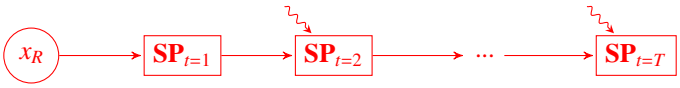

Figure 1: The policy graph structure for **HTP**.

The graph structure for **HTP** is given in Figure 1. There are T nodes. The probability associated with each arc is 1, so that $\phi_{R,1} = \phi_{t,t+1} = 1$ for all $t = 1, \ldots, T-1$. The random variable ω_t is shown as squiggly lines in the policy graph. The first node $t = 1$ is deterministic. The circular root node contains the initial condition $x_R = \bar{x}$. The subproblem for each node t is:

$$\begin{aligned} \mathbf{SP}_t(x, \omega_t): \min_{u,x'} \quad & C_t(x, u, x', \omega_t) \\ \text{s.t.} \quad & x' = f_t(x, u, \omega_t) \\ & u \in \mathcal{U}_t(x, \omega_t) \\ & x' \in \mathcal{X}_t. \end{aligned}$$

We use the policy graph because, by construction, it provides a decomposition of the problem into a structure that is naturally amenable to solution via SDDP-type algorithms, such as the one implemented in `SDDP.jl` (Dowson & Kapelevich, 2021). Under various technical assumptions (including convexity, relatively complete recourse, and the existence of an optimal policy with finite cost), `SDDP.jl` will find an ϵ-optimal policy almost surely in a finite number of iterations.

For more details on the convergence properties, see the work of Philpott & Guan (2008); Dowson (2020); Shapiro & Ding (2020).

2.3. *Infinite horizon*

The need to choose a suitable terminal cost function C_T can be mitigated by solving the discounted infinite-horizon problem:

$$\begin{aligned}
\textbf{HTP} - \infty : \min \quad & \mathbb{E}[\textstyle\sum_{t=1}^{\infty} \rho^t C_t(X(t), U(t))] \\
\text{s.t.} \quad & X(t) = f_t(X(t-1), U(t), \omega(t)) && t = 1, 2, \dots \\
& X(0) = \bar{x} \\
& U(t) \in \mathcal{U}_t(x(t-1), \omega(t)) && t = 1, 2, \dots \\
& X(t) \in \mathcal{X}_t && t = 1, 2, \dots .
\end{aligned}$$

Infinite-horizon discounted cost models can be represented by policy graphs that contain a cycle. An example is shown in Figure 2. In each step, the probability of transitioning from node t to $t+1$ (or from node T to node 1) is ρ where $\rho < 1$. Most commonly, T is chosen so that one loop of the graph represents one year, and the nodes within each loop allow for seasonality in demand and inflows over the course of each year. Importantly, the subproblems $\mathbf{SP}_t$ do not change from the finite-horizon model, only the structure of the graph changes, so that $\phi_{R,1} = 1$, $\phi_{T,1} = \rho$, and $\phi_{t,t+1} = \rho$ for all $t = 1, \dots, T-1$.

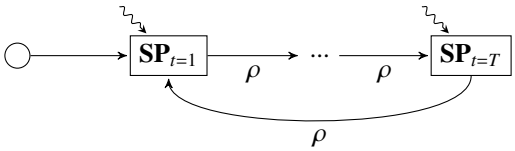

Figure 2: The policy graph structure for **HTP** $- \infty$.

2.4. *Investments*

We can extend the **HTP** $- \infty$ problem to include an investment variable u_{inv} as follows:

$$\begin{aligned}
\textbf{INV} - \textbf{HTP} - \infty : \min \quad & c_{inv}{}^{\top} u_{inv} + \mathbb{E}[\textstyle\sum_{t=1}^{\infty} \rho^t C_t(X(t), U(t))] \\
\text{s.t.} \quad & X(t) = f_t(X(t-1), U(t), \omega(t)) && t = 1, 2, \dots \\
& X(0) = \bar{x} \\
& U(t) \in \mathcal{U}_t(x(t-1), u_{inv}, \omega(t)) && t = 1, 2, \dots \\
& X(t) \in \mathcal{X}_t(u_{inv}) && t = 1, 2, \dots .
\end{aligned}$$

Here u_{inv} is the vector of new capacity that is installed before the start of stage $t = 1$. The newly installed capacity modifies the feasibility set of the state variables $\mathcal{X}_t$, and it may also modify the feasibility set of the control variable $\mathcal{U}_t$. Here c_{inv} is the investment cost I per unit of u_{inv} [$/MW], assuming that the invested capacity has an infinite lifetime. If the lifetime is finite, say τ years, we assume that the capacity is reinvested every τ years and set:

$$c_{inv} = I + \sum_{i=1}^{\infty} I \cdot \beta^{i \cdot \tau} = \frac{I}{1 - \beta^{\tau}}.$$

The policy graph in Figure 2 can be amended to accommodate capacity expansion decisions at the beginning of the infinite operating horizon. This gives the policy graph shown in Figure 3. The transition matrix is $\phi_{R,inv} = \phi_{inv,1} = 1$, $\phi_{T,1} = \rho$, and $\phi_{t,t+1} = \rho$ for all $t = 1, \ldots, T-1$. The decisions in $\mathbf{SP}_{inv}$ determine the system capacities, and the decisions in $\mathbf{SP}_t$ are operating decisions optimizing the control of the resulting system over an infinite horizon.

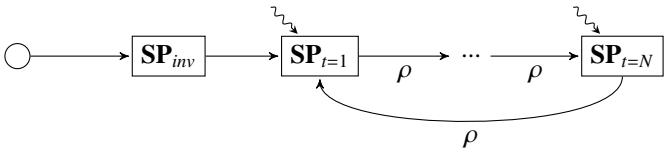

Figure 3: The policy graph structure for **INV** – **HTP** – ∞.

In order to pass the capacity decisions from $\mathbf{SP}_{inv}$ into the cyclic policy graph, we denote the existing reservoir state variables as x_{res} and we add a state variable, x_{inv}, which is set by the investment control variable u_{inv}. Subproblem $\mathbf{SP}_{inv}$ is as follows:

$$\begin{aligned} \mathbf{SP}_{inv}(x): \min_{u,x'} \quad & {c_{inv}}^{\top} u_{inv} \\ \text{s.t.} \quad & u_{inv} \geq 0 \\ & x'_{res} = x_{res} \\ & x'_{inv} = x_{inv} + u_{inv} \end{aligned}$$

The new x_{inv} state variable changes the constraints of each subproblem $\mathbf{SP}_t$. For dispatchable plants, investment increases the maximum output in any dispatch period. The same holds for grid investments, where the maximum power flow on a transmission line is increased. The effect of investment on the output of intermittent capacity such as wind or solar power is less straightforward. In the next section, we describe a methodology that approximates wind generation from invested wind capacity for use in a medium term hydro-scheduling model.

3. Modelling variable renewable investments in the subproblems

In the previous section we showed how capacity investments can be modelled in SDDP.jl using a policy graph. In this section we show how to translate this investment decision into capacity constraints on generation in each stage problem. When generation comes from intermittent sources the constraints are stochastic. We illustrate how we approximate these constraints using wind as an example. The same method will apply for solar power.

In each stage problem, wind generation $w_m(h,t)$ (MW from wind farm m in hour h for week t of the year) is a decision variable that contributes to the electricity supply with a short-run marginal cost equal to 0. The output of wind farm m in week t and hour h depends upon its capacity K_m, but also the (random) amount of wind available at that time. This can be represented by a random factor $\lambda_{m,h,t}(\omega_t)$ that limits the output through:

$$w_m(h,t,\omega_t) \leq \lambda_{m,h,t}(\omega_t) \cdot K_m,$$

where the inequality constraint allows curtailment of wind generation at zero cost.

In our model we approximate the distribution of $\lambda_{m,h,t}(\omega_t)$ by an empirical distribution from Y years of historical wind generation data. The wind farms for which these data are available are grouped into geographical regions r, and $\lambda_{m,h,t}(\omega_t)$ is assumed to be the same value $\lambda_{r,h,t}(\omega_t)$ for each wind farm m in region r. In any region r and hour h in week t, $\lambda_{r,h,t}(\omega_t)$ can take on one of Y values with equal probability. Each value in this distribution (corresponding to year $\omega_t \in \{1,2,\ldots,Y\}$) is obtained by dividing the observed wind generation in year ω_t in hour h in week t in region r by the installed wind capacity in that region. This model assumes wind generation in any hour is independent of that in any previous hour. This ignores the short-term temporal dependence of wind generation from hour to hour that might become important if investment in short-term storage were being considered. The standard approach (see e.g. Liu et al. (2017)) in this case is to use "representative days" of wind generation, which model the dynamics at the expense of introducing some perfect foresight into the decision making.

