



Co-movements between forward prices and resource availability in hydro-dominated electricity markets

Andreas Kleiven¹ · Simon Risanger¹ · Stein-Erik Fleten¹

Received: 5 December 2022 / Accepted: 7 July 2023
© The Author(s) 2023

Abstract

Hydropower producers estimate the opportunity value of their water, known as a water value, by comparing current prices to future opportunities. When hydropower dominates the energy mix, the system's hydrological state predominantly governs supply and thus prices. Despite this intuitive relationship, industry practice is to assume that inflow to reservoirs and prices are independent when they establish operational policies 1–2 years ahead. To investigate the impact of this assumption, we formulate the hydropower scheduling problem as a Markov decision process and develop a novel price model that considers the joint dynamics of forward prices and inflows. We find that producers underestimate their water value when they ignore co-movements between price and inflow. The dependency makes producers more willing to postpone generation and tolerate slightly higher spillage risk. This is because high inflow periods tend to observe low prices and the reservoir capacity is limited. Nevertheless, a case study of a hydropower plant with industry data suggests modest economic losses in practice. Our numerical results suggest a potential gain of 0.17% in expected revenue and approximately unchanged revenue variance if producers consider the co-movements when establishing an operational policy.

Keywords Commodity pricing · Markov decision process (MDP) · Hydropower scheduling · Stochastic dual dynamic programming (SDDP)

Simon Risanger and Stein-Erik Fleten contributed equally to this work.

✉ Andreas Kleiven
andreas.kleiven@ntnu.no

Simon Risanger
simon.risanger@ntnu.no

Stein-Erik Fleten
stein-erik.fleten@ntnu.no

¹ Norwegian University of Science and Technology, Trondheim, Norway

1 Introduction

Reservoir management is fundamental to ensure efficient allocation of resources in electricity markets dominated by hydropower [22, 44]. Decision-makers want an operational policy that maximizes expected rewards or minimizes expected costs over a given time horizon, considering production and reservoir constraints and uncertainty. As a result, they must solve a sequential decision problem under uncertainty. From a system perspective, the goal is to minimize costs [37], possibly with a risk measure [41]. Individual producers, on the other hand, aim to maximize profits [8, 17, 20, 26]. Inflow to reservoirs is a common source of uncertainty, and in liberalized wholesale markets, prices are also uncertain. A reasonable assumption in such markets is that producers are price-takers and cannot influence the price by their decisions. However, inflows impact prices. According to both economic theory and empirical observations, low inflows to reservoirs correlate with high prices and vice versa in hydro-dominated systems. This is because hydropower has no fuel price and producers instead calculate a water value to determine the expected price to sell energy [38]. When a system has abundant water in its reservoirs, the water value decreases because there is sufficient supply and even a risk of spillage. Conversely, a system with limited water has scarce supply and risk not being able to meet demand. Despite this intuitive relationship, a common industry practice is to assume independent stochastic variables when determining an operational policy. One reason is that forecasting prices and forecasting inflow are functionally distinct activities, performed by different professions. This reason is mirrored by [23], studying a problem involving stochastic wind infeed and electricity prices. For hydropower, it is difficult to identify and estimate underlying co-movements because system prices typically correlate with aggregated system inflows rather than local ones. Still, aggregate inflows are a result of the local inflows at producer locations. Consequently, there is little information about the value and effect of including co-movements between prices and inflow in medium-term hydropower planning.

[30] argue that correlations between accumulated inflow over a whole season and prices are more important than correlations between weekly prices and inflow because of storage capacity, but the authors do not provide any numerical results to support the argument. Common industry practice is to partially ignore correlations when optimizing release decisions. A common process is to first create simultaneous scenarios for price and inflow from a power market model where inflow is the input and price the output [45]. In the second step, typically these simultaneous scenarios are used as input to a more detailed optimization of reservoirs in a smaller geographical area. However, the state-of-the art algorithm used to optimize release decisions consider price and inflow as independent stochastic processes [20]. Indeed, no works properly investigate the impact of ignoring co-movements in hydropower planning. The closest effort is [32], which present a case study of a Norwegian hydropower plant and include a correlation coefficient between the stochastic price and inflow processes. Their findings suggest notable reductions in expected revenues. Nevertheless, we find that their treatment of the

topic is incomplete, and we aim to address the gaps. First, they investigate a case study and report only case-specific findings. In this paper, we present analytical insights on including co-movements. Second, [32] do not explain or outline the intuition behind the results. This is important because we find that their case study contains some unintuitive results, as we discuss later in Sect. 5. By contrast, we explain all findings according to intuition and analytical insights. Third, the modeling of the co-movements by a correlation coefficient estimated using a fixed Pearson correlation coefficient is simplistic. Our work, on the other hand, introduces a sophisticated price process to consider the dependency between price and inflow that vary over the year. Contrary to [32], we find no evidence of a substantial impact on expected revenues.

