Hydropower Uncertainty in Capacity Expansion: Calibration Against a Hydro-Thermal Scheduling Model

Madison L. Zegeer[a*], Jannik Haas[a], Andy Philpott[b], Rebecca A.M. Peer[a]

[a]*Department of Civil and Environmental Engineering, University of Canterbury, Christchurch 8041, New Zealand*
[b]*Department of Engineering Science, University of Auckland, Auckland 1010, New Zealand*

**Corresponding author. MAIL. 20 Kirkwood Avenue, Christchurch 8041, New Zealand*

ABSTRACT

The global transition to a zero-carbon energy system requires modelling tools that translate decarbonisation targets into actionable investment and operational strategies. Hydropower's operational flexibility is essential for integrating variable renewables, yet many capacity expansion models (CEMs) rely on deterministic, perfect-foresight assumptions, which can misrepresent system flexibility and produce unrealistic dispatch patterns. We present a model validation framework to evaluate whether a reduced-order hydro uncertainty formulation provides a faithful approximation to stochastic hydro-thermal operational model. GEMSTONE, a two-stage stochastic capacity expansion model with scenario-dependent availability factors and end-of-season storage-gates, is benchmarked against JADE, a high-fidelity SDDP-based hydro-thermal scheduling model, across a range of representative inflow scenarios. Results demonstrate that the storage-gate formulation closely replicates total system operating costs (within $\pm 1\%$) but suppress scenario-specific operational decisions to approximate uncertainty and maintain non-anticipativity. Increasing gate flexibility improves alignment with inflow scenarios but generates perfect-foresight policies. The calibrated reduced-order storage-gate formulation efficiently captures key hydropower uncertainty dynamics. Future work will integrate this approach into high-resolution CEMs with high renewable investments to support large-scale decarbonisation planning. The framework offers energy system modelers a computationally efficient, transferable method for accurately representing hydropower uncertainty in capacity expansion planning studies.

Keywords (5): OR in energy; Stochastic programming; Hydrothermal scheduling; Seasonal storage modelling; Capacity expansion planning

HIGHLIGHTS

- Introduces hydropower uncertainty into capacity expansion via storage-gate formulation
- Calibrates storage-gates against hydro-thermal scheduling model (SDDP) benchmark
- Evaluates performance using system costs and storage operation dynamics
- Demonstrates that storage-gates efficiently capture hydropower uncertainty
- Provides a transferable reduced-order framework for capacity expansion planning

1 Introduction

The global transition to a sustainable, zero-carbon energy system demands not only ambitious policy targets but also advanced modelling tools capable of translating high-level decarbonisation goals into actionable investment in renewable technologies and their operational strategies (Xie et al., 2024). Capacity Expansion Models (CEMs) play a pivotal role in determining the optimal timing, scale, and location of future energy infrastructure investments (Luss, 1982). Recent developments in CEMs have improved high temporal resolution, spatial granularity, and modular scenario analysis, enabling detailed investment planning across regional and global scales (Craig et al., 2022; Babrowski et al., 2014; Giannakidis et al., 2015). However, due to model complexity, computational intensity, and the time pressure to produce relevant policy, many CEMs are deterministic and rely on perfect foresight assumptions (Yue et al., 2018; Haas et al., 2017; Pfenninger et al., 2014; Hall & Buckley, 2016). When CEMs do explore uncertainty, it is typically through large scenario spaces and sensitivity analyses that vary input assumptions but retain deterministic, perfect foresight (Frey et al., 2025; Kachirayil et al., 2022). Perfect foresight is often embedded within model formulations and can significantly influence investment outcomes, yet its implications are less scrutinised than uncertainties in technology or policy assumptions (Analysis of Foresight in Long-Term Energy System Models, 2021). Consequently, the extent to which perfect foresight distorts operational behaviour under uncertainty remains insufficiently quantified in many high-resolution CEMs.

This limitation is particularly important in zero-carbon CEMs where decarbonisation is achieved through investment of low-emission generation and storage technologies such as solar, wind, hydropower, geothermal, batteries, biomass, and/or carbon capture storage (CCS). Large-scale investment in wind and solar generation plays a large role in many investment portfolios due to their low operational costs (Wetzel et al., 2024). As dependence on these variable renewable resources increases, energy storage becomes essential for balancing generation and consumption (Levin et al., 2023; Haas et al., 2019). Hydropower's unique operational characteristics, such as reservoir management and rapid ramping, are critical for balancing these variable renewable resources (Martínez-Jaramillo et al., 2023; Otero et al., 2023). However, CEM deterministic formulations with perfect foresight can lead to unrealistic dispatch patterns and misrepresent system flexibility requirements (Pfenninger et al., 2014; Gøtske & Victoria, 2021), potentially leading to misleading investment signals.

Accurate representation of hydropower systems is further complicated by the need to model intertemporal reservoir dynamics under uncertain inflows. Hydropower generation is highly sensitive to local climatic and geographic factors, such as topography, temperature, snowfall, and precipitation patterns (Osman et al., 2023; Christensen & Kjellström, 2020). Changes in precipitation regimes or glacial melt can either enhance or constrain reservoir inflows, substantially altering generation capacity (Cronin et al., 2018; Yalew et al., 2020; Nohara et al., 2006; Culbertson et al., 2016; Van Vliet et al., 2013). Simplified reservoir dynamics tend to underestimate the need for flexibility or overstate system reliability (Sterl et al., 2021; Kittel et al., 2024). In practice, system operators currently rely on probabilistic forecasts (e.g., 10–30-day outlooks) that introduce forecast errors and necessitate adaptive decision-making dynamics, which deterministic, perfect foresight frameworks cannot represent (Schmidt, 2025). As a result, current CEMs may inadequately capture the operational and structural uncertainties that influence real-world decision-making, where objectives extend beyond pure cost minimisation (Trutnevyte, 2016).

1.1 Literature Review

Stochastic modelling approaches for explicitly incorporating hydropower uncertainty into *operational* power systems planning are well established. Comprehensive reviews document their application to high-renewable systems (Zhou et al., 2016) and widespread use in stochastic optimisation frameworks (Moya et al., 2022; Reddy et al., 2017) and uncertainty-aware operational planning (Fan et al., 2022). Stochastic Dual Dynamic Programming (SDDP) (Pereira & Pinto, 1991) has become the standard for long-term hydro-thermal scheduling (Soroudi, 2013; Su et al., 2025), offering robust policies under variable inflow conditions (Philpott & De Matos, 2012; Fullner & Rebennack, 2023). SDDP models simulate power system operations by incorporating a range of possible future scenarios, thereby capturing the stochastic nature of renewable generation and demand variability (Kiszka & Wozabal, 2025).

Recent CEMs increasingly incorporate uncertainty directly into investment decisions (Jaehnert et al., 2013; Hole et al., 2023), but their computational intensity limits application in large-scale, high-resolution settings (Conejo & Baringo, 2016).
This has motivated reduced-order approximations of intertemporal water constraints. Common heuristics include imposing historical lower-bound constraints on reservoir levels to prevent unrealistic drawdown during dry periods. These methods can improve solution feasibility and mimic long-term conservation (Gøtske et al., 2024), but they still allow anticipative model behaviour and do not fully capture the value of storage under correlated, multi-year inflow uncertainty.

The GEMSTONE model (Ferris & Philpott, 2021) advances these approaches by integrating multi-year inflow uncertainty and limiting anticipative behaviour through scenario-dependent availability factors and hydro reservoir storage governed by end-of-season storage-*gates*, defined as bands around an optimised, *non-anticipative* seasonal target. This allows reservoir levels to vary by scenario within prescribed upper and lower limits, providing a computationally tractable representation of intertemporal hydropower operation. While GEMSTONE represents an improved approximation of hydropower uncertainty, it remains a reduced-order model of the true stochastic process.

1.2 Knowledge Gap

Despite these advances, critical gaps remain. Model comparison studies show that structural differences between power system models can lead to divergent results under harmonised inputs (Gils et al., 2022), but such studies do not isolate the impact of specific assumptions. In particular, the credibility of perfect foresight in representing intertemporal system dynamics (especially in hydro-dominated systems) has not been explicitly tested. No existing work has systematically evaluated whether reduced-order hydropower representations (like storage-gates) can reproduce the operational behaviour of a high-fidelity stochastic scheduling model (e.g., SDDP). We lack quantitative evidence on how much operational realism is sacrificed by these simplifications.

1.3 Contributions

This study addresses these gaps by systematically benchmarking a reduced-order representation of hydropower uncertainty against a high-fidelity SDDP-based hydro-thermal operational model within a capacity expansion context. By explicitly linking stochastic operational detail with computationally tractable expansion modelling, it provides a framework for evaluating the extent to which simplified formulations can reproduce key operational outcomes. The novel contributions of this work are threefold: (i) the development and application of a reduced-order hydropower representation (seasonal storage-gates) within a CEM framework to capture intertemporal uncertainty; (ii) Benchmarking performance of this reduced-order model against a high-fidelity SDDP hydro-thermal model, quantifying how well it reproduces system costs and reservoir trajectories; and (iii) the identification of conditions under which such approximations provide reliable representations of intertemporal storage behaviour.