Since each weekly stage problem in the SDDP model is to be solved many times, hourly demand and hourly wind generation are approximated by piecewise constant load duration curves. The standard approach to estimating these for any year ω_t assumes a level of wind capacity K_r for each region $r \in \mathcal{R}$ and subtracts the wind generation $\lambda_{r,h,t}(\omega_t)K_r$ from the demand $D_r(h,t)$ in

that region to give a net demand:

$$\hat{D}_r(h,t,K_r,\omega_t) = D_r(h,t) - \lambda_{r,h,t}(\omega_t)K_r,$$

for region r in hour h in week t corresponding to the wind in historical year ω_t. Here $\lambda_{r,h,t}(\omega_t)$ is the load factor computed from historical wind generation in region r divided by capacity in year ω_t. The hourly demand $D_r(h,t)$ could also be stochastic, drawn from some distribution. In any case the total system net demand $\sum_{r\in\mathcal{R}} \hat{D}_r(h,t,K_r,\omega_t)$ for any week t is sorted to give a load duration curve with 168 data points that is approximated by load blocks $b \in \{1,2,\ldots,B\}$, each having constant net load equal to average net load over the hours in its block. The blocks are ordered so that net load decreases with b.

It is easy to see that different choices of wind capacity K_r will produce different values of net demand. Moreover the duration curve obtained by sorting $\sum_{r\in\mathcal{R}} \hat{D}_r(h,t,K_r,\omega_t)$ will sort the observations from different hours and historical years into a possibly different order for each choice of K_r. A brute-force approach to capacity optimization would enumerate all potential choices of K_r, $r \in \mathcal{R}$, and solve an SDDP model with this choice.

In our model, the values of K_r, $r \in \mathcal{R}$ are represented by state variables x_{inv}, which means that they are optimized by SDDP rather than enumerated ex-ante. Thus the net load duration curve will change both in average net load in each load block and in the hours allocated to each block as new values of x_{inv} are computed. We approximate this process by assuming that the number of hours $T(t,b)$ in each load block are invariant as x_{inv} varies, even though they might correspond to changing chronological periods. The average wind generation in each block is assumed to change linearly with K_r. To estimate the slope of this line we first construct wind energy duration curves for week t for a finite number of sample values of K_r, assuming in each case that load block b contains $T(t,b)$ hours. Since our wind investments are assumed to be allocated to regions using proportions α_r, the values of K_r chosen are $\alpha_r K$, where K is the total wind investment for the system.[1]

As observed above the chronological hours allocated to b will depend on K, so we denote these hours by $b(K)$. The average power available from wind energy in block b in region r given system investment K is then

$$\bar{w}_r(b,t) = \frac{\alpha_r K \sum_{h\in b(K)} \lambda_{r,h,t}(\omega_t)}{T(t,b)}.$$

[1] For New Zealand we choose $K = 1,2,3,4,5$ GW.

Since we do not have $b(K)$ for all possible values of K, we approximate $\bar{w}_r(b,t)$ by a linear function of K with slope $\mu_{r,b,t}$ and zero intercept, estimated by regressing $\bar{w}_r(b,t)$ against K using the sample values of K for which $\bar{w}_r(b,t)$ has been computed exactly. We can then model wind generation in block b and week t using a variable $w_r(t,b)$ in SDDP and the investment state x_{inv} to represent K, and the constraint:

$$w_r(t,b) \leq \mu_{r,b,t} \cdot x_{inv}.$$

The approximation can be illustrated by showing how it applies to some regions in New Zealand, as a preview to the next section where we discuss the New Zealand system in more detail. Figure 4 shows the fitted linear functions for six random weeks in one historical year representing wind generation for the Otago region of New Zealand for the first (peak), third and fifth load block. A linear regression gives a good representation of the wind generation in each

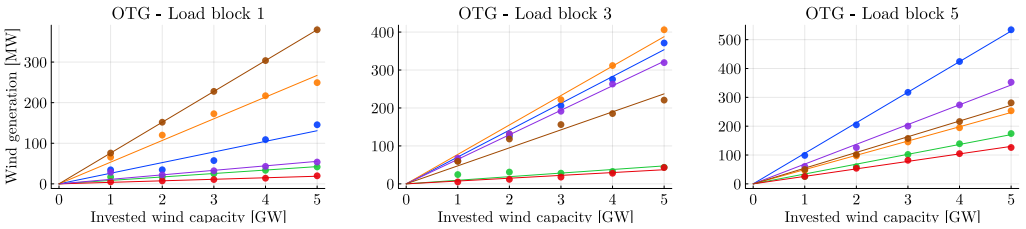

(a) Linear wind representation in Otago for load block 1

(b) Linear wind representation in Otago for load block 3

(c) Linear wind representation in Otago for load block 5

Figure 4: Otago wind generation (MW) versus K (GW), where load block sizes $T(b,t)$ are determined using $\bar{K}$=2.5 GW. Different colours correspond to different weeks of the year, six random weeks in one historical year are shown.

load block at least at the five data points we have chosen. In general, the average generation across the weeks shifts upwards as b increases. In fact, for all the regions, the yearly mean wind generation is always increasing going from the first block to block two, three, four and five.

Figure 5 shows the data points and the corresponding linear regression for the Wellington region. This has a poorer fit, and a few of the estimated slopes in load block 1 are negative. They end up negative because $\bar{w}_r(b,t)$ is discontinuous in K (possibly) jumping downwards when $b(K)$ changes. The impact of this is strongest in the Wellington region because the value of K_r/K is large compared to other regions in the model.

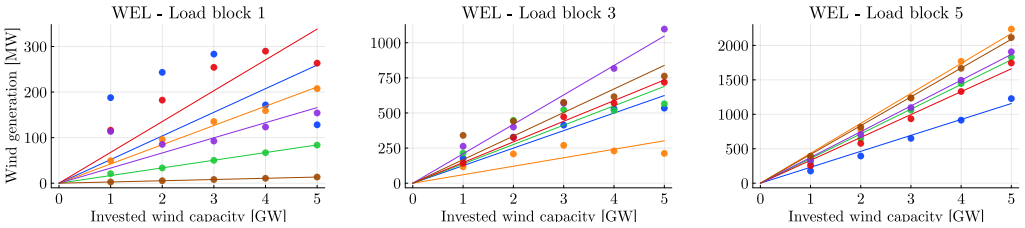

(a) Linear wind representation in Wellington for load block 1

(b) Linear wind representation in Wellington for load block 3

(c) Linear wind representation in Wellington for load block 5

Figure 5: Wellington wind generation (MW) versus K (GW), where load block indices are determined using $\bar{K}$=2.5 GW. Different colours correspond to different weeks of the year, six random weeks in one historical year are shown.

4. A New Zealand case study

In this section, we show the results of applying our model to the New Zealand electricity system. We do this by amending JADE, an existing open-source SDDP model of the New Zealand electricity system (Electric Power Optimization Centre, 2021). JADE is written in Julia (Bezanson et al., 2017) using the JuMP (Lubin et al., 2023) and SDDP.jl (Dowson & Kapelevich, 2021) packages, and it is distributed by the New Zealand Electricity Authority. For this paper, we extended JADE to add the investment decisions and a wind generation model as described in the previous sections.

The New Zealand electricity system is spread over two islands as shown in Figure 6. The details of this system can be found in, for example, Philpott et al. (2019) and so we only give a brief description here. As in Philpott & Downward (2023), JADE in the current paper approximates the full 250-node transmission network by an 11-node model as shown in the right-hand panel. Generation capacities and demand in each island for 2017 are also shown in Figure 6. The islands are joined by a HVDC cable of capacity 1050 MW (shown by the dashed line in Figure 6).

JADE has the ability to model arbitrarily many state variables but following Philpott & Downward (2023) we limit these to storage volumes in Lake Taupo at node WTO, storage volumes in Lakes Tekapo, Ohau and Pukaki at node CAN, and storage volumes in Lake Hawea and Lakes Manapouri-Te Anau at node OTG. The proposed Lake Onslow reservoir is located in node OTG. The details of the power plants associated with these reservoirs are given in Appendix B, and the mathematical formulation of JADE is provided in Appendix A.

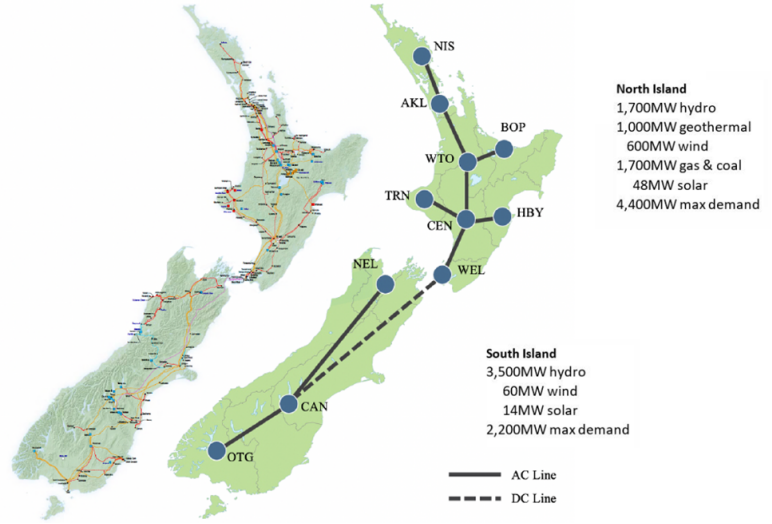

Figure 6: The New Zealand electricity system, showing major transmission lines and the 11 node approximation used in the JADE model.