In this work, we first characterize the impact of co-movements on the policy in a two-stage setting. From there, we develop a stochastic model with co-dynamics between price and inflow to reservoirs, so we can investigate the effect in a multi-stage setting. We formulate a Markov Decision process (MDP) of the medium-term hydropower planning problem and solve it using stochastic dual dynamic programming (SDDP), the state-of-the-art algorithm for seasonal hydropower scheduling [37, 39]. Thus we can analyze numerically the effect of co-movements between exogenous factors on the optimal policy.

Our work contributes to the extensive literature that studies the hydropower scheduling problem [20, 25, 28, 41]. In a broader perspective, we also contribute to research on general insights regarding correlations in stochastic optimization [7, 24]. An example of such work is [1], which uses a distributionally robust approach to analyze losses associated with ignoring correlations. Also related are studies that analyze the effect of correlation between hydro and wind [4, 9, 29, 36, 43] and between electricity prices and wind [34, 36], and between electricity demand and fuel prices [15]. Our specific contribution is to characterize the impact of co-movements on the policy of the hydropower scheduling problem in a two-stage setting.

The joint model for price and inflow is also a specific contribution to models for commodity pricing. We develop an additive three-factor stochastic model where the term structure of electricity prices is partly explained by inflow to reservoirs in the system, partly by a latent factor representing long-term dynamics, and partly by another latent factor for short-term deviations that cannot be explained by inflow to reservoirs. Our model can be interpreted as an additive discrete-time extension of the two-factor model by [40] with two latent factors and one factor that represents some physical phenomena explaining prices. Furthermore, we develop a risk-neutral version that can value hydropower cashflows. Our model allows for a linear and seasonal dependence structure between stochastic variables, and we present a calibration procedure for how to estimate the joint dynamics under a risk-neutral measure when futures contracts only exist for the calibration of a subset of the stochastic factors.

By applying our stochastic model to a hydropower case study, we find that at a given time and state, producers are more willing to postpone production and tolerate a higher probability of spillage if they incorporate co-movements when establishing their production policy. If a high-inflow state realizes in the next stage, a low-price state is more likely, which makes the expected revenue in the high-inflow state less when

incorporating correlations. Therefore, since the reservoir capacities are limited, the producer is more willing to spill the incoming water and keep the reservoir at a higher level. Despite theoretical differences, our result suggests that the impact is modest on practical applications. Using industry data from a Norwegian hydropower plant, we find that the expected revenue decreases between 0.17% and 0.3% for different set-ups, including different parameter estimates of the stochastic model, and different reservoir and production capacities. Nevertheless, this is valuable information for reservoir operators. They can continue their current practices with more evidence backing their assumption of independent price and inflows does not incur large costs.

The rest of the paper is structured as follows. In Sect. 2, we formulate the hydropower scheduling problem as a Markov decision process. We characterize the optimal policy and present a simple example that illustrates the added value of correlations in a two-stage setting. Section 3 briefly discusses the SDDP algorithm, before we present the stochastic model with joint dynamics of prices and inflows to reservoirs in Sect. 4. Section 5 presents a case study, including calibration of the stochastic model, Markov-chain discretization, and numerical results. Concluding remarks are provided in Sect. 6.