This motivates the central research questions of this study:

(1) To what extent can seasonal storage-gate formulations in expansion planning models reproduce the system operating costs obtained from a detailed hydropower operations model?
(2) To what extent does the seasonal storage-gate formulation reproduce the scenario-specific reservoir storage trajectories while preserving non-anticipativity relative to detailed operational benchmark?
(3) How does seasonal storage-gate flexibility influence system costs and anticipativity?

If the approximation is validated, this reduced-order formulation would be highly transferable into high-resolution CEM frameworks such as PyPSA (Schumm et al., 2026) or REMix (Canessa et al., 2025) enabling these models to represent seasonal storage effects without incurring the computational burden of a full-scale a SDDP formulation. Moreover, the reduced-order structure enables efficient exploration of alternative and previously unseen inflow sequences through parameterisation informed by SDDP results, providing interpretable insights into reservoir behaviour under uncertainty. The proposed approach therefore supports more robust long-term planning decisions in hydro-dominated power systems under climatic uncertainty.

The methodology is applied to Aotearoa New Zealand which contains a hydro-dominated system with large seasonal reservoirs with significant carryover storage which provides a useful testbed for examining whether such aggregated seasonal formulations capture operational behaviour. Furthermore, Aotearoa New Zealand is committed to net-zero emissions by 2050 and 100% renewable electricity by 2030 (Ministry for the Environment, 2019; International Energy Agency, 2023). Aotearoa New Zealand also has unique operational challenges including long lead times for renewables and storage, intermittent wind and solar generation causing supply-demand imbalances (e.g., the “duck curve”), the phase-out of fossil-fuel dispatchable power, potential hydro shortfalls in dry years, and rising electricity demand from economic growth and electrification of transport and heating (Stelling et al., 2026; Mason et al., 2013; Purdie, 2022; Caruso et al., 2017). Therefore, Aotearoa New Zealand creates a relevant context for assessing the implications of modelling assumptions on long-term planning outcomes.

This results section is structured around the three research questions. **Section 3.1** evaluates the ability of the reduced-order CEM (GEMSTONE) to reproduce the operating costs of the high-fidelity SDDP benchmark (JADE). **Section Error! Reference source not found.** examines the extent to which seasonal storage-gates reproduce scenario-specific reservoir storage trajectories. Section **Error! Reference source not found.** analyses the sensitivity of system outcomes to the level of seasonal storage flexibility allowed by the gate formulation.

2 Methods

2.1 Modelling framework

We adopt a model validation framework to evaluate whether a reduced-order seasonal storage-gate formulation provides a faithful approximation of stochastic hydro-thermal operational dynamics. Specifically, we benchmark GEMSTONE (Ferris & Philpott, 2021), a two-stage stochastic CEM that utilises scenario-dependent availability factors and hydro reservoir storage governed by end-of-season storage-*gates*, against JADE (Downward & Philpott, 2021) a high-fidelity SDDP-based hydro-thermal operational model.

The methodology framework is shown in **Figure 1.** Structural consistency is ensured by aligning input data: load blocks partitions, demand profiles, transmission limits, network representations, short run marginal costs (SRMC), value of lost load (VOLL), and existing capacities in the system. This alignment ensures that both frameworks solve the same underlying operational problem. The models are then evaluated under a layered scenario design comprising of (i) common stochastic inflow sequences (i.e. representative hydrological years) and (ii) structural sensitivity (i.e. variations in seasonal storage-gate width). This framework enables systematic evaluation of representation-driven deviations with validation metrics of system cost and storage behaviour.

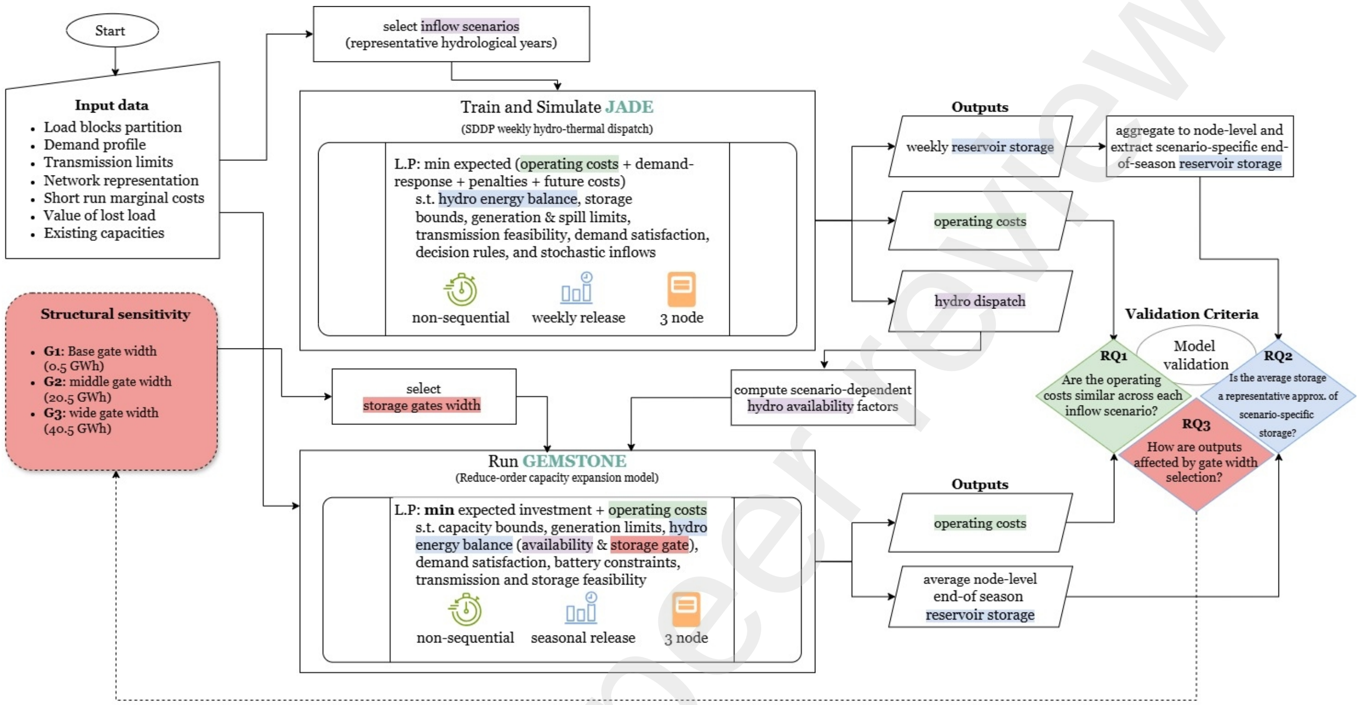

Figure 1 *Methodology overview for Gemstone–Jade model validation*

2.1.1 Reference model: JADE (existing tool)

JADE (Julia DOASA Environment) is a flexible application of the Dynamic Outer Approximation Sampling Algorithm (DOASA), which is an implementation of the SDDP algorithm designed for medium-term hydro-thermal operational scheduling in electricity systems. JADE is a well-established and trusted model employed by the New Zealand Electricity Authority for operational planning and analysis. Accordingly, it is regarded as a reliable benchmark representing baseline real-world system behaviour and is the anchor to our calibration and validation of GEMSTONE. The core purpose of JADE is to compute optimal generation policies that minimise the operational cost of a hydro-thermal system; this includes fuel costs, emission costs, variable operation and maintenance costs, demand response costs, contingent storage penalties, flow penalties, and future cost terms. The decision variables represent operational and state-related choices at each weekly stage of the model. The primary state variable is the volume of water stored in each reservoir at the end of every week; this storage level influences future decisions and is updated dynamically based on inflows, hydro generation, and spill. Operational decision variables include the amount of electricity generated by each thermal and hydro plant, the volume of water released or spilled from reservoirs, and the power flows across transmission lines. Constraints are subject to (i) hydro generation cannot exceed plant capacity or specific power limits, (ii) spill cannot exceed capacity or environmental limits, and (iii) transmission flows do not exceed line capacities. Detailed documentation of JADE can be found in Downward & Philpott (2021)

2.1.2 Reduced-order model: GEMSTONE

GEMSTONE (Ferris & Philpott, 2021) was selected for three key reasons. First, GEMSTONE captures renewable uncertainty through scenario-dependent availability factors, which are based on a suite of representative years. and the reduced-order structure enables efficient exploration of alternative and previously unseen inflow sequences through parameterisation informed by JADE results. Second, hydro reservoir storage is governed by end-of-season storage- *gates*. These gates are defined as bands around an optimised, non-anticipative seasonal target, allowing scenario-contingent reservoir levels to vary within prescribed upper and lower limits. These gates approximate optimal reservoir operator behaviour without the computational burden of full dispatch modelling and limit the deviation of seasonal storage from this average, thereby preventing unrealistic foresight while retaining operational flexibility. Third, if GEMSTONE's approximation is validated such a reduced-order formulation would be highly transferable into high-resolution CEM frameworks such as PyPSA (Schumm et al., 2026) or REMix (Canessa et al., 2025), enabling these models to represent seasonal storage effects without incurring the computational burden of full-scale a SDDP formulation. Detailed documentation of GEMSTONE can be found in Ferris & Philpott (2021).