Our use of JADE enables a straightforward comparison to be made with the wind investment decisions optimized by Philpott & Downward (2023) using enumeration in JADE. They considered three alternative scenarios for a fully renewable electricity system in 2035, in which all capacity decisions were determined up front apart from wind capacity that was optimized. In all three scenarios existing non-renewable electricity plants were closed. In the first case a large pumped hydro storage facility *Lake Onslow* (1.5 GW and 5 TWh), is built to balance the system. The second alternative includes a zero-emission green peaker plant, and the third case assumes existing renewable capacity. We refer to these three cases as *Onslow*, *Peakers* and *Wind only* from now on. The investments in new wind capacity for each case were then found by solving JADE a number of times with different capacity levels and selecting the capacity choice yielding a generation weighted average price (GWAP) equal to the LCOE for the new wind power plants. The results of these experiments provide a useful benchmark for the investments generated by our method.

Of the 11 regions shown in Figure 6, seven have consented wind farms. We allocate new wind capacity proportional to the existing capacity in these regions, as shown in Table 1. This means that the state variable x_{inv} is implemented as the national wind capacity K, and the regional

shares of this capacity are fixed to $\alpha_r K$.

Region	α_r	Representative wind power plant	Capacity [MW]
CAN	0.0814	White Hill	58
CEN	0.0627	Tararua Stage 3	93
HBY	0.0914	Te Uku	28
OTG	0.1926	White Hill	58
TRN	0.000	Waipipi	133
WEL	0.5670	West Wind	142.6
WTO	0.0049	Te Uku	28

Table 1: Representative wind farms and α_r for each region with wind.

Recall that the levelized cost of energy (LCOE) for a power plant is defined as:

$$\text{LCOE} = \frac{\text{Lifetime cost of capacity}}{\text{Lifetime energy produced}}.$$

If we denote the capital cost of 1 MW of wind by I, its lifetime by τ years, its capacity factor by η, and H as the number of hours in a year (8760), then (ignoring maintenance costs):

$$\text{LCOE} = \frac{I}{H \cdot \eta \cdot (1 + \beta + \beta^2 + \ldots \beta^{\tau-1})}.$$

Thus:

$$I = \text{LCOE} \cdot H \cdot \eta \cdot \frac{1 - \beta^\tau}{1 - \beta}.$$

Following Philpott & Downward (2023), we assume LCOE = 65 NZD/MWh. The national wind capacity factor η is estimated from the wind profiles to be 0.355, and we assume $\tau = 20$. When $\beta = 0.9$, this gives an investment cost for our infinite-horizon model of $I = 1.78 \times 10^6$ NZD/MW. The fuel cost of the green peaker plant is set to 160 NZD/MWh.

For the stagewise-independent uncertainty at each operational node, we use the 30 years of historical weekly inflow data from the period 1990-2019. During training, these are adjusted using the Dynamic Inflow Adjustment technique of Philpott & Pritchard (2018) in order to better model the temporal correlation of historical inflows. We initialize the reservoir level state variable with the reservoir levels as they were on 1 January 2020.

5. Results

In this section, we present the results of applying our model to the New Zealand case study briefly introduced in the previous section. The main results presented in this paper are the resulting invested wind capacities for different model runs with our altered model.

5.1. Comparison to Philpott & Downward (2023)

In each of the three case studies, we trained our model for 1,500 SDDP iterations, each iteration of which includes one forward pass and one backward pass. Solution times are on the order of 24 hours, using Julia v1.10, `SDDP.jl` v1.7.0, Gurobi 11, and a single core of a standard laptop. We chose 1,500 iterations based on the results of Figure 7, which shows the lower bound and the estimated rolling mean of the simulation values from the forward pass for all of the three cases.

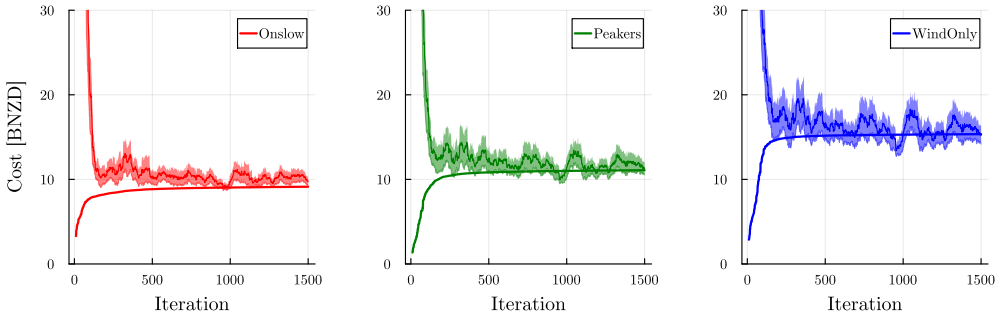

Figure 7: Plots of the lower bound and the estimated rolling mean of the last 100 simulation values from the forward pass for each of the three cases when the model is trained with 1,500 iterations

To compare the investment decisions from our model with the levels published in Philpott & Downward (2023) which, recall, were found by re-solving JADE with different investment levels and choosing the best, Figure 8 plots the first-stage investment decision plotted every 25 iterations, as well as the the benchmark investment levels from Philpott & Downward (2023) (shown by the horizontal lines). In general, the results in Figure 8 indicates that our method for integrating the investment decision works, with resulting wind investments close to the ones found in the previous study. We would not expect the optimal investments in our JADE model to match exactly the optimal investments in the JADE model used in Philpott & Downward (2023) for several reasons. First, the linear wind representation described in Section 3 is a simplification of

the affect that wind investment will have on residual load, compared with Philpott & Downward (2023) where the net load duration curve is re-computed for each level of investment. Second, our model represents tranches of load shedding as a proportion of original load, whereas Philpott & Downward (2023) treats these as proportions of residual load (after wind has been subtracted). Finally, Philpott & Downward (2023) uses round numbers when enumerating different levels of wind investment, whereas the investment decision in our model is a continuous variable.

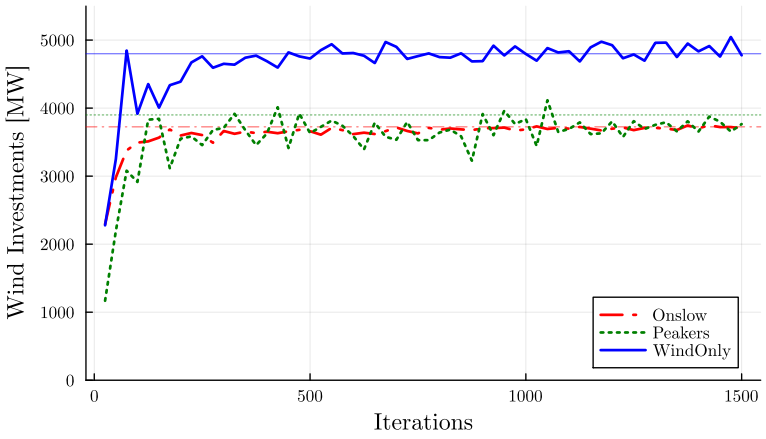

Figure 8: Investment decisions for new wind capacity in the three different cases against the number of SDDP training iterations. Horizontal, dotted lines are those found by Philpott & Downward (2023). The investment decisions are evaluated every 25 iterations.

The variability in the level of wind investment, particular for the *Peakers* case, is an artifact of the SDDP algorithm's cutting plane formulation with multiple near-optimal primal solutions. As shown in Figure 9, the first-stage ($\mathbf{SP}_{inv}$) objective value is flat in the region of the optimal investments. We could rectify this using a regularization term in the first-stage objective function so that, for example, the optimal decision is the minimal quantity of wind investment for a particular cost.

5.2. Incorporating HVDC and peaker capacity

One of the strengths of integrating the investment decisions in the model compared to the method from Philpott & Downward (2023) is that one can easily optimize more investment types, at the cost of increased computational effort. To illustrate this, we conduct an experiment where

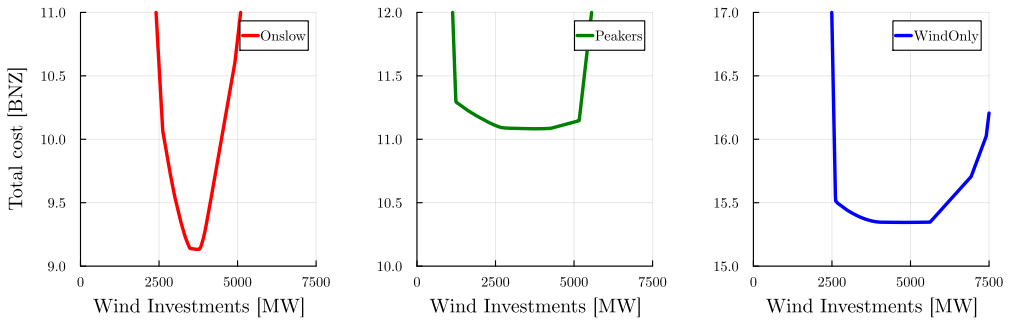

Figure 9: Objective function value of the $\mathbf{SP}_{inv}$ node for different levels of wind investment. Note that the y-axis scales have the same scale but are shifted vertically.

the wind capacity is split into North Island and South Island investments. In addition, we let the model optimize the capacity of the green peaker in the *Peakers* case, and the HVDC transmission capacity between the North Island and the South Island of New Zealand in all of the cases. Thus, the x_{inv} state variable is now a vector with four elements. Solving such a case by enumeration is much more difficult.