2 Hydropower scheduling as a Markov decision process

Hydropower planning is a sequential decision problem under uncertainty. Endogenous states include the amount of water in the reservoirs, and exogenous states include price and inflow. Given the current state and an action of the producer, the state transition is independent of previous actions and states. State transitions therefore satisfy the Markov property, and the problem can be formulated as a Markov decision process (MDP). We denote π as policy, s as endogenous states, ω as exogenous random variables, X_t^π maps the next state when following policy π at time t , and s_t^π is the state reached at time t when following policy π . The problem can be formulated as

$$V_1^\rho(s_0, \omega_1) = \max_{\pi \in \Pi} \mathbb{E} \left[\sum_{t \in \mathcal{T}} \gamma^t r(X_t^\pi(s_t^\pi, \omega_t), \omega_t) \mid \omega_1, s_1 \right], \quad (1)$$

where the expected value at $t = 1$ is the expected reward, r , from following policy $\pi \in \Pi$, where Π is the feasible set of policies. Representation (1) uses discrete time where the planner makes decisions at every $t \in \mathcal{T}$. The set $\mathcal{T} = \{1, \dots, T\}$ contains the stages of the MDP. The endogenous state, s_t , includes the amount of water in the reservoir at time t . The exogenous state vector $\omega_t = (\omega_t^O, \omega_t^C)$ consists of factors that determines immediate rewards, ω_t^O , and components that impact the feasible stage t action set, ω_t^C . Superscripts O and C denote *objective* uncertainty and *constraint* uncertainty, respectively. Superscript ρ denotes that ω_t^C and ω_t^O are correlated. We adopt the notation $V_1^0(s_0, \omega_1)$ and $\alpha_1^0(s_1, \omega_1)$ as the value function and continuation function assuming independent ω_t^C and ω_t^O , respectively. Using Bellman's principle of optimality, we can reformulate problem (1) recursively as a stochastic dynamic program

$$V_t^p(s_{t-1}, \omega_t) = \max_{(x_t, s_t, v_t) \in \mathcal{X}_t(s_{t-1}, \omega_t^C)} r(x_t, \omega_t^O) + \gamma \alpha_{t+1}^p(s_t, \omega_t), \tag{2}$$

$$\alpha_{t+1}^p(s_t, \omega_t) = \mathbb{E}[V_{t+1}(s_t, \omega_{t+1}) \mid \omega_t], \quad t = 1, \dots, T, \tag{3}$$

with $\alpha_{T+1}^p(s_T, \omega_T) = 0$. The set $\mathcal{X}_t(s_{t-1}, \omega_t^C)$ defines the solution space. The continuation, or the expected profit-to-go, function α_{t+1}^p provides the expected value at time t by following the optimal policy from time $t + 1$ to T . We want to investigate the effect of a negative relationship between ω_t^O and ω_t^C , i.e., $\text{Cov}(\omega_t^O, \omega_t^C) < 0$, on the optimal policy of Problem (1). This correlation represents the influence that inflows, ω_t^C , have on prices, ω_t^O .

A producer’s reward is simply the revenue from selling produced energy, x_t , at price ω_t^O :

$$r(x_t, \omega_t^O) = \omega_t^O x_t. \tag{4}$$

Constraints (5a) to (5f) define the decision space. State s_t is the volume of water in the reservoir and (5a) defines the transition from an incoming state s_{t-1} to an outgoing state s_t . It depends on production, x_t , spillage, v_t , and inflow, ω_t^C . There is no penalty for spillage except the opportunity cost of lost sales. Restriction (5b) ensures that the state is within the bounds of the reservoir, where R denotes the maximum capacity. Production is either limited by the reservoir contents, (5c), or the maximum production capacity, (5d), where G denotes maximum production capacity. Finally, (5e) and (5f) ensure that production and spillage cannot be negative, respectively.

$$\mathcal{X}_t(s_{t-1}, \omega_t^C) = \{ \begin{aligned} & s_t + x_t + v_t = s_{t-1} + \omega_t^C \quad (z_1) \end{aligned} \tag{5a}$$

$$-x_t - v_t \leq R - s_{t-1} - \omega_t^C \quad (z_2) \tag{5b}$$

$$x_t \leq s_{t-1} + \omega_t^C \quad (z_3) \tag{5c}$$

$$x_t \leq G \quad (z_4) \tag{5d}$$

$$x_t \geq 0 \quad (z_5) \tag{5e}$$

$$v_t \geq 0 \quad (z_6) \tag{5f}$$

}.

Proposition 2.1 states the properties of the value function V_t^p and continuation function α_t^p . These properties make it possible to characterize the optimal policy for the hydropower scheduling problem defined by (2) to (3) in Proposition 2.2.