2.1.3 Model characteristics similarities and differences

Methodologically, GEMSTONE and JADE are both centralised models, aligning with the objectives of central planners (e.g., governments or integrated utilities) rather than competitive market analysis. Both models focus on the electricity sector, emphasising generation, storage, and dispatch. Nonetheless, JADE does not incorporate investment decision-making; it focuses solely on dispatch and scheduling within a fixed infrastructure. GEMSTONE, in contrast, includes a static investment module for capacity expansion planning. Therefore, GEMSTONE is calibrated to operate in dispatch-only mode to enable direct comparison with JADE (additional calibration of inputs and structural consistency is outlined in **Section 2.2**)

Temporally, both models use load duration curves and screening curve approaches but differ in temporal granularity. GEMSTONE divides the year into four seasons with ten load blocks each while JADE operates at a weekly resolution (52 weeks, three load blocks per week). Demand is matched across models despite the differing load-block structures and is specified for each node.

Spatially, both GEMSTONE and JADE adopt a multi-node representation of the New Zealand electricity system comprising the South Island (SI), North Island (NI), and Upper North Island (HAY)). Hydropower generation is located only in the SI and NI. The majority of hydropower storage capacity is situated in the SI, with major reservoirs including Lake Tekapo, Lake Pūkaki, Lake Ōhau, Lake Hāwea, and Lake Manapouri–Te Anau. On the NI, there are just two primary reservoirs: y Lake Taupō and Lake Waikaremoana. The storage capacity of the SI (3045 GWh) is therefore much larger than NI (741.4 GWh).

There are also operational similarities and differences between JADE and GEMSTONE. Given its operational focus, JADE represents hydropower with detailed geographic reservoir schematics and explicit river chains. It also can include environmental flow requirements, contingent storage penalties, and decision rules linking generation and spill to reservoir storage levels. In contrast, GEMSTONE's spatial representation is aggregated to node-level reservoirs in the SI and NI, thereby collapsing individual reservoir dynamics into node-level storage capacity. In this representation, interactions between reservoirs within cascaded river systems (such as upstream–downstream release timing and intermediate storage constraints) are implicitly relaxed and represented only through scenario-dependent hydropower availability factors aggregate storage limits, which may reduce the model's ability to capture spatial coordination of water releases along river chains. GEMSTONE further simplifies spatial representation through a DC load flow approximation with simplified (or neglected) transmission losses. By comparison, JADE represents transmission using a radial network structure with piecewise-linear losses and directional capacities, allowing a more granular representation of network constraints and dynamic inefficiencies.

As previously discussed, GEMSTONE and JADE are significantly different in terms of scenario planning. GEMSTONE uses a two-dimensional scenario index to represent seasonal hydro inflow variability and short-term renewable variability. JADE models inflow uncertainty through historical inflow sequences, with options for stagewise independence or Dependent Inflow Adjustment (DIA). Because GEMSTONE must determine a single set of seasonal storage targets that are feasible across all inflow realisations, it cannot replicate the scenario-specific storage trajectories produced by JADE; instead, it approximates their aggregate seasonal behaviour.

Both models rely on linear programming formulations but differ in decomposition strategies. JADE uses stochastic dual dynamic programming (SDDP) (Pereira & Pinto, 1991) to solve large-scale stochastic scheduling problems efficiently. GEMSTONE applies heuristic decomposition, separating investment and operational decisions, which improves tractability but does not guarantee exact optimality. In terms of implementation, GEMSTONE is developed in GAMS (GAMS Development Corporation, 2023) and solved using CPLEX (Cplex, 2009), with input data managed via structured Excel sheets. JADE is developed in Julia (Bezanson et al., 2017) using JuMP (Lubin et al., 2023) and the SDDP.jl (Dowson & Kapelevich, 2021) packages, also solved with CPLEX (Cplex, 2009), with input data in CSV format and flexible simulation options.

2.2 Inputs and structural consistency (calibration)

Model structures are aligned through consistent input data, including load-block partitions, demand profiles, transmission limits, network representations, SRMC, VOLL, and existing system capacities. This ensures that both frameworks represent the same underlying operational problem. Because JADE is a fixed-capacity operational model, investment decisions were disabled in GEMSTONE by fixing existing capacities and setting all investment variables to zero. Existing capacities of hydro, thermal, wind and geothermal were selected by the user. Annualised capital and fixed O&M costs were also removed from GEMSTONE to ensure that price formation in both models reflects only short-run marginal costs (fuel, variable O&M, and carbon costs). GEMSTONE was also further simplified by removing the Conditional Value at Risk (CVaR) and emission costs.

To ensure that differences in results arise solely from the reduced-order representation rather than divergent modelling assumptions, GEMSTONE was first calibrated directly to JADE. GEMSTONE was trained on hydro dispatch trajectories generated by JADE (without availability factors), reproducing identical system costs across scenarios. Once calibration was achieved, scenario-dependent availability factors and seasonal storage-*gates* were reintroduced in GEMSTONE.

Technology-specific availability factors were harmonised to reflect JADE's representation. Wind and geothermal output were fixed using uniform load factors of 0.35 and 0.85 respectively. Run-of-river and reservoir hydro availability factors were parameterised using JADE dispatch data. Seasonal storage bounds were subsequently tightened using JADE-derived operating envelopes. Structural simplifications were introduced where necessary to match JADE's base-case configuration. Battery variables and constraints were removed from GEMSTONE, as JADE does not include storage technologies of this type. Transmission flow constraints were relaxed and losses set to zero in both models to simplify network representation. Load shedding and associated penalty terms (VOLL) were disabled. Demand and existing capacities for each region (SI, NI, HAY) were aggregated temporally from JADE into GEMSTONE but can be modified if desired.

Collectively, these adjustments ensure structural equivalence between the two models. Remaining differences in results can therefore be attributed primarily to differences in temporal resolution and seasonal storage representation, rather than discrepancies in input data or system configuration.

2.3 Stochastic inflow scenarios

Previous hydropower studies either select extreme years with the highest and lowest generation (Amorim et al., 2020; Savelsberg et al., 2018), or include all years above or below a threshold relative to the long-term average (Demissie & Solomon, 2016) Falchetta et al., 2020). Gil et al. (Gil et al., 2021) instead select weather years using variability-based criteria, preserving most inter-annual and intra-annual hydro inflow variability. Here, we adopt an alternative approach by selecting a set of historical representative years.

A set of inflow scenarios (2005–2017) was selected to represent a range of hydrological conditions. **Figure 2** summarises the hydrological classification and corresponding seasonal inflow profiles for the SI as it the majority of hydropower storage capacity. The top panel shows the assigned year type (wet, normal, dry, very dry) for each year (2005–2017). Hydrological year types are identified directly from inflow data and corresponding demand. The lower panel reports seasonal inflow indices normalised relative to the 2005–2017 mean. Thin lines indicate individual years, while bold lines show the mean seasonal profile for each hydrological class. Seasons are defined as summer (weeks 1–13), autumn (weeks 14–26), winter (weeks 27–39), spring (weeks 40–52).

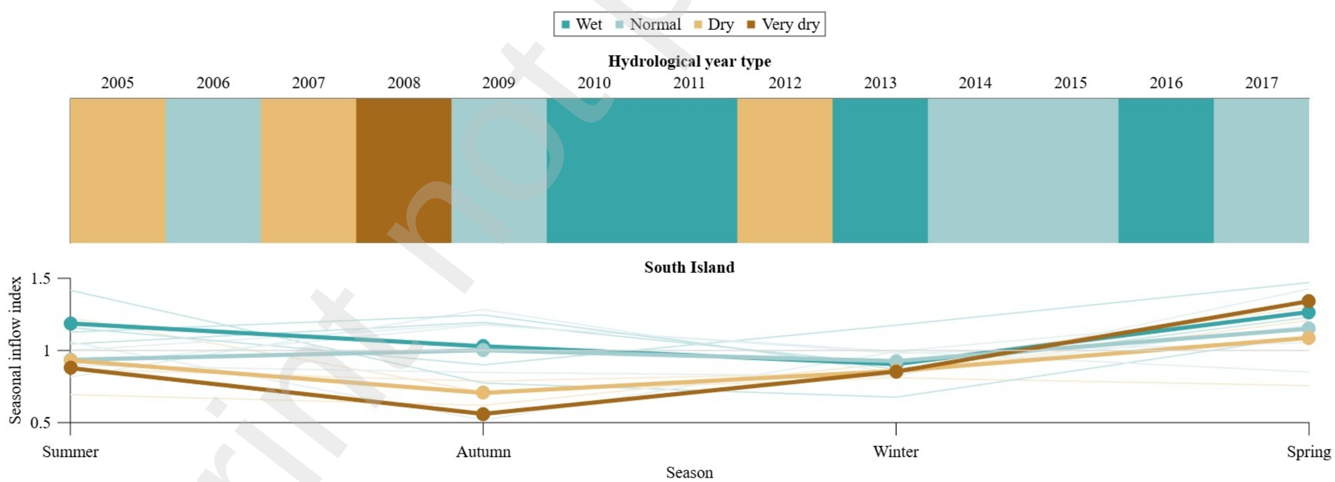

Figure 2 *Hydrological year classification and corresponding seasonal inflow profiles for the SI (2005–2017).*

JADE was trained on these inflow sequences and the policy obtained was then simulated on each historical year to generate in-sample estimates for each of the scenarios. Each simulation was then used to parametrise the hydro scenario-dependent availability factors in GEMSTONE.