Investments and costs	*Onslow*	*Peakers*	*Wind only*
Wind - North Island [MW]	1318	1475	2063
Wind - South Island [MW]	0	0	0
HVDC capacity [MW]	658	0	281
Green peaker [MW]	-	816	-
CAPEX [MNZD]	4689	3680	4846
OPEX [MNZD]	1878	2069	3276
Total [MNZD]	6567	5749	8122

Table 2: Resulting investments and costs from the extended investment model. *CAPEX* is the cost of the deterministic investment node. *OPEX* is the SDDP lower bound less the CAPEX cost.

Table 2 summarizes the results from the extended investment model after training with 2,000 iterations. The total wind investment is lower in all cases than the values in Figure 8. Multiple factors are contributing to this outcome. In our previous model we assumed that the capacity on the HVDC connection between the South Island and the North Island was sufficient to meet all transfers. This enabled a comparison with Philpott & Downward (2023) as shown in Figure 8.

The new model optimizes the HVDC capacity, with the existing capacity as the starting point. Thus, optimal wind investments are lower when wind investments needs to come with following grid investments to access storage in the South Island. In this case, shedding more load is cost efficient compared to extensive grid expansion. Also, as load centers are located on the North Island (and the capacity factor of the representative wind power plants on the North Island is higher than on the South Island), all wind investments are made in the North Island only.

Table 2 shows that increased HVDC capacity is most beneficial in the *Onslow* case, because the Lake Onslow pumped hydro storage is located in the South Island and requires grid capacity to provide its flexibility to the North Island.

In the *Peakers* case, the optimal value of peaking capacity determined by the new model is significantly higher than that assumed in Philpott & Downward (2023). The *Wind only* case has the highest CAPEX, because of the larger wind investments, and the highest OPEX, because of load shedding, of the three cases. The investment costs from building Lake Onslow (estimated in MBIE: New Zealand Battery Project (2022) to be 15,700 MNZD) are not included in the CAPEX cost for the *Onslow* case. Thus the total cost for the *Onslow* case is 22,267 MNZD, which is almost four times the cost for the cheapest alternative, namely the *Peakers* case.

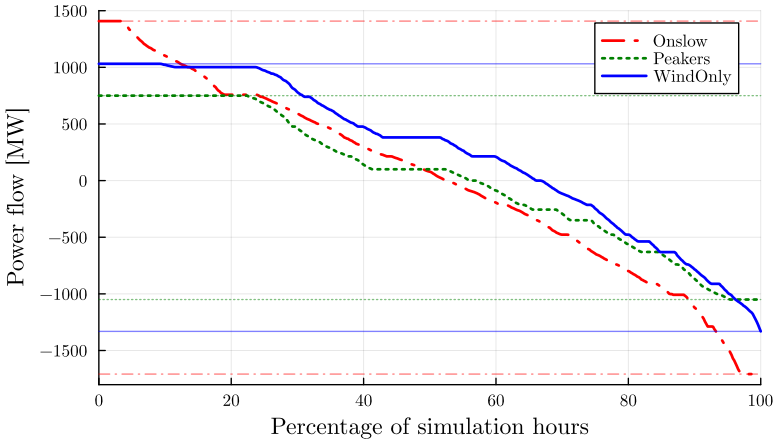

Figure 10: Flow-duration curve for flow on the HVDC line from WEL (North Island) to CAN (South Island) in the three different cases when simulated over the 31 historical inflow years 1990–2020. The horizontal lines show the maximum capacity of the lines in each case, which is the sum of the existing capacity and the invested grid capacity.

Figure 10 shows the utilization of the HVDC cable in each investment case over the 31

historical inflow years 1990–2020. The resulting HVDC capacities are shown as the horizontal lines, while the three plotted lines display the duration curve for the load flow on the HVDC connection with positive load flow defined as flow from the Wellington region (North Island) to the Canterbury region (South Island). The green, horizontal line is equal to the existing capacity because no grid investments are made in the *Peakers* case. The x-axis gives the percentage of simulated hours where load flow is equal to or greater than y MW. In all cases the utilization of the transmission line goes both ways, with the largest flows in the *Onslow* case, where North to South flow is needed for pumping and reverse flow transports Onslow generation north. In the two other cases, HVDC investments are lower, and the full capacity from the North Island to the South Island is utilized more than the full reverse flow.

Figure 10 is complemented by Figure 11, which visualizes the utilization of the six (five without Lake Onslow) main reservoirs in the system. The figure shows that four of the five reservoirs are operated with lower storage levels when other sources of flexibility are in place, like Lake Onslow or a green peaker plant, hence reducing the risk of spillage.

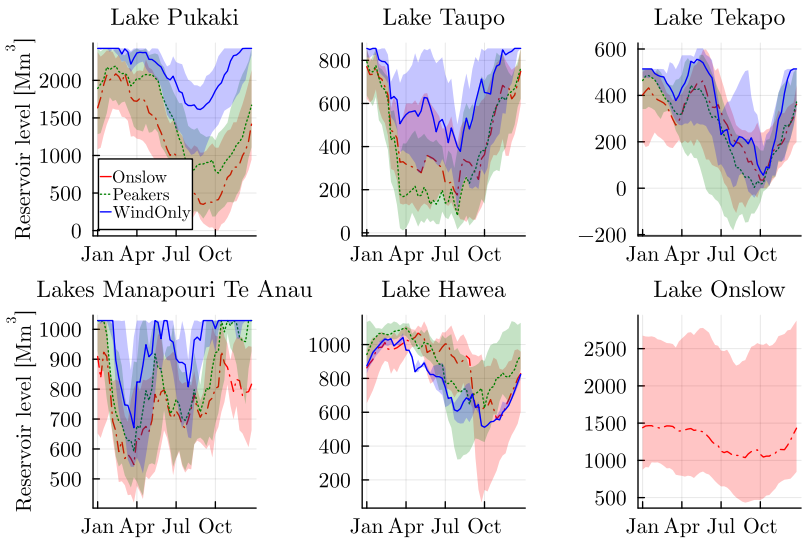

Figure 11: Hydro reservoir levels relative to consented minimum operating levels for the three different cases when simulated over the 31 historical inflow years 1990–2020. The lines are the median reservoir level and each coloured band depicts the 25th and 75th percentile levels. Some reservoirs can have negative storage when the reservoir level falls below the consented operating minimum (incurring a cost penalty).

Figure 12 plots trajectories of marginal water values in the hydro reservoirs. These values can

be taken as an indication of spot energy prices when these are determined by energy constraints rather than capacity constraints. Marginal water values are very low in the *Wind only* case, and higher in the other two cases. In the *Wind only* case, prices are set by periods when the wind is not blowing and load must be shed because conventional capacity and/or grid capacity is insufficient, even though energy in reservoir storage is plentiful giving low marginal water values. In the other two cases, we have increased dispatchable capacity, and prices are set by anticipated energy shortages that are reflected by higher marginal water values. More detailed plots of the storage trajectories and marginal water values for the 31 historical inflow years from 1990–2020 can be found in Appendix C and Appendix D.

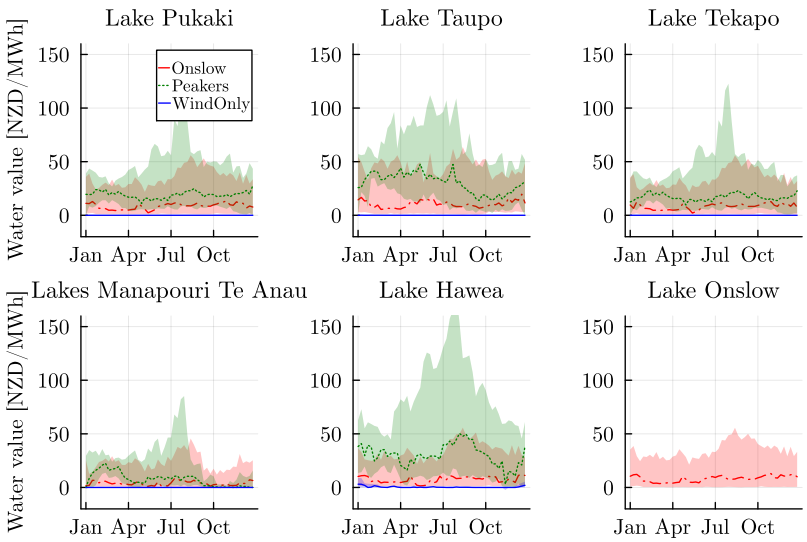

Figure 12: Marginal water values for different reservoirs in the three different cases when simulated over the 31 historical inflow years 1990–2020. The lines are the median marginal water value and each coloured band depicts the 25th and 75th percentiles.