Proposition 2.1 For a given $(t, \omega_t^O, \omega_t^C) \in \mathcal{T} \times \mathbb{R} \times \mathbb{R}_+$, the value function $V_t^\rho(\cdot, \omega_t)$ and continuation function $\alpha_{t+1}^\rho(\cdot, \omega_t)$ are concave, and the marginal water value, $\frac{\partial \alpha_{t+1}^\rho(\cdot, \omega_t)}{\partial s_t}$, is decreasing in the endogenous state s_t .

Proof See A.1. □

Proposition 2.2 The optimal generation policy of Problem 2–3 is given by

$$x_t^*(s_{t-1}, \omega_t) = \begin{cases} 0 & \text{if } \omega_t^O < \frac{\partial \alpha_{t+1}^\rho(\cdot, \omega_t)}{\partial s_t} \Big|_{s_{t-1} + \omega_t^C} \\ \min\{G, K, s_{t-1} + \omega_t^C\} & \text{if } \omega_t^O > \frac{\partial \alpha_{t+1}^\rho(\cdot, \omega_t)}{\partial s_t} \Big|_{s_{t-1} + \omega_t^C} \end{cases} \quad (6)$$

where

$$K = s_{t-1} + \omega_t^C - \arg \min_{s_t} \left| \omega_t^O - \frac{\partial \alpha_{t+1}^\rho(\cdot, \omega_t)}{\partial s_t} \Big|_{s_t} \right|$$

Proof See A.2. □

2.1 Covariance in two-stage hydropower scheduling

To study the effect of covariance on the optimal MDP policy, we consider reservoir management in two stages where the producer must determine whether to produce now or later. In other words, the producer faces problem (2) to (3) with $T = 2$. Proposition 2.3 states that a decision-maker underestimates the marginal water value if it ignores the negative correlation between objective uncertainty (price) and constraint uncertainty (inflow).

Proposition 2.3 If $T = 2$, $G = R \in (0, \infty)$, $\text{Cov}(\omega_2^O, \omega_2^C) < 0$, and $s_0 + \omega_1^C < R$, then $\frac{\partial \alpha_2^\rho(\cdot, \omega_1)}{\partial s_1} \Big|_{s_0 + \omega_1^C} > \frac{\partial \alpha_2^0(\cdot, \omega_1)}{\partial s_1} \Big|_{s_0 + \omega_1^C}$.

Proof See A.3. □

Proposition 2.2 combined with Proposition 2.3 implies $x_1^{\rho*} \leq x_1^{0*}$ and $v_2^{\rho*} \geq v_2^{0*}$ in a two-stage setting. A producer will therefore produce the same or less today at optimum if it considers correlations. Hence, correlations provide an incentive to postpone production. If a high inflow realizes, prices tend to be low, and the incoming water is not as valuable compared to an independent stochastic model. This is because the reservoir capacity is limited. Example 2.1 illustrates the phenomenon.

Example 2.1 Added value of including correlations under policy computation

This example demonstrates the impact of ignoring correlations on the optimal production schedule and profit losses. We consider problem (2) to (3) when $T = 2$

and maximum production drains the whole reservoir, $G = R$. Let the process for the stochastic variables be

$$\omega_t^O = \chi_t + \mu^O \tag{7}$$

$$\omega_t^C = v_t + \mu^C \tag{8}$$

$$\chi_t = \phi^O \chi_{t-1} + \sigma^O \epsilon_t^O \tag{9}$$

$$v_t = \phi^C v_{t-1} + \sigma^C \epsilon_t^C \tag{10}$$

where

$$(\epsilon_t^O, \epsilon_t^C) \sim \mathcal{N} \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \right]$$

with numerical values and nomenclature outlined in Table 1.

Figure 1 displays the value of the continuation function and its derivative over different states. We observe that at maximum reservoir, the continuation function must be the same, since any potential new inflow cannot be stored. At low reservoirs, we observe that the continuation value differs. The independent policy sees higher value in the future, since negative correlation is ignored. When there is zero spillage risk, the difference in value is equal to the covariance between price and inflow. This is illustrated in the lower right window in Fig. 1, which shows the value at level 60 where spillage risk is close to zero. As a result, the slope of the curve based on the independent policy must decay at a lower rate. This is illustrated in Fig. 1b, which shows the marginal water value. Hence, the first stage generation decision differs with and without correlation.