Total annual operating costs from GEMSTONE and JADE are compared on an inflow scenario-by-scenario basis to captures aggregate hydro–thermal trade-offs and provide a consistent metric for operational behavioural comparison. Results are shown in **Section 3.1.**

2.4 Reservoir storage policies

GEMSTONE and JADE differ fundamentally in temporal resolution and state representation. As a result, validation focuses on behavioural consistency rather than direct equivalence of state variables. JADE weekly reservoir storage outputs are aggregated to node-level (SI and NI) and final week of each season (weeks 13, 26, 39, and 52) is extracted to influence the boundaries of end-of-season storage in GEMSTONE. GEMSTONE produces average node-level end-of-season reservoir storage to enforce non-anticipative seasonal storage targets, requiring feasibility across the full inflow distribution, whereas the JADE solves a scenario-conditional problem. As a result, exact correspondence in storage trajectories is neither expected nor required.

In practice, hydro operators follow seasonal rule curves with soft bounds, allowing limited deviations while penalizing large departures via spill risk, security-of-supply, or regulations (Electricity Authority, 2026). Without gates, GEMSTONE fully arbitrages summer inflows to winter demand, anticipates future inflows, overvalues hydro flexibility, suppresses thermal capacity, and understates dry-year risk. Storage-gates constrain this behaviour, bounding seasonal deviations to reflect operational practice. Gate width thus serves as a proxy for operational risk tolerance and rule-curve flexibility, not physical storage capacity. Results are shown in **Section Error! Reference source not found..**

2.5 Structural sensitivity (gate width variation)

GEMSTONE represents seasonal storage decisions as first-stage variables in a stochastic optimisation framework. Seasonal storage targets must be feasible across all inflow scenarios, enforcing non-anticipativity. However, strictly fixed seasonal storage levels could be overly restrictive, while fully scenario-specific storage decisions would eliminate the value of a shared seasonal policy. To balance realism and tractability, seasonal storage decisions are allowed to vary within a symmetric band around a prescribed seasonal average level, defined by an exogenous storage-gate width parameter. The gate width therefore governs the degree of inter-seasonal flexibility permitted. A sensitivity analysis is therefore conducted on the storage-gate width parameter.

The gate width is progressively increased, and total annual operating cost is recorded. Total system cost here serves purely as a diagnostic metric of whether the seasonal storage constraint is binding, not as the objective that the gate is designed to optimise. GEMSTONE's storage-gate is a structural modelling device to approximate inter-seasonal flexibility in a tractable two-stage framework; it is not a policy lever. The minimum gate width at which system costs stabilise is identified as the point beyond which the seasonal storage constraint is non-binding. From analysis, three test cases ranging from tight to wide (**Table 1)** are selected to demonstrate the impact on total annual operating costs per inflow scenario to assess the impact of storage-gate formulation limitation on scenario-specific storage adjustments. Results are shown in **Section Error! Reference source not found..**

Table 1 *Storage-gate widths scenarios for structural sensitivity analysis.*

Scenario Name	Gate width (GWh)
G1: Base	$\pm$ 0.5
G2: Middle	$\pm$ 20.5
G3: Wide	$\pm$ 40.5

Additionally, because the storage-gate width is defined in absolute terms and applied uniformly across regions, it induces asymmetric flexibility when expressed relative to regional storage capacity. Given the much smaller seasonal storage volume in the NI, a fixed gate width constitutes a larger share of its capacity than in the SI, leading to a more pronounced increase in operational flexibility in the NI. This scaling effect is evident when gate widths are expressed as a percentage of regional storage capacity (**Table 2**) where identical gate widths represent higher relative flexibility in the NI.

Table 2 *Storage-gate widths expressed as a percentage of seasonal storage capacity in the South Island (SI) and North Island (NI), illustrating how a uniform absolute gate width induces different relative flexibility across regions.*

Gate width (GWh)	% of SI capacity	% of NI capacity
$\pm$ 0.5	~0.02%	~0.067%
$\pm$ 20.5	~0.67%	~2.77%
$\pm$ 40.5	~1.33%	~5.46%

3 Results and discussion

3.1 Storage-gate formulation can recover near-identical operating costs to high-fidelity SDDP benchmark

Total annual operating costs from GEMSTONE and JADE were compared on an inflow scenario-by-scenario basis to capture aggregate hydro–thermal trade-offs and provide a consistent metric for operational behavioural comparison. **Figure 3** summarises the cost-based comparison between GEMSTONE and JADE across all scenarios. The vertical axis enumerates individual inflow scenarios, while the horizontal axis reports the percentage difference in total operating cost relative to the JADE baseline. In the base case, the two models produce nearly identical operating costs, with a maximum deviation of 0.94%, indicating a high degree of numerical agreement between the formulations.

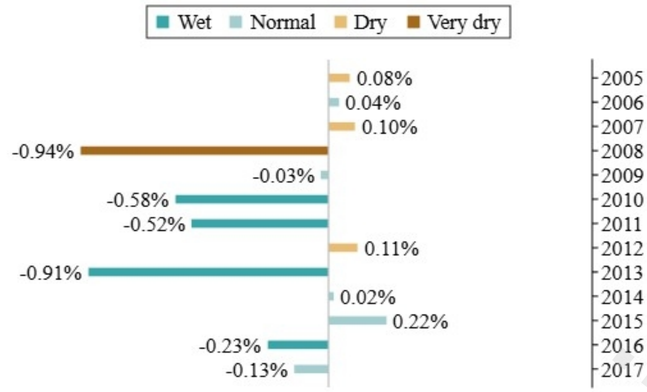

Figure 3 *Percentage difference in total operating cost between GEMSTONE and JADE across each scenario, shown relative to the JADE baseline. Bars are colour-coded by hydrological year type.*

The close cost alignment observed in **Figure 3** indicates that, despite GEMSTONE's reduced temporal resolution, its dispatch decisions closely replicate those of JADE. Minor deviations arise from seasonal aggregation, which abstracts from the weekly operational detail embedded in JADE's scheduling framework. Additional residual differences in operational costs arise from spatial differences in GEMSTONE and JADE. In JADE, reservoirs are subject to explicit geographic and hydraulic constraints, including flow routing along hydropower cascades and scheme-specific water balances. By contrast, GEMSTONE aggregates multiple schemes into a single regional storage node per island. This abstraction relaxes intra-regional constraints (for example, separating Lake Taupō and Lake Waikaremoana in the NI) and effectively allows water to move freely within an island. In reality, flow routing, cascade dynamics, and scheme-level balancing constraints restrict such transfers. The aggregation therefore introduces additional internal flexibility that can underestimate operating costs, particularly under stressed conditions.

GEMSTONE's node-level storage-gate formulation, although computationally convenient, enables physically unrealistic water reallocation within the node that is infeasible in JADE (due to explicit geographic and hydraulic constraints), resulting in slightly lower operating costs especially in very dry years (i.e. 2008) and in some wet years (i.e. 2013). In addition, GEMSTONE represents thermal generation costs at an aggregated technology level (diesel, CCGT, OCGT, coal), which encourages preservation of hydropower as a hedge against high-cost thermal dispatch later in the season. This behaviour elevates the opportunity costs (shadow value) of stored energy under stressed conditions and explains the small, systematic cost advantage observed in some scenarios.

Nevertheless, they dispatch broadly similar volumes of hydropower across scenarios and preserve the economically relevant intertemporal trade-offs under wet, normal, and dry conditions. Prior studies (Li et al., 2025) has shown that aggregate storage representations are often sufficient to capture the economic value of hydropower in long-term planning contexts, even if water balance constraints and scarcity signals are preserved. Our results extend this literature by demonstrating that a seasonal storage-gate formulation can recover near-identical operating costs to a high-fidelity stochastic scheduling model across a wide range of inflow scenarios.

At the same time, the residual discrepancies observed in extremely dry and wet years align with known limitations of spatial and temporal aggregation. Aggregated reservoir modelling, by definition, relax geographic and hydraulic constraints, leading to optimistic (overestimated) flexibility, which become particularly relevant (active constraints) under stress conditions. The slightly lower costs observed in GEMSTONE during dry and wet years reflect this effect and thus expected. Importantly, however, these deviations remain small relative to total system costs, suggesting that the loss of spatial detail does not materially distort costs.

In short, this close alignment in system-level costs suggests that GEMSTONE preserves the economically relevant dispatch behaviour captured by the high-fidelity JADE model with high accuracy. Minor residual differences are driven by spatial node-level aggregation.