When demand cannot be met, the system must shed load. Figure 13 shows the probability of shedding more than y percent of load in a random week. In all three cases, more than 10 percent of the 1612 simulated weeks have some load shedding. Note that load-shedding includes demand-side flexibility, which we model by a step-wise sequence of tranches that are available at different costs, shown in Appendix B, Table B.6. Moreover, the quantity of shed load is a relatively small quantity ($< 0.4\%$) of total system demand.

The *Wind only* case sheds load in a large proportion of weeks, in periods when the amount of

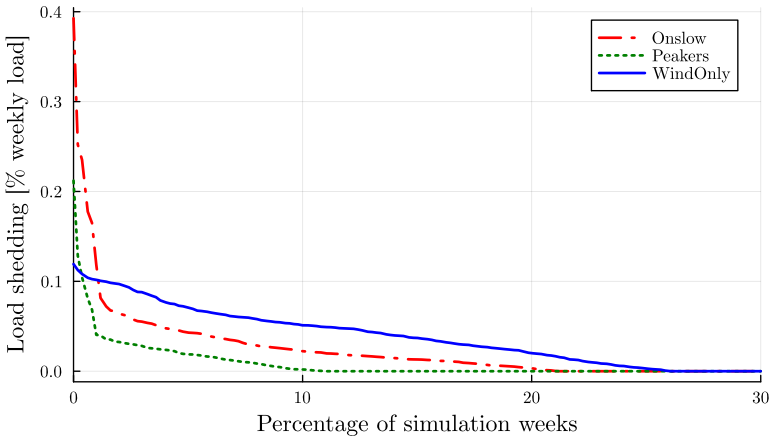

Figure 13: Duration curves of the weekly load shedding in the three cases when simulated over the 31 historical inflow years 1990–2020.

wind generation is unable to meet demand even if all other generators are at full capacity. This means prices are set by shortage costs while marginal water values are generally lower (since generation or grid capacity constraints mean extra water cannot be used to offset this shortage). The prices in these periods provide revenue for wind investment; in other periods prices are zero.

The *Peakers* case provides dispatchable generation capacity to reduce the number of shortage periods. The green trajectory in Figure 13 shows that this results in less load shedding from capacity constraints. Some of the investment that was made in the *Wind only* case is diverted to peakers. The availability of more dispatchable plant makes reservoir operation less conservative making energy shortages more likely, and marginal water values increase.

The *Onslow* case also provides dispatchable generation capacity to reduce the number of shortage periods so the red trajectory in Figure 13 lies below the blue, at least for nearly every week in the simulation. Observe that there are a small number of weeks in the simulations when the load shedding is much higher. As observed in Philpott & Downward (2023), these cases occur when Onslow generation is required but the lake has been emptied by a dry period in the previous year. Marginal water values and load shedding become very high in these weeks.

6. Modelling extensions

The model and solution approach we present in this paper is a good starting point for future research, both the general approach for similar applications, and for the specific JADE model. In this section we list some potential extensions.

6.1. Conditional starting state

Our model requires that we initialize the starting state of the reservoirs with a value for $x(0)$. Even though the policy is cyclic and stationary, the choice of this value impacts the optimal investment decision because we use a discounted objective function instead of long-run average cost. (As an extreme case, consider that the optimal investment decision if we knew the reservoirs were starting empty would be different to the investment decision if we knew the reservoirs were starting full.)

We could improve our model by adding a new node $\mathbf{SP}_{res}$ that sets the outgoing reservoir levels to a random value, passes the investment decision through unchanged, and has no objective function or control variables. The policy graph structure is given in Figure 14.

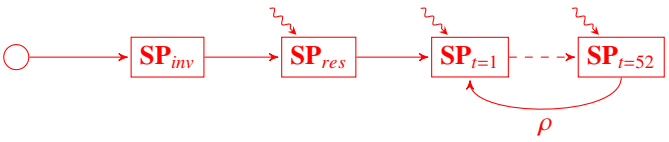

Figure 14: The policy graph structure for randomized starting state.

The random variable in $\mathbf{SP}_{res}$ could be initialized with historical data for the level of the reservoirs on January 1st of each year. If the historical distribution is different to the distribution of reservoir levels in the optimal policy, we could train and simulate the policy to obtain new distribution that is used as input to create a new policy graph.

6.2. Dynamic investments

An alternative to the conditional starting state would be would be to incorporate dynamic investment into our model. One way to achieve this would be the policy graph shown in Figure 15. There are first Y weeks of an operational subproblem in which the agent can make and adjust the investment decisions, followed by a cycle of 52 nodes, corresponding to the infinite horizon problem that we use in our current model. Because capacity investments have a lead

time from design to delivery, the subproblems would need to be adjusted to model a pipeline of investments, rather than a single u_{inv} decision that can be implemented in a single week.

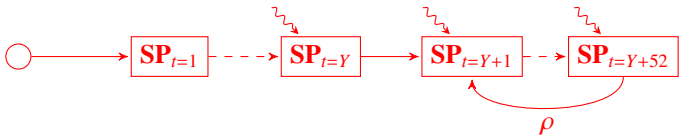

Figure 15: The policy graph structure for a dynamic investment problem.

6.3. Wind and demand modelling

An important avenue of research would be to improve the model we used to convert hourly wind generation data into the weekly load-block data that can be incorporated into an SDDP-type algorithm. An improved wind model could also look to model uncertain demand, for example, by making d_t random in (1). More broadly, one could investigate better statistical correlations of wind and inflow uncertainty due, for example, to climatic conditions such as the El Niño Southern Oscillation.

6.4. Batteries

The wind model we describe in this paper gives an approximation of the duration curves for demand net of wind computed by Philpott & Downward (2023) for each investment choice. Both papers ignore the temporal variation of wind over the short term. Although peaking capacity can be included in the model to cover the peak part of the net load duration curve, the optimization of battery operations (charging and discharging) requires a time series of net demand. This is typically provided to a planning model in the form of representative days of net demand. Future work could study the effect on the policies obtained from using these rather than (approximate) net load duration curves.

6.5. Case study

For this particular case study in New Zealand, it would be useful to relax the assumption that new generation is installed in proportion to existing assets, and expand the model to include an investment variable at each node in the transmission network. Each investment variable increases the dimension of the state space, which reduces the computational tractability.

7. Conclusions

In this paper, we have shown how capacity expansion investments can be included in an infinite-horizon operational SDDP model by augmenting the policy graph to include an investment node. We derived a linear representation of wind generation in different load blocks to link the investment decisions with the subproblem constraints. Additionally, by applying our method to the New Zealand energy system, we have shown that our proposed method is flexible, and that it can be used to analyse a variety of options for expansion within the same system. Because our model is based on the policy graph and SDDP.jl, it did not require us to code a customized decomposition algorithm, which makes our model easier to adopt and implement than alternatives in the literature, and we provided a number of ideas for future extensions.

The results from our case study show that, with the assumptions we made, a peaker plant located in the North Island is a more cost efficient supplement to wind investments than the *Lake Onslow* pumped hydro storage project.[2]

Acknowledgements

J. Hole was supported by the Norwegian Water Resources and Energy Directorate (NVE) and the Department of Electric Energy at the Norwegian University of Science and Technology (NTNU). J.Hole and A. Philpott acknowledge financial support from the NZ MBIE Catalyst grant UOCX2117 - New Zealand-German Platform for Green Hydrogen Integration (HINT). The work is associated with the two Centres for Environment-friendly Energy Research FME CINELDI (RCN grant 257626) and FME NTRANS (RCN grant 296205).

References

Backe, S., Skar, C., del Granado, P. C., Turgut, O., & Tomasgard, A. (2022). Empire: An open-source model based on multi-horizon programming for energy transition analyses. *SoftwareX*, *17*, 100877.

Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. (2017). Julia: A Fresh Approach to Numerical Computing. *SIAM Review*, *59*, 65–98.

Bruno, S., Ahmed, S., Shapiro, A., & Street, A. (2016). Risk neutral and risk averse approaches to multistage renewable investment planning under unncertaitny. *European Journal of Operational Research*, (pp. 979–989).

[2] As we were finishing writing this article, the New Zealand Government announced that they are terminating the NZ Battery Project which was investigating the feasibility of Lake Onslow (Inside Government NZ, 2023).

Bødal, E., Botterud, A., & Korpås, M. (2022). Capacity expansion planning with stochastic rolling horizon dispatch. *Electric Power Systems Research*, *205*, 107729. doi:10.1016/j.epsr.2021.107729.