Table 1 Numerical values and nomenclature for Example 2.1

Parameter	Value	Unit	
ω_1^O	20	€/MWh	Initial price
ω_1^C	20	MWh	Initial inflow
μ^O	30	€/MWh	Price mean
μ^C	20	MWh	Inflow mean
ϕ^O	0.9		Memory coefficient price
ϕ^C	0.5		Memory coefficient inflow
σ^O	10	€/MWh	Price volatility
σ^C	6	MWh	Inflow volatility
ρ	-0.5		Correlation coefficient
$s_0 + \omega_1^C$	85	MWh	Initial reservoir
R	100	MWh	Maximum reservoir
G	100	MWh	Maximum generation




Fig. 1 Continuation value $\alpha_2^i(\cdot, \omega_1)$ and derivative of the continuation value $\frac{\partial \alpha_2^i(\cdot, \omega_1)}{\partial s_t}$ as a function of reservoir volume s_t , with (black) and without (red) correlation, $i \in \{0, \rho\}$. We observe the continuation value with and without correlations as parallel lines with an offset of $Cov(w^O, w^C)$ when the reservoir is low. The two values get closer as the reservoir level increases. The marginal water value is higher when co-movements are accounted for. This is valid for high reservoir levels, when there is spillage risk. It leads, in principle, to delayed production releases for high reservoir levels (color figure online)

From Proposition 2.2, when $G > s_t + \omega_t^C > K$, the optimal first stage generation decision is

$$K = s_{t-1} + \omega_t^C - \arg \min_{s_t} \left| \omega_t^O - \frac{\partial \alpha_{t+1}^\rho(\cdot, \omega_t)}{\partial s_t} \Big|_{s_t} \right|. \tag{11}$$

Table 2 outlines the optimal solution with zero and -0.5 correlation. The second row shows the spill probability in the last stage at optimum. An optimal policy based on the joint distribution with correlation -0.5 delays some production and tolerates a higher spill probability compared to the policy based on prices and inflow being independent. Table 3 show the expected revenue, where the upper left entry is the optimal policy value of the independent model relative to the optimal value of the joint model with $\rho = -0.5$. An independent policy overestimates the expected revenue by 1.3%. Simulating the policy on the joint model with $\rho = -0.5$ leads to an expected loss of 0.03%, compared to the optimal policy for the joint model, which is to produce less in the first stage.

The state space provides a clearer comprehension of the situations where correlations matter. Figure 2a shows the difference in policy performance as a function

Table 2 Optimal policies for Example 2.1

	$i = 0$	$i = \rho$	
$x_0^{*,i}$	15.0	13.2	First-stage decision
$P(v_1^{*,i} > 0)$	6.8%	9.2%	Spill probability

The first and second columns show the solution under zero and -0.5 correlation, respectively

Table 3 Difference in expected revenue of policies relative to $V_1^0(s_0, \omega_1)$

Policy computation	Policy simulation	
	$\rho = 0$	$\rho = -0.5$
$\rho = 0$	1.3%	- 0.03%
$\rho = -0.5$	-	0 %

The first column displays the correlation coefficient used under policy computation, and the last two column headers display the correlation coefficient used when simulating respective policies. The entry policy computation $\rho = -0.5$ and policy simulation $\rho = 0$ is left open since this case is not relevant