3.2 Storage-gates formulation suppresses scenario-specific decisions while preserving non-anticipativity

Figure 4 compares weekly reservoir storage trajectories produced by the high-fidelity JADE operations model with the seasonal storage-gate formulation in GEMSTONE. The horizontal axis shows the week of the year, while the vertical axis reports stored energy as a percentage of total capacity. The figure is organised into three panels showing total system storage (left), South Island (SI) storage (middle), and North Island (NI) storage (right). For each inflow scenario, JADE generates a full weekly storage trajectory for each reservoir. Storage from the SI reservoirs (Lake Tekapo, Lake Pūkaki, Lake Ōhau, Lake Hāwea, and Lake Manapouri–Te Anau) and NI reservoirs (Lake Taupō and Lake Waikaremoana) is aggregated to island level. The blue points show JADE storage levels at the end of each season (weeks 13, 26, 39, and 52) across the 13 inflow scenarios. GEMSTONE's end-of-season storage-gates at these same seasonal boundaries are shown as red error bars.

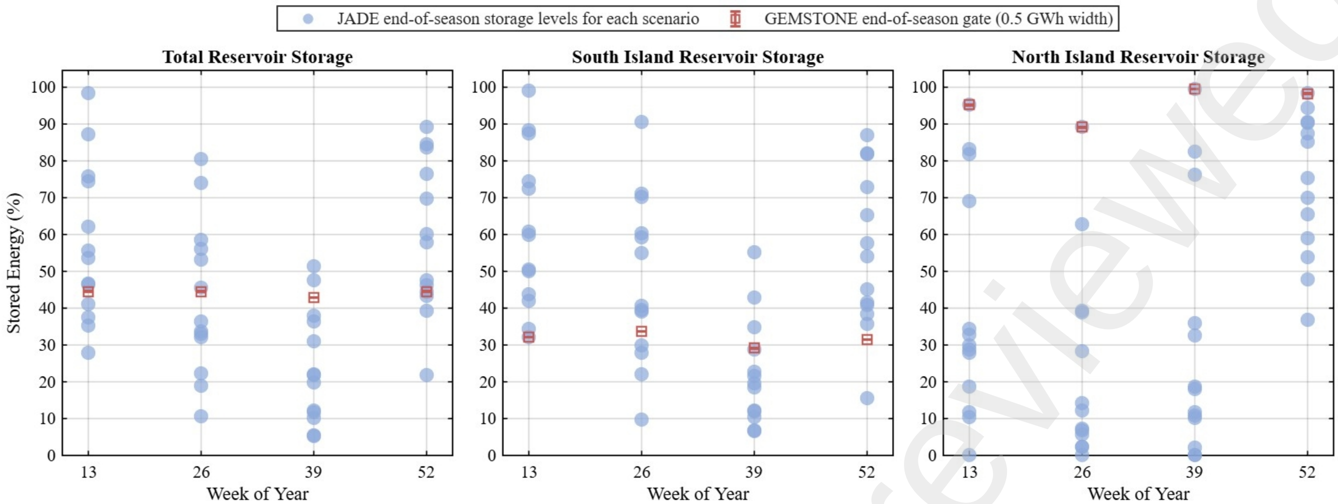

Figure 4 *Comparison of reservoir storage trajectories between JADE and the corresponding seasonal storage representation in GEMSTONE (storage-gate* $= 0.5\ GWh\ width$*). Panels display total system storage (left), South Island (SI) storage (middle), and North Island (NI) storage (right). Blue points represent end-of-season storage levels from JADE for 13 inflow scenarios (weeks 13, 26, 39, and 52), aggregated from individual reservoirs to island level. Red error bars show the corresponding end-of-season storage-gates imposed in GEMSTONE.*

Figure 4 shows that seasonal storage-gate formulation in GEMSTONE produces nearly identical end-of-season storage outcomes across inflow scenarios, in contrast to the wide dispersion generated by the high-fidelity JADE operations model. Across the four seasonal boundaries (weeks 13, 26, 39, and 52), JADE produces a broad distribution of storage outcomes reflecting hydrological variability and operational responses to inflow conditions. Total system storage ranges from approximately 5–100% of capacity across scenarios. The dispersion is particularly pronounced in the SI, where storage spans roughly 10–100% across seasons, reflecting the dominance of large hydro reservoirs and their role in seasonal balancing. The NI also exhibits substantial variability, with storage outcomes ranging from near empty to almost full depending on the inflow scenario.

By contrast, GEMSTONE produces almost identical end-of-season storage levels across all scenarios. With the exogenous storage-gate width set to 0.5 GWh, the model consistently selects storage levels close to the midpoint of each seasonal storage-gate and scenario-specific deviations are negligible. These patterns stem directly from GEMSTONE's hydrological formulation. GEMSTONE enforces non-anticipative seasonal storage targets, requiring feasibility across the full inflow distribution, whereas the JADE solves a scenario-conditional problem. As a result, exact correspondence in storage trajectories is neither expected nor required. Seasonal storage decisions are chosen non-anticipatively in the first stage, while the tight storage-gate (0.5 GWh) supress meaningful scenario-specific drawdown or refill. The model therefore maintains conservative storage levels to ensure feasibility under the dry inflow realisations, smoothing out differences between wet and dry years.

These results demonstrate that the non-anticipative seasonal storage-gate formulation suppresses the scenario-specific reservoir responses observed in the operations model, causing storage trajectories to converge to nearly identical values across inflow scenarios. While this abstraction removes short-term hydrological responsiveness, it also serves a modelling purpose: by constraining end-of-season storage within predefined ranges, the formulation limits perfect foresight and approximates risk-averse reservoir management within a computationally tractable capacity-expansion framework.

3.3 Trade-off between scenario alignment and perfect foresight under increasing gate flexibility

The results in **Section Error! Reference source not found.** demonstrate that a tight seasonal storage-gate formulation (0.5 GWh gate width) suppresses scenario-specific reservoir responses in GEMSTONE. A natural question is therefore whether relaxing the gate constraint can recover the inter-annual variability observed in the high-fidelity JADE operations model. To investigate this, the seasonal storage-gate width was progressively increased, and the resulting costs were evaluated against JADE benchmarks. This experiment examines the trade-off between operational realism and the modelling objective of limiting perfect foresight in the capacity-expansion formulation.

Figure 5 illustrates the relationship between storage-gate width (GWh) and total system operating cost in GEMSTONE, averaged across scenarios. The storage-gate defines the allowable deviation of scenario-specific seasonal storage transfers from first-stage seasonal targets. Across the tested range (0.5–100 GWh), total system costs decrease monotonically as the gate is widened. Cost reductions are largest at small gate widths and diminish as the gate increases. Beyond approximately 40.5 GWh, the gate ceases to bind and total system operating cost stabilises, with no further reductions observed.

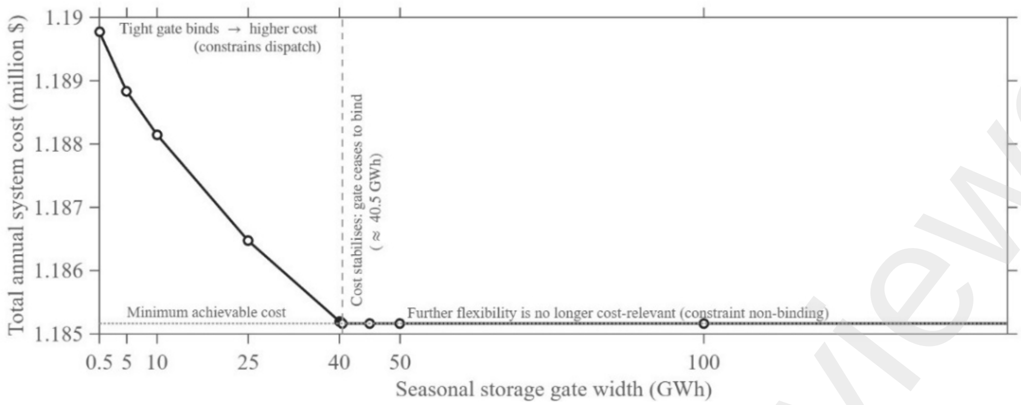

Figure 5 *Total annual system operating cost (averaged across inflow scenarios) as a function of seasonal storage gate width.*

The observed cost–gate relationship in **Figure 5** has two clear regions: (i) a constraint-active region in which widening the storage-gate reduces operating cost, and (ii) a constraint-inactive region beyond ~40.5 GWh, in which further widening yields no additional savings. This plateau occurs because GEMSTONE's coarse temporal structure (four seasons × limited load blocks) and seasonal energy-balance constraints ultimately limit the operational exploitation of extra flexibility. Once the gate exceeds what the model can physically dispatch within its aggregated blocks and transmission limits, additional width no longer creates new cost-reducing dispatch patterns.