Clarke, L., Wei, Y.-M., De La Vega Navarro, A., Garg, A., Hahmann, A., Khennas, S., Azevedo, I., Löschel, A., Singh, A., Steg, L., Strbac, G., & Wada, K. (2022). Climate change 2022: Mitigation of climate change. In P. Shukla, J. Skea, R. Slade, A. Al Khourdajie, R. van Diemen, D. McCollum, M. Pathak, S. Some, P. Vyas, R. Fradera, M. Belkacemi, A. Hasija, G. Lisboa, S. Luz, & J. Malley (Eds.), *Climate Change 2022: Mitigation of Climate Change* chapter Energy Systems. Cambridge, UK and New York, NY, USA: Cambridge University Press. doi:10.1017/9781009157926.008.

Dominguez, R., Conejo, A. J., & Carrion, M. (2016). Investing in generation capacity: A multi-stage linear-decision-rule approach. *IEEE Transactions on Power Systems*, *31*, 4784–4794.

Downward, A., Baucke, R., & Philpott, A. (2020). JuDGE.jl: a Julia package for optimizing capacity expansion. Optimization Online. https://optimization-online.org/?p=16779.

Dowson, O. (2020). The policy graph decomposition of multistage stochastic optimization problems. *Networks*, *76*, 3–23. doi:https://doi.org/10.1002/net.21932.

Dowson, O., & Kapelevich, L. (2021). SDDP.jl: a Julia package for stochastic dual dynamic programming. *INFORMS Journal on Computing*, *33*, 27–33.

Electric Power Optimization Centre (2021). JADE.jl: A Julia DOASA Environment. https://github.com/EPOC-NZ/JADE.jl. Accessed: 2023-07-07.

Inside Government NZ (2023). Minister sinks 'hugely wasteful' Lake Onslow hydro scheme. https://insidegovernment.co.nz/minister-sinks-hugely-wasteful-lake-onslow-hydro-scheme. [Accessed 2023-12-06].

Jacobson, A., Pecci, F., Sepulveda, N., Xu, Q., & Jenkins, J. (2023). A computationally efficient benders decomposition for energy systems planning problems with detailed operations and time-coupling constraints. *INFORMS Journal on Optimization*, . doi:https://doi.org/10.1287/ijoo.2023.0005. [Ahead of print].

Kaut, M., Midthun, K., Werner, A., Tomasgard, A., & Hellemo, M., L.and Fodstad (2014). Multi-horizon stochastic programming. *Computational Management Science*, *11*, 179–193.

Lara, C. L., Siirola, J. D., & Grossmann, I. E. (2020). Electric power infrastructure planning under uncertainty: stochastic dual dynamic integer programming (sddip) and parallelization scheme. *Optimization and Engineering*, *21*, 1243–1281.

Liu, Y., Sioshansi, R., & Conejo, A. J. (2017). Multistage stochastic investment planning with multiscale representation of uncertainties and decisions. *IEEE Transactions on Power Systems*, *33*, 781–791.

Lubin, M., Dowson, O., Dias Garcia, J., Huchette, J., Legat, B., & Vielma, J. P. (2023). JuMP 1.0: Recent improvements to a modeling language for mathematical optimization. *Mathematical Programming Computation*, . doi:10.1007/s12532-023-00239-3.

Maceira, M., Penna, D., Diniz, A., Pinto, R., Melo, A., Vasconcellos, C., & Cruz, C. (2018). Twenty years of application of stochastic dual dynamic programming in official and agent studies in brazil-main features and improvements on the newave model. In *2018 Power Systems Computation Conference (PSCC)* (pp. 1–7). IEEE.

Maceira, M. E. P., Duarte, V., Penna, D., Moraes, L., & Melo, A. (2008). Ten years of application of stochastic dual dynamic programming in official and agent studies in brazil-description of the newave program. *16th PSCC, Glasgow,*

Scotland, (pp. 14–18).

MBIE: New Zealand Battery Project (2022). *https://www.mbie.govt.nz/building-and-energy/energy-and-natural-resources/low-emissions-economy/nz-battery/.*

New Zealand Electricity Authority (2023). Electricity Authority Wholesale Market Tools. https://www.emi.ea.govt.nz/Wholesale/Tools/JADE. [Accessed 2023-12-11].

Newham, N. (2008). *Power system investment planning using stochastic dual dynamic programming*. Ph.D. thesis The University of Canterbury, New Zealand.

Pereira, M. V. F., & Pinto, L. M. V. G. (1991). Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, *52*, 359–375.

Philpott, A., & Downward, A. (2023). Using JADE to analyse pumped storage in New Zealand. Technical report, Electric Power Optimization Centre. https://www.epoc.org.nz/papers/JADEOnslow.pdf.

Philpott, A., & Guan, Z. (2008). On the convergence of stochastic dual dynamic programming and related methods. *Operations Research Letters*, *36*, 450–455.

Philpott, A., & Pritchard, G. (2018). EMI-DOASA. Technical report, Electric Power Optimization Centre. https://www.epoc.org.nz/papers/EMI-DOASA.pdf.

Philpott, A., Read, E., Batstone, S., & Miller, A. (2019). The New Zealand electricity market: Challenges of a renewable energy system. *IEEE Power and Energy Magazine*, *17*, 43–52.

PSR-Inc. (2023). OptGen. https://www.psr-inc.com/en/software/. [Online; accessed 2023-11-22].

Rebennack, S. (2014). Generation expansion planning under uncertainty withh emissions quotas. *Electric Power Systems Research*, (pp. 78–85).

Shapiro, A., & Ding, L. (2020). Periodical multistage stochastic programs. *SIAM Journal on Optimization*, *30*, 2083–2102.

Thomé, F., Perez, R. C., Okamura, L., Soares, A., & Binato, S. (2019). Stochastic multistage co-optimization of generation and transmission expansion planning. *arXiv preprint arXiv:1910.01640*, .

Wu, Z., Zeng, P., & Zhang, X.-P. (2016). Two-stage stochastic dual dynamic programming for transmission expansion planning with significant renewable generation and nk criterion. *CSEE Journal of power and energy systems*, *2*, 3–10.

Appendix A. JADE formulation

JADE seeks a policy of electricity generation that meets demand and minimizes the expected cost of thermal generation fuel consumed plus any costs of load reduction. All data are deterministic except for weekly inflows that are assumed to be stagewise independent. The resulting stochastic dynamic programming model is defined as follows. Let $x_j(t)$ denote the storage in reservoir j at the *end* of week t, and let the Bellman function $C_t(\bar{x}, \omega(t))$ be the minimum discounted expected fuel cost to meet electricity demand in weeks $t, t+1, \ldots$, when reservoir storage $x_j(t-1)$ at the start of week t is equal to $\bar{x}_j$ and the inflow to reservoir j in week t is known to be $\omega_j(t)$.

In JADE, a weekly discount factor $\rho = \beta^{1/52}$ is used when going from any stage to the next, where $\beta < 1$ is the annual discount factor. This implies that in each stage there is a probability of $1 - \rho$ of transitioning from stage t to a zero-node 0 where $C_0(\bar{x}, \omega(0)) = 0$.

The Bellman function $C_t(\bar{x}, \omega(t))$ for week t is the optimal solution value of the mathematical program:

$$
\begin{aligned}
C_t(\bar{x}, \omega(t)) = & \\
\min \quad & \textstyle\sum_{i \in \mathcal{N}} \sum_b T(b,t) \left(\sum_{m \in \mathcal{F}(i)} \phi_m f_m(b,t) + \sum_{l \in \mathcal{L}(i)} \psi_{lb} z_i(l,b,t) \right) \\
& + \rho \cdot \mathbb{E}[C_{t+1}(x(t), \omega(t+1))] \\
\text{s.t.} \quad & g_i(y(b,t)) + \textstyle\sum_{m \in \mathcal{F}(i)} f_m(b,t) + \\
& \qquad \textstyle\sum_{m \in \mathcal{H}(i)} \gamma_m h_m(b,t) + \sum_{l \in \mathcal{L}(i)} z_i(l,b,t) = D_i(b,t), && i \in \mathcal{N} \\
& x(t) = \bar{x} - S \textstyle\sum_b T(b,t) \left(A\, h(b,t) + A\, s(b,t) - \omega(t) \right) \\
& 0 \le f_m(t) \le a_m, && m \in \mathcal{F}(i),\ i \in \mathcal{N} \\
& 0 \le h_m(t) \le b_m, \quad 0 \le s_m(t) \le c_m, && m \in \mathcal{H}(i) \\
& 0 \le x_j(t) \le r_j, && j \in \mathcal{J},\ i \in \mathcal{N},\ y \in Y.
\end{aligned}
$$

This description uses the following indices:

Index	Refers to
t	index of week
i	node in transmission network
b	index of load block
m	index of plant
j	index of reservoir
$\mathcal{N}$	set of nodes in transmission network
$\mathcal{F}(i)$	set of green peaker plants at node i
$\mathcal{H}(i)$	set of hydro plants at node i
$\mathcal{L}(i)$	set of load types at node i
$\mathcal{J}$	set of reservoirs.