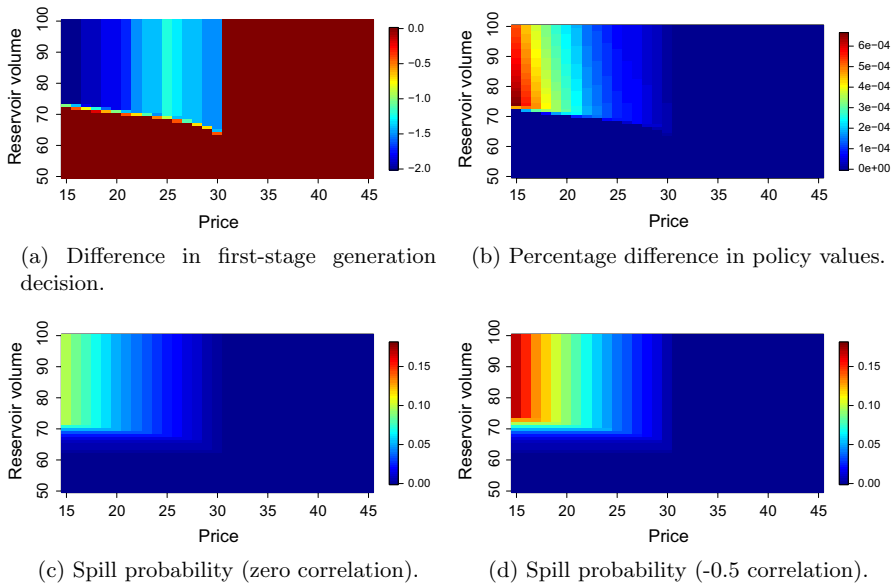


Fig. 2 Comparing policies assuming $\rho = 0$ and $\rho = -0.5$. Upper left is difference in first-stage decision. Upper right is difference in policy values when evaluated on the joint distribution. The lower panel compares spill probabilities

of the current reservoir volume and the current price. The policies are the same for a current price above 30. This is because $\mathbb{E}(\omega_2^2 | \omega_1^0 > 30) < \omega_1^0$, and it is optimal to produce at maximum, which means that correlations are irrelevant. Correlations matter when $\omega_1^0 < 30$ and $P(v_2^{*,i} > 0) > 0$. Here, the current water has a higher value in the future if we consider correlations. It is therefore optimal to produce less today. If the spillage probability is sufficiently small in the second stage, i.e., the reservoir level is low, and $\omega_1^0 < 30$, it is optimal to store water because $\mathbb{E}(\omega_2^0 | \omega_1^0 < 30) > \omega_1^0$ according to the price model. In this case, the first-stage decision is approximately equivalent with and without correlation, as seen in Fig. 2a. We emphasize that the figures are based on simulations and hence there is zero spillage probability in this region. Although, in theory, there will be a spillage probability since the inflow

domain is all real numbers, and the correlated policy will produce slightly less also in this region of the state space.

Figure 2 outlines the policies, where Fig. 2b displays the potential for additional expected revenue from including correlations under policy computation. For example, if the price is around 15 EUR/MWh and the reservoir volume at 75 MWh, the policy with correlation gets around 0.06% higher expected revenue. Figures 2c to 2d compare spillage in the second stage with and without correlation. A policy with correlation has a higher spill probability in the region where the correlated policy delays production.

3 Solution approach

Unlike the two-stage application in Example 2.1, hydropower producers may have reservoirs that can hold water months or years into the future. Medium-term hydropower operations, for instance, plan 1–2 years into the future with time discretized into weeks [20]. This becomes a multistage stochastic problem. The hydropower scheduling problem (2) to (3) suffers from the curses of dimensionality. A particular challenge is that the state and decision variables are continuous and consequently provide infinite combinations. Although we can discretize them, the problem becomes computationally intractable before the granularity becomes representative for an actual hydropower application. Instead, the industry standard is to solve the problem using SDDP, introduced by [37], which approximates the continuation function by a piecewise linear function defined by Benders cuts. The MDP can therefore have continuous states and actions. SDDP creates a valid outer approximation of the expected future value function under certain assumptions [13]:

- A1. The value function $V_t(\cdot, \omega_t)$ is concave (for maximization) with respect to s_t for all stages $t \in \mathcal{T}$ for a fixed ω_t .
- A2. The uncertainty sample space Ω_t is finite for all stages $t \in \mathcal{T}$.
- A3. The feasible region, $\mathcal{X}_t(\cdot, \omega_t)$, is non-empty and an optimal solution exists for every obtainable incoming state variable s_{t-1} and sample ω_t for all stages $t \in \mathcal{T}$.
- A4. Realizations of uncertainty are independent of other stages, i.e., $\mathbb{P}(\omega_t \in \Omega_t | \omega_{t-1}) = \mathbb{P}(\omega_t \in \Omega_t)$, for all stages $t \in \mathcal{T}$. This is commonly referred to as stage-wise independence in the SDDP literature.