An important implication is that the storage-gate does not continuously "add value": it matters only until it ceases to be the marginal binding limitation on inter-seasonal hydro use. The plateau therefore provides a practical operational interpretation: it identifies a level of seasonal flexibility beyond which GEMSTONE behaves as if seasonal storage were effectively unconstrained (within the structure of GEMSTONE). Beyond this point, the persistence of differing seasonal storage trajectories without cost changes indicates multiple cost-equivalent storage pathways. This is expected in the storage-gate formulation: once the inter-seasonal coupling constraint stops binding, storage becomes partially "free" in the objective, so alternative feasible storage paths can exist that preserve the same thermal dispatch and total cost.

To assess these flexibility regimes further, we compare GEMSTONE and JADE outcomes under representative base (tight) gate, middle and wide storage-gate widths, evaluating total operating costs for each inflow scenario individually rather than by scenario averages.

Figure 6 examines how this relaxation affects scenario-level performance by comparing GEMSTONE operating costs with those from JADE for each inflow year scenarios under a base gate width of 0.5 GWh, middle gate width of 20.5, and wide gate width of 40.5 GWh (after which costs plateau in a constraint-inactive region shown in **Figure 5**)

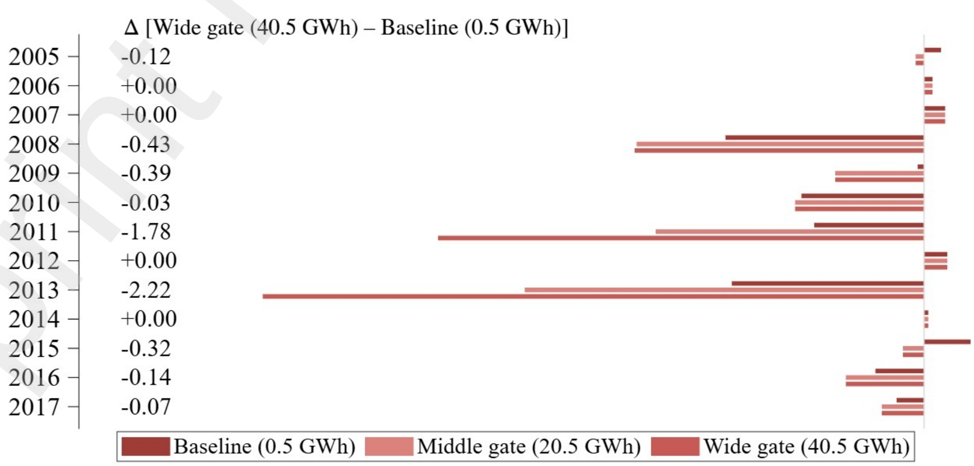

Figure 6 *Percentage difference in total operating cost between GEMSTONE (.5, 20.5, 40.5 GWh storage-gate) and JADE across scenarios, shown relative to the JADE baseline.*

The tight gate enforces limited scenario-specific deviation from the first-stage targets (low clairvoyance and low flexibility), producing GEMSTONE costs that closely track JADE's operational benchmark. In contrast, the loosening the gate allows substantially greater inter-seasonal shifting and anticipativity to use all available water in wet or dry years, enabling GEMSTONE to achieve the same or lower costs than JADE by effectively relaxing the non-anticipative constraint. Costs are further reduced in wide gate scenarios. The wide gate configuration yields lower costs within ±3 % (maximum –3.13 % in 2013), while the tight gate configuration produces systematically higher costs confirming that excess flexibility allows GEMSTONE to "outperform" the detailed operational benchmark due to higher anticipativity to use all available water in wet or dry years.

These results highlight a structural trade-off in the storage-gate formulation. Increasing the gate width improves agreement with the operations model by allowing reservoirs to respond more flexibly to inflow variability. However, once the constraint becomes sufficiently wide, the model effectively anticipates future inflows when determining seasonal storage levels. This behaviour reintroduces the perfect foresight that the storage-gate abstraction was designed to limit. Consequently, while wider gates recover scenario-specific behaviour, they undermine the risk-averse operational approximation embedded in the seasonal storage formulation.

3.4 Limitations and outlook

This study demonstrates that GEMSTONE can closely replicate JADE's operational outcomes under calibrated conditions, validating its potential for long-term planning in hydro-dominated systems, with minor residual differences that are driven by temporal and spatial aggregation. The applicability of this reduced-order representation may also depend on system structure: spatial aggregation of reservoirs to node-level storage is more likely to perform well in systems dominated by large seasonal reservoirs but may be less accurate in systems with long cascaded river chains or strong upstream–downstream operational coupling. Non-anticipative seasonal storage targets combined with the narrow storage-gate limit scenario-specific drawdown and refill, effectively constraining the model's ability to adjust storage dynamically in response to realized inflows. As a result, the formulation suppresses the high-frequency, scenario-dependent storage adjustments that are resolved explicitly in JADE. However, increasing the storage-gate width allows GEMSTONE to exploit anticipativity and produce unrealistically low operating costs relative to JADE.

Future work should explore hybrid approaches that integrate GEMSTONE's reduced-order investment framework with more dynamic reservoir representations, such as informed release regression models or scenario-tree formulations. The storage-gate formulation could also be refined by defining gate widths as region-specific percentages of storage capacity, allowing the level of seasonal flexibility to better reflect underlying reservoir characteristics. Further research should also refine representations of future demand structures including electrification trends at the national level (Canessa et al., 2026) and prosumers at a local level (Steidl et al., 2026) as well as future weather patterns (Wessel et al., 2022; Suomalainen et al., 2022). Finally, exploring a broader set of capacity configurations including higher projected demand levels and systems with larger shares of solar and wind generation (Stelling et al., 2026) would help assess the robustness of the proposed formulation under more diverse future energy system conditions.

Despite these limitations, this comparative analysis offers valuable insights for planners and policymakers developing robust decarbonisation strategies. GEMSTONE's computational efficiency makes it well-suited for scenario-rich capacity expansion studies, while JADE provides operational precision for reliability assessments. Together, these models contribute toward integrated planning tools that balance scalability and realism which is an essential capability for guiding investment and operational decisions in the transition to sustainable energy systems.

4 Conclusions and future work

In this work, we adopt a model validation framework to assess whether a reduced-order seasonal storage-gate formulation provides a faithful approximation of stochastic hydro–thermal operational dynamics. The results demonstrate that seasonal storage-gates provide an effective reduced-order representation of hydro uncertainty in planning models. When calibrated appropriately, gate-based constraints can capture the essential dynamics of reservoir drawdown and refill under stochastic inflows, improving upon simpler approaches that rely solely on historical lower bounds (Gøtske et al., 2024). However, the choice of gate width is critical. Narrow gates restrict the model's ability to respond dynamically to inflow variability, while excessively wide gates permit anticipatory behaviour that exploits future information, leading to unrealistically optimistic operating outcomes. Careful calibration therefore plays a central role in balancing behavioural realism and computational tractability.

Under calibrated conditions, GEMSTONE closely replicates JADE's operational outcomes, with residual differences largely attributable to temporal and spatial aggregation. Nonetheless, the combination of non-anticipative seasonal storage targets and narrow storage-gates necessarily limits scenario-specific drawdown and refill behaviour, constraining the model's ability to adjust storage dynamically in response to realised inflows. As a result, high-frequency, scenario-dependent storage adjustments (explicitly resolved in JADE) are suppressed in the reduced-order formulation. Conversely, increasing gate width relaxes these constraints but allows anticipativity to enter the solution, producing operating costs that fall below those observed in the benchmark SDDP model.

Future work should explore hybrid modelling strategies that retain GEMSTONE's tractable investment framework while introducing more dynamic representations of reservoir operation, such as informed release regression models or limited scenario-tree formulations. Further extensions could refine representations of future demand structures, including national-level electrification trajectories, local-scale prosumer behaviour and evolving climate-driven weather patterns. Finally, evaluating a broader set of capacity configurations (such as higher-demand futures and systems with larger shares of solar and wind generation) would provide additional insight into the robustness of reduced-order hydro representations under increasingly stressed system conditions.

Data availability

Detailed documentation of JADE can be found in Downward & Philpott (2021). Detailed Documentation of GEMSTONE can be found in Ferris & Philpott (2021). Both models are open source. All other data support the findings of this study are included within the article (and any supplementary information files).

Acknowledgments

We appreciate the support of the UC Sustainable Development Goals Scholarship. We thank the Catalyst: Strategic Fund, administered by the Ministry of Business Innovation and Employment of New Zealand, and the German Federal Ministry of Education and Research (grant number 03SF0690) for supporting the project HINT (New Zealand-German Platform for Green Hydrogen Integration).

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process.

During the preparation of this work the author(s) used Enterprise Bing Copilot in order to assist with language editing. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the published article.