The parameters are:

Symbol	Meaning	Units
ϕ_m	short-run marginal cost of peaker plant m	\$/MWh
ψ_{lb}	cost of shedding load type l in load block b	\$/MWh
γ_m	conversion factor for water flow into energy	MWs/m^3
$D_i(b,t)$	electricity demand in node i in block b, week t	MW
$T(b,t)$	number of hours in load block b in week t	h
S	number of seconds per hour (3600)	
a_m	thermal plant capacity	MW
b_m	hydro plant capacity	m^3/s
c_m	spillway capacity	m^3/s
r_j	reservoir capacity	m^3
Y	feasible set of transmission flows	
A	incidence matrix of river chain	

The variables are:

Symbol	Meaning	Units
ρ	weekly discount factor	
$f_m(b,t)$	generation of green peaker plant m in load block b in week t	MW
$z_i(l,b,t)$	shed load of type l in load block b in node i in week t	MW
$x_j(t)$	storage in reservoir j at end of week t	m^3
$\bar{x}_j$	known storage in reservoir j at start of week t	m^3
$h(b,t)$	vector of hydro releases in block b, week t	m^3/s
$s(b,t)$	vector of hydro spills in block b, week t	m^3/s
$\omega(t)$	inflow (assumed constant over the week)	m^3/s
$y(b,t)$	flow in transmission lines in load block b in week t	MW
$g_i(y)$	sum of flow into node i when transmission flows are y	MW

The water-balance constraints in the storage reservoirs at the end of week t are represented by:

$$x(t) = \bar{x} - S \sum_b T(b,t)(A\,h(b,t) + A\,s(b,t) - \omega(t)),$$

where $x_j(t)$ is the storage in reservoir j at the end of week t, $s_j(b,t)$ denotes the rate of spill (in m^3/second) in load block b in week t, and $\omega_j(t)$ is the uncontrolled rate of inflow into reservoir

j in week t. We multiply all of these by S to convert to m^3/hour, and then by $T(b,t)$ to give m^3 in each load block. All these are subject to capacity constraints. (In some cases we also have minimum flow constraints that are imposed by environmental resource consents.) The parameter γ_m, which varies by generating station m, converts flows of water $h_m(t)$ into electric power. The same variables and constraints can be used to model pumping of water into a reservoir, except the value of the parameter γ_m is negative to reflect that energy is consumed as water is pumped into a higher reservoir.

Appendix B. System data

We summarize here the features of the New Zealand system represented in JADE. A full data set can be downloaded from New Zealand Electricity Authority (2023). For our study, all existing fossil-fuel plant have been removed from the model.

Generator	Region	Capacity [MW]	Specific power (MW/cumec)
Arapuni	WTO	192	0.462
Aratiatia	WTO	78	0.284
Atiamuri	WTO	84	0.196
Karapiro	WTO	96	0.264
Maraetai	WTO	352	0.526
Matahina	BOP	80	0.595
Ohakuri	WTO	112	0.284
Rangipo	CEN	120	1.960
Tokaanu	CEN	240	1.750
Waikaremoana	HBY	140	3.535
Waipapa	WTO	54	0.139
Whakamaru	WTO	124	0.316

Table B.3: North Island hydro generation stations optimized in JADE model

Generator	Region	Capacity [MW]	Specific power (MW/cumec)
Aviemore	CAN	220	0.310
Benmore	CAN	540	0.818
Clyde	OTG	464	0.518
Cobb	NEL	32	4.405
Coleridge	CAN	39	1.009
Manapouri	OTG	842	1.531
Ohau_A	CAN	264	0.501
Ohau_B	CAN	212	0.417
Ohau_C	CAN	212	0.417
Roxburgh	OTG	320	0.402
Tekapo_A	CAN	27	0.232
Tekapo_B	CAN	154	1.285
Waitaki	CAN	105	0.162
Onslow_Pump	OTG	1500	-7.027
Onslow_Gen	OTG	1500	5.417

Table B.4: South Island hydro generation stations optimized in JADE model

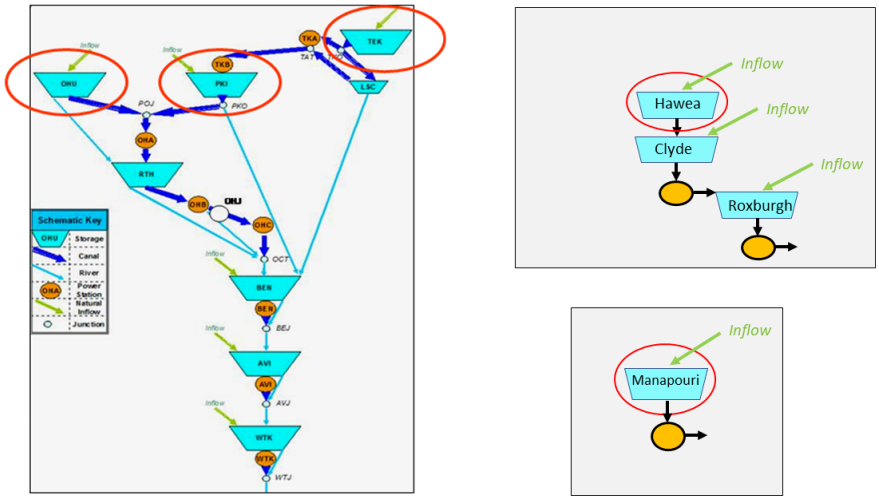

Figure B.16: South Island river chains optimized in JADE. Storage in circled reservoirs are state variables. South Island hydro plant not shown here are assumed to be run-of-river.

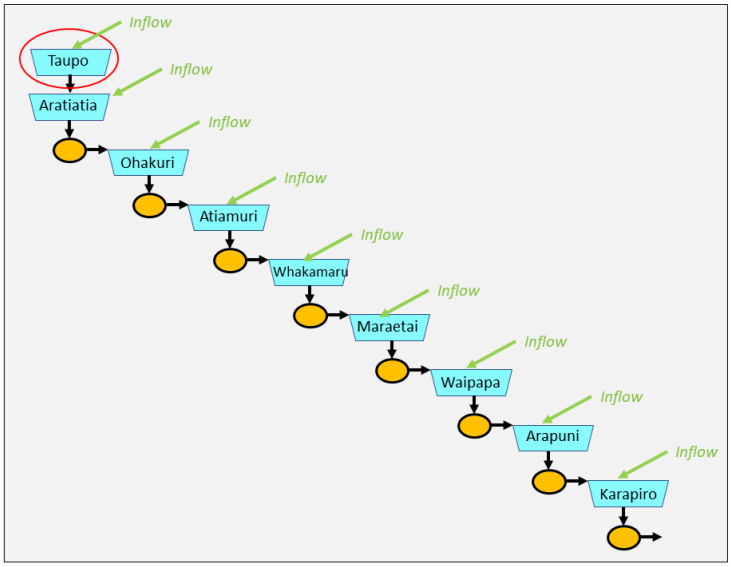

Figure B.17: North Island river chain (Waikato river) optimized in JADE. Storage in circled reservoir (Taupo) is a state variable. North Island hydro plant not shown here are assumed to be run-of-river.

Region	Generation (MW)
AKL	39.54
WTO	864.97
BOP	240.23
CEN	30.16
TRN	59.91
CAN	15.80
NEL	22.24
OTG	40.23

Table B.5: Small fixed generation (including geothermal) for each region where this exists.

Sector	Segment	Sector Proportion	Segment Proportion	Cost [NZD/MWh]
Industrial	Low	0.1	0.25	530
Industrial	Medium	0.1	0.25	740
Industrial	High	0.1	0.5	3180
Commercial	Low	0.3	0.25	5290
Commercial	Medium	0.3	0.25	5290
Commercial	High	0.3	0.5	10580
Residential	Low	0.6	0.25	10580
Residential	Medium	0.6	0.25	10580
Residential	High	0.6	0.5	10580

Table B.6: The cost of load shedding in the JADE model

Appendix C. Storage trajectories

In Figures C.18, C.19, and C.20, we plot the reservoir storage for the six largest reservoirs over the 31 years of historical simulation.

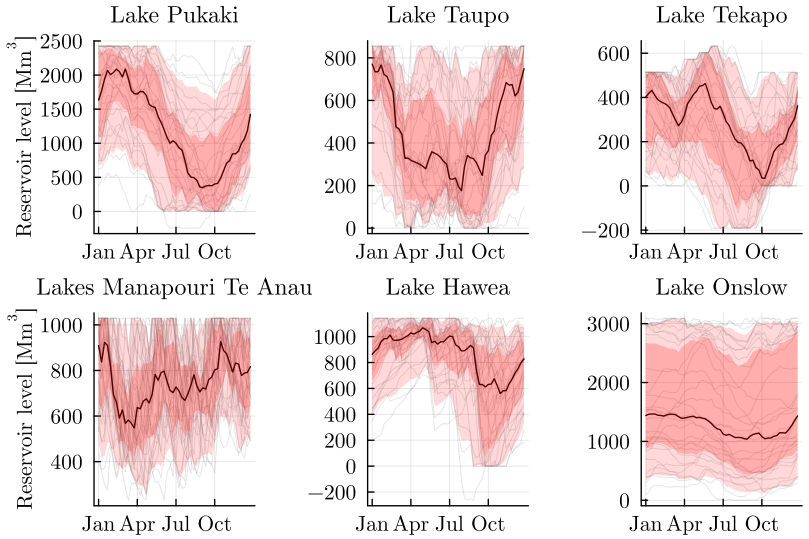

Figure C.18: Reservoir levels relative to the operating minimum for the historical inflow years 1990–2020 for the *Onslow* case. Light gray lines are the 31 trajectories from each simulated year, the solid line is the median, and the shaded bands are the 10–90 and 25–75 percentiles.