References

Amorim, F., Simoes, S. G., Siggini, G., & Assoumou, E. (2020). Introducing climate variability in energy systems modelling. *Energy*, *206*, 118089. https://doi.org/https://doi.org/10.1016/j.energy.2020.118089

Analysis of Foresight in Long-Term Energy System Models. (2021). www.epri.com

Babrowski, S., Heffels, T., Jochem, P., & Fichtner, W. (2014). Reducing computing time of energy system models by a myopic approach. *Energy Systems*, *5*(1), 65–83. https://doi.org/10.1007/s12667-013-0085-1

Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. (2017). Julia: A Fresh Approach to Numerical Computing. *SIAM Review*, *59*(1), 65–98. https://doi.org/10.1137/141000671

Canessa, R., Osorio-Aravena, J. C., Peer, R., Gils, H. C., Wetzel, M., Gulagi, A., Breyer, C., & Haas, J. (2026). Projected hourly and regional energy demand for power heat and transport in New Zealand to 2050. *Scientific Data*, *13*(1), 193. https://doi.org/10.1038/s41597-025-06511-6

Canessa, R., Peer, R., Wetzel, M., Gils, H. C., Gulagi, A., Breyer, C., Osorio-Aravena, J. C., & Haas, J. (2025). From hydro nation to solar giant: New Zealand's renewable build-out and hydrogen export scenarios. *Environmental Research: Energy*, *2*(4), 45019. https://doi.org/10.1088/2753-3751/ae2871

Caruso, B. S., King, R., Newton, S., & Zammit, C. (2017). Simulation of Climate Change Effects on Hydropower Operations in Mountain Headwater Lakes, New Zealand. *River Research and Applications*, *33*(1), 147–161. https://doi.org/https://doi.org/10.1002/rra.3056

Christensen, O. B., & Kjellström, E. (2020). Partitioning uncertainty components of mean climate and climate change in a large ensemble of European regional climate model projections. *Climate Dynamics*, *54*(9–10), 4293–4308. https://doi.org/10.1007/s00382-020-05229-y

Conejo Antonio J. and Baringo, L. and K. S. J. and S. A. S. (2016). Generation Expansion Planning. In *Investment in Electricity Generation and Transmission: Decision Making under Uncertainty* (pp. 61–114). Springer International Publishing. https://doi.org/10.1007/978-3-319-29501-5_3

Cplex, I. I. (2009). V12. 1: User's Manual for CPLEX. *International Business Machines Corporation*, *46*(53), 157.

Craig, M. T., Wohland, J., Stoop, L. P., Kies, A., Pickering, B., Bloomfield, H. C., Browell, J., De Felice, M., Dent, C. J., Deroubaix, A., Frischmuth, F., Gonzalez, P. L. M., Grochowicz, A., Gruber, K., Härtel, P., Kittel, M., Kotzur, L., Labuhn, I., Lundquist, J. K., … Brayshaw, D. J. (2022). Overcoming the disconnect between energy system and climate modeling. In *Joule* (Vol. 6, Number 7, pp. 1405–1417). Cell Press. https://doi.org/10.1016/j.joule.2022.05.010

Cronin, J., Anandarajah, G., & Dessens, O. (2018). Climate change impacts on the energy system: a review of trends and gaps. *Climatic Change*, *151*(2), 79–93. https://doi.org/10.1007/s10584-018-2265-4

Culbertson, A. M., Martin, J. F., Aloysius, N., & Ludsin, S. A. (2016). Anticipated impacts of climate change on 21st century Maumee River discharge and nutrient loads. *Journal of Great Lakes Research*, *42*(6), 1332–1342. https://doi.org/https://doi.org/10.1016/j.jglr.2016.08.008

Demissie, A. A., & Solomon, A. A. (2016). Power system sensitivity to extreme hydrological conditions as studied using an integrated reservoir and power system dispatch model, the case of Ethiopia. *Applied Energy*, *182*, 442–463. https://doi.org/https://doi.org/10.1016/j.apenergy.2016.08.106

Downward, A., & Ade, A. P. (n.d.). *JADE Documentation A Julia DOASA Environment*.

Dowson, O., & Kapelevich, L. (2021). <tt>SDDP.jl</tt> : A Julia Package for Stochastic Dual Dynamic Programming. *INFORMS Journal on Computing*, *33*(1), 27–33. https://doi.org/10.1287/ijoc.2020.0987

Electricity Authority (EA). (2026). *Electricity Market Information*. Https://Www.Emi.Ea.Govt.Nz.

Falchetta, G., Kasamba, C., & Parkinson, S. C. (2020). Monitoring hydropower reliability in Malawi with satellite data and machine learning. *Environmental Research Letters*, *15*(1), 014011. https://doi.org/10.1088/1748-9326/ab6562

Fan, H., Wang, C., Liu, L., & Li, X. (2022). Review of Uncertainty Modeling for Optimal Operation of Integrated Energy System. *Frontiers in Energy Research*, *Volume 9-2021*. https://doi.org/10.3389/fenrg.2021.641337

Ferris, M. C., & Philpott, A. (2021). *Renewable electricity capacity planning with uncertainty at multiple scales*. https://doi.org/https://doi.org/10.1007/s10287-023-00472-0

Frey, U. J., Cao, K.-K., Sasanpour, S., Buschmann, J., & Breuer, T. (2025). The benefits of exploring a large scenario space for future energy systems. *Nature Communications*. https://doi.org/10.1038/s41467-025-67593-9

Fullner, C., & Rebennack, S. (2023). *STOCHASTIC DUAL DYNAMIC PROGRAMMING AND ITS VARIANTS-A REVIEW*.

GAMS Development Corporation. (2023). *General Algebraic Modeling System (GAMS)* (43.3.0).

Giannakidis, G., Labriet, M., Ó Gallachóir, B., & Tosato, G. (Eds.). (2015). *Informing Energy and Climate Policies Using Energy Systems Models* (Vol. 30). Springer International Publishing. https://doi.org/10.1007/978-3-319-16540-0

Gil, E., Morales, Y., & Ochoa, T. (2021). Addressing the effects of climate change on modeling future hydroelectric energy production in Chile. *Energies*, *14*(1). https://doi.org/10.3390/en14010241

Gils, H. C., Gardian, H., Kittel, M., Schill, W. P., Zerrahn, A., Murmann, A., Launer, J., Fehler, A., Gaumnitz, F., van Ouwerkerk, J., Bußar, C., Mikurda, J., Torralba-Díaz, L., Janßen, T., & Krüger, C. (2022). Modeling flexibility in energy systems — comparison of power sector models based on simplified test cases. *Renewable and Sustainable Energy Reviews*, *158*. https://doi.org/10.1016/j.rser.2021.111995

Gøtske, E. K., Andresen, G. B., Neumann, F., & Victoria, M. (2024). Designing a sector-coupled European energy system robust to 60 years of historical weather data. *Nature Communications*, *15*(1), 10680. https://doi.org/10.1038/s41467-024-54853-3

Gøtske, E. K., & Victoria, M. (2021). Future operation of hydropower in Europe under high renewable penetration and climate change. *IScience*, *24*(9). https://doi.org/10.1016/j.isci.2021.102999

Haas, J., Cebulla, F., Cao, K., Nowak, W., Palma-Behnke, R., Rahmann, C., & Mancarella, P. (2017). Challenges and trends of energy storage expansion planning for flexibility provision in low-carbon power systems – a review. In *Renewable and Sustainable Energy Reviews* (Vol. 80, pp. 603–619). Elsevier Ltd. https://doi.org/10.1016/j.rser.2017.05.201

Haas, J., Nowak, W., & Palma-Behnke, R. (2019). Multi-objective planning of energy storage technologies for a fully renewable system: Implications for the main stakeholders in Chile. *Energy Policy*, *126*, 494–506. https://doi.org/https://doi.org/10.1016/j.enpol.2018.11.034

Hall, L. M. H., & Buckley, A. R. (2016). A review of energy systems models in the UK: Prevalent usage and categorisation. *Applied Energy*, *169*, 607–628. https://doi.org/10.1016/j.apenergy.2016.02.044

Hole, J., Philpott, A., & Dowson, O. (2023). *Capacity planning of renewable energy systems using stochastic dual dynamic programming*. https://doi.org/10.1016/j.ejor.2024.12.031

International Energy Agency, I. (2023). *Managing the Seasonal Variability of Electricity Demand and Supply*. www.iea.org

Jaehnert, S., Wolfgang, O., Farahmand, H., Völler, S., & Huertas-Hernando, D. (2013). Transmission expansion planning in Northern Europe in 2030—Methodology and analyses. *Energy Policy*, *61*, 125–139. https://doi.org/10.1016/j.enpol.2013.06.020

Kachirayil, F., Weinand, J. M., Scheller, F., & McKenna, R. (2022). Reviewing local and integrated energy system models: insights into flexibility and robustness challenges. *Applied Energy*, *324*, 119666. https://doi.org/https://doi.org/10.1016/j.apenergy.2022.119666

Kiszka, A., & Wozabal, D. (2025). Stochastic dual dynamic programming for optimal power flow problems under uncertainty. *European Journal of Operational Research*, *321*(3), 814–836. https://doi.org/https://doi.org/10.1016/j.ejor.2024.09.045

Kittel, M., Roth, A., & Schill, W.-P. (2024). *Coping with the Dunkelflaute: Power system implications of variable renewable energy droughts in Europe*. http://arxiv.org/abs/2411.17683