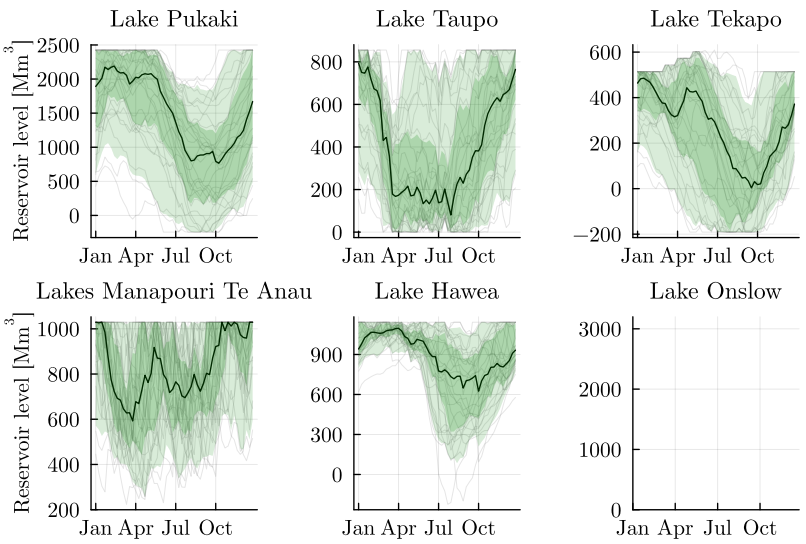

Figure C.19: Reservoir levels relative to the operating minimum for the historical inflow years 1990–2020 for the *Peakers* case. Light gray lines are the 31 trajectories from each simulated year, the solid line is the median, and the shaded bands are the 10–90 and 25–75 percentiles.

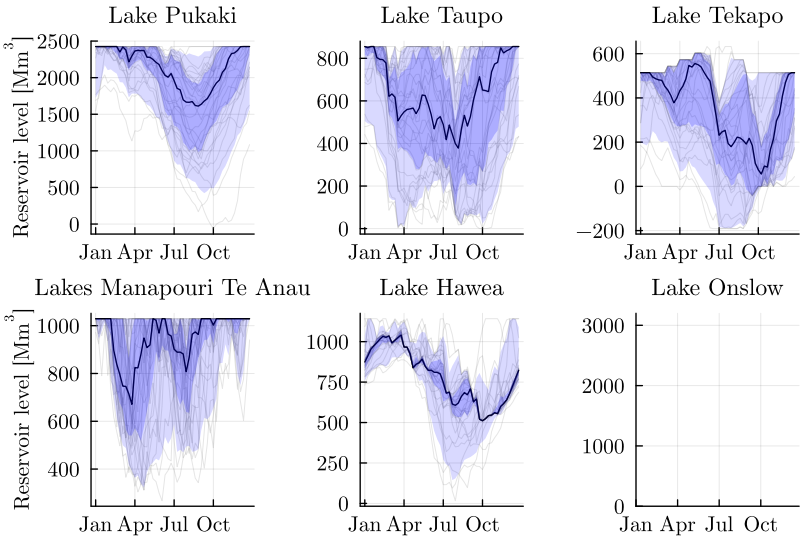

Figure C.20: Reservoir levels relative to the operating minimum for the historical inflow years 1990–2020 for the *Wind only* case. Light gray lines are the 31 trajectories from each simulated year, the solid line is the median, and the shaded bands are the 10–90 and 25–75 percentiles.

Appendix D. Marginal water values

In Figures D.21, D.22, and D.23, we plot the marginal water values for the six largest reservoirs over the 31 years of historical simulation.

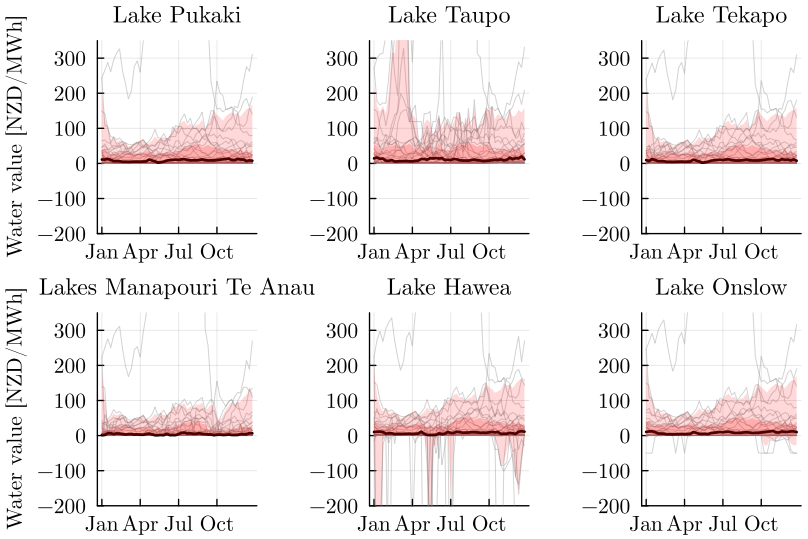

Figure D.21: Marginal water values for the historical inflow years 1990–2020 for the *Onslow* case. Light gray lines are the 31 trajectories from each simulated year, the solid line is the median, and the shaded bands are the 10–90 and 25–75 percentiles.

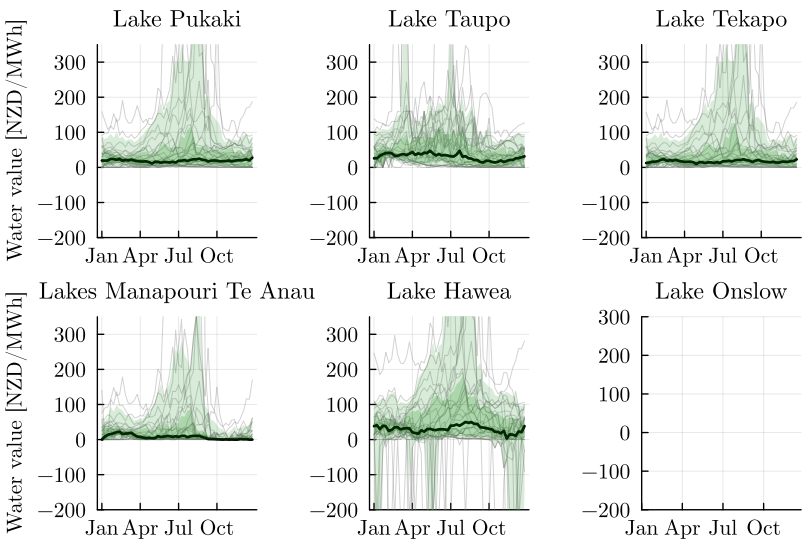

Figure D.22: Marginal water values for the historical inflow years 1990–2020 for the *Peakers* case. Light gray lines are the 31 trajectories from each simulated year, the solid line is the median, and the shaded bands are the 10–90 and 25–75 percentiles.

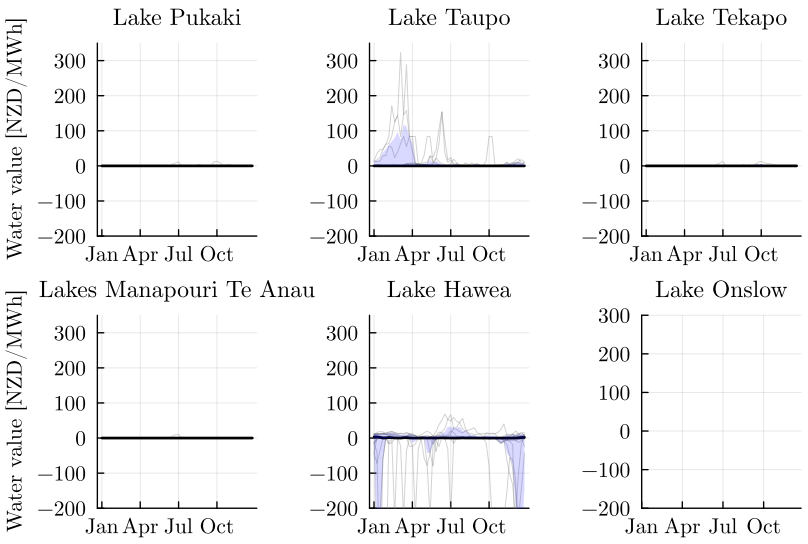

Figure D.23: Marginal water values for the historical inflow years 1990–2020 for the *Wind only* case. Light gray lines are the 31 trajectories from each simulated year, the solid line is the median, and the shaded bands are the 10–90 and 25–75 percentiles.