Levin, T., Bistline, J., Sioshansi, R., Cole, W. J., Kwon, J., Burger, S. P., Crabtree, G. W., Jenkins, J. D., O'Neil, R., Korpås, M., Wogrin, S., Hobbs, B. F., Rosner, R., Srinivasan, V., & Botterud, A. (2023). Energy storage solutions to decarbonize electricity through enhanced capacity expansion modelling. *Nature Energy*, *8*(11), 1199–1208. https://doi.org/10.1038/s41560-023-01340-6

Li, X., Song, J., Ma, Y., Zhu, Z., Liu, H., & Wei, C. (2025). Capacity planning for hydro-wind-photovoltaic-storage systems considering high-dimensional uncertainties. *Energy Informatics*, *8*(1). https://doi.org/10.1186/s42162-024-00462-9

Lubin, M., Dowson, O., Garcia, J. D., Huchette, J., Legat, B., & Vielma, J. P. (2023). JuMP 1.0: recent improvements to a modeling language for mathematical optimization. *Mathematical Programming Computation*, *15*(3), 581–589. https://doi.org/10.1007/s12532-023-00239-3

Luss, H. (1982). Operations Research and Capacity Expansion Problems: A Survey. *Operations Research*, *30*(5), 907–947. https://doi.org/10.1287/opre.30.5.907

Martínez-Jaramillo, J. E., van Ackere, A., & Larsen, E. (2023). Long term impacts of climate change on the transition towards renewables in Switzerland. *Energy*, *263*, 126089. https://doi.org/https://doi.org/10.1016/j.energy.2022.126089

Mason, I., Page, S., & Williamson, A. (2013). Security of supply, energy spillage control and peaking options within a 100% renewable electricity system for New Zealand. *Energy Policy*, *60*. https://doi.org/10.1016/j.enpol.2013.05.032

Ministry for the Environment. (2019). *Climate Change Response Zero Carbon Amendment Act 2019*.

Moya, B., Moreno, R., Püschel-Løvengreen, S., Costa, A. M., & Mancarella, P. (2022). Uncertainty representation in investment planning of low-carbon power systems. *Electric Power Systems Research*, *212*. https://doi.org/10.1016/j.epsr.2022.108470

Nohara, D., Kitoh, A., Hosaka, M., & Oki, T. (2006). Impact of Climate Change on River Discharge Projected by Multimodel Ensemble. *Journal of Hydrometeorology*, *7*, 1076–1089. https://doi.org/10.1175/JHM531.1

Osman, A. I., Chen, L., Yang, M., Msigwa, G., Farghali, M., Fawzy, S., Rooney, D. W., & Yap, P.-S. (2023). Cost, environmental impact, and resilience of renewable energy under a changing climate: a review. *Environmental Chemistry Letters*, *21*(2), 741–764. https://doi.org/10.1007/s10311-022-01532-8

Otero, N., Horton, P., Martius, O., Allen, S., Zappa, M., Weschler, T., & Schaefli, B. (2023). Impacts of hot-dry conditions on hydropower production in Switzerland. *Environmental Research Letters*, *18*(6). https://doi.org/10.1088/1748-9326/acd8d7

Pereira, M. V. F., & Pinto, L. M. V. G. (1991). Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, *52*(1), 359–375. https://doi.org/10.1007/BF01582895

Pfenninger, S., Hawkes, A., & Keirstead, J. (2014). Energy systems modeling for twenty-first century energy challenges. In *Renewable and Sustainable Energy Reviews* (Vol. 33, pp. 74–86). Elsevier Ltd. https://doi.org/10.1016/j.rser.2014.02.003

Philpott, A. B., & De Matos, V. L. (2012). Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion. *European Journal of Operational Research*, *218*(2), 470–483. https://doi.org/10.1016/j.ejor.2011.10.056

Purdie, J. (2022). Modelling climate change impacts on inflows, lake storage and spill in snow-fed hydroelectric power catchments, Southern Alps, New Zealand. *Journal of Hydrology (New Zealand)*, *61*(2), 151–178. https://search.informit.org/doi/10.3316/informit.622279673546423

Reddy, S. S., Sandeep, V., & Jung, C.-M. (2017). Review of stochastic optimization methods for smart grid. *Frontiers in Energy*, *11*(2), 197–209. https://doi.org/10.1007/s11708-017-0457-7

Savelsberg, J., Schillinger, M., Schlecht, I., & Weigt, H. (2018). The impact of climate change on Swiss hydropower. *Sustainability (Switzerland)*, *10*(7). https://doi.org/10.3390/su10072541

Schmidt, F. (2025). *On long-duration storage, weather uncertainty and limited foresight*. http://arxiv.org/abs/2505.12538

Schumm, L., Haas, J., Peer, R., & Sterner, M. (2026). The role of hydrogen offtaker regulation in highly renewable electricity systems. *Energy*, *342*, 139513. https://doi.org/https://doi.org/10.1016/j.energy.2025.139513

Soroudi, A. (2013). Robust optimization based self scheduling of hydro-thermal Genco in smart grids. *Energy*, *61*, 262–271. https://doi.org/10.1016/j.energy.2013.09.014

Steidl, S., Zabala Figueroa, A., Peer, R., Wilkening, H., Alhamwi, A., Medjroubi, W., Langnickel, H., Ruf, H., & Haas, J. (2026). High-resolution synthetic residential multi-energy load profiles under electrification scenarios in New Zealand and Germany. *Environmental Research: Energy*, *3*(1), 14501. https://doi.org/10.1088/2753-3751/ae3e90

Stelling, P., Brent, A. C., & Burmester, D. (2026). The impact of 100% renewable electricity on hydropower generation in Aotearoa New Zealand. *Renewable Energy Focus*, *56*, 100769. https://doi.org/https://doi.org/10.1016/j.ref.2025.100769

Sterl, S., Fadly, D., Liersch, S., Koch, H., & Thiery, W. (2021). Linking solar and wind power in eastern Africa with operation of the Grand Ethiopian Renaissance Dam. *Nature Energy*, *6*(4), 407–418. https://doi.org/10.1038/s41560-021-00799-5

Su, C., Li, L., Zhang, T., Sui, Q., & Yang, Y. (2025). A real-time scheduling framework of cascade hydropower-photovoltaic power complementary systems based on model predictive control. *Applied Energy*, *392*, 126023. https://doi.org/https://doi.org/10.1016/j.apenergy.2025.126023

Suomalainen, K., Wen, L., Sheng, M. S., & Sharp, B. (2022). Climate change impact on the cost of decarbonisation in a hydro-based power system. *Energy*, *246*. https://doi.org/10.1016/j.energy.2022.123369

Trutnevyte, E. (2016). Does cost optimization approximate the real-world energy transition? *Energy*, *106*, 182–193. https://doi.org/https://doi.org/10.1016/j.energy.2016.03.038

Van Vliet, M. T. H., Franssen, W. H. P., Yearsley, J. R., Ludwig, F., Haddeland, I., Lettenmaier, D. P., & Kabat, P. (2013). Global river discharge and water temperature under climate change. *Global Environmental Change*, *23*(2), 450–464. https://doi.org/10.1016/j.gloenvcha.2012.11.002

Wessel, J., Kern, J. D., Voisin, N., Oikonomou, K., & Haas, J. (2022). Technology Pathways Could Help Drive the U.S. West Coast Grid's Exposure to Hydrometeorological Uncertainty. *Earth's Future*, *10*(1), e2021EF002187. https://doi.org/https://doi.org/10.1029/2021EF002187

Wetzel, M., Ruiz, E. S. A., Witte, F., Schmugge, J., Sasanpour, S., Yeligeti, M., Miorelli, F., Buschmann, J., Cao, K.-K., Wulff, N., Gardian, H., Rubbert, A., Fuchs, B., Scholz, Y., & Gils, H. C. (2024). REMix: A GAMS-based framework for optimizing energy system models. *Journal of Open Source Software*, *9*(99), 6330. https://doi.org/10.21105/joss.06330

Xie, L., Majumder, S., Huang, T., Zhang, Q., Chang, P., Hill, D. J., & Shahidehpour, M. (2024). The role of electric grid research in addressing climate change. *Nature Climate Change*, *14*(9), 909–915. https://doi.org/10.1038/s41558-024-02092-1

Yalew, S. G., van Vliet, M. T. H., Gernaat, D. E. H. J., Ludwig, F., Miara, A., Park, C., Byers, E., De Cian, E., Piontek, F., Iyer, G., Mouratiadou, I., Glynn, J., Hejazi, M., Dessens, O., Rochedo, P., Pietzcker, R., Schaeffer, R., Fujimori, S., Dasgupta, S., … van Vuuren, D. P. (2020). Impacts of climate change on energy systems in global and regional scenarios. *Nature Energy*, *5*(10), 794–802. https://doi.org/10.1038/s41560-020-0664-z

Yue, X., Pye, S., DeCarolis, J., Li, F. G. N., Rogan, F., & Gallachóir, B. (2018). A review of approaches to uncertainty assessment in energy system optimization models. *Energy Strategy Reviews*, *21*, 204–217. https://doi.org/10.1016/j.esr.2018.06.003

Zhou, Z., Liu, C., & Botterud, A. (2016). *Stochastic Methods Applied to Power System Operations with Renewable Energy: A Review*. https://doi.org/10.2172/1307655