Proposition 2.1 proves that the value function $V_t(\cdot, \omega_t)$ is concave, and hence it satisfies A1. We ensure that our scenario generation procedure makes a finite sample space for all stages and hence satisfies A2. Moreover, we select parameters so the feasible region satisfies assumption A3. Our application is a liberalized electricity market where prices are stage-wise dependent random variables, which violate assumption A4. SDDP can circumvent this issue by either including a state variable that follows an autoregressive time series (TS-SDDP) or discretize the random data process to a Markov chain (MC-SDDP) [25]. Because the objective multiplies price with a decision variable, making price a state variable, as required by TS-SDDP, introduces a bilinear term in

the objective that violates assumption A1. A hybrid approach, which acts as a wrapper around the SDDP algorithm, can ensure a concave value function [20]. There are also extensions of the algorithm that can handle the bilinear term in the objective where the value function is a saddle function [11]. The main drawback of MC-SDDP is that a fixed uncertainty partition is required, which may lead to bad performance on some scenarios. However, the MC-SDDP implementation is less complex, it converges faster, and it has been shown to provide policies that perform well when applied to hydro-power planning [25]. Therefore, we argue that MC-SDDP is suitable for analyzing the effect of correlation between price and inflow on hydropower operations policies.

Because the continuation function is concave, SDDP replaces it with the variable θ_t , defined by the hyperplanes or cuts

$$\theta_t \geq \alpha_t^k + \beta_t^k s_t. \tag{12}$$

Parameters α_t^k and β_t^k are the intercept and slope of the hyperplanes, where k indexes iterations. SDDP solves (2) to (3) multiple times in an iterative fashion, where each iteration generates a cut, i.e., parameters α_t^k and β_t^k . Let K denote the current iteration. As K increases, the outer approximation of the continuation function receives more hyperplanes which improves the approximation of the concave continuation function and has convergence guarantees [19, 39]. Problem (2) to (3) is a subproblem that exists for each stage $t \in \mathcal{T}$. The iterative scheme to solve the multistage stochastic problem by SDDP includes a forward simulation and a backward pass. The former starts at $t = 1$ with a pre-determined initial state variable. It samples ω_t , solves (2) to (3), and saves the resulting state variable, \hat{s}_t . The algorithm proceeds to the next stage, $t + 1$, samples ω_{t+1} , and solves (2) to (3) again. The process continues until the terminal stage, T , where the algorithm has a list of simulated values of \hat{s}_t for all $t \in \mathcal{T}$. The backward pass starts at $t = T - 1$ and moves backwards through the list. At stage t we solve (2) to (3) for all ω_{t+1} with fixed \hat{s}_t and obtain objective value $\hat{V}_{t+1, \omega_{t+1}}$ and subgradient $\hat{\lambda}_{t+1, \omega_{t+1}}$ of the objective function with respect to s_t . Cut parameters at iteration K is then calculated according to (13) and (14). These cuts are included in the next forward simulation at iteration $K + 1$ and provide a more detailed approximation of the continuation function.

$$\beta_i^{K+1} = \mathbb{E}[\hat{\lambda}_{t+1, \omega_{t+1}}^K] \tag{13}$$

$$\alpha_i^{K+1} = \mathbb{E}[\hat{V}_{t+1, \omega_{t+1}}^K] - \beta_i^{K+1} \hat{s}_t \tag{14}$$

Note that we use average cuts, which explains the expectation operator in (13) and (14). Multi-cut options, which include all the cuts, i.e., removes the expectation operator in (13) and (14), are also possible but increase the subproblem's size. The SDDP algorithm terminates when new cuts add negligible reductions to the estimated upper bound, reaches an iteration limit, or satisfies a statistical bounds criterion. When SDDP terminates, it has approximated the continuation function for each stage $t \in \mathcal{T}$. The policy is simply the arguments that maximize (2) at stage t with the continuation function approximation. Because the contribution of our work is in the

price process we use, and not the SDDP method itself, we refer to for instance [37, 13], or [25] for more information on SDDP and pseudocode.

4 Stochastic modelling

This section presents a model of spot prices with connection to the hydrological system state and local inflow. Local inflow often possesses strong seasonal variations. Therefore, similar to [20], we normalize inflow to reservoirs and impose an autoregressive model of order one, AR-1, to the seasonality-adjusted inflow series. This model is often used by practitioners. However, inflow generally has positive asymmetry, which motivates the use of skewed distributions, e.g. the lognormal distribution [42]. A comparison of the model fit for our Gaussian inflow model and its lognormal counterpart is provided in Appendix B. Motivated by the empirical fit, we specify a Gaussian model for inflow, but emphasize that our joint model allows for general inflow models. We let $\bar{\mu}_t$ denote the historical average of local inflow and $\bar{\sigma}_t$ the historical standard deviation of local inflow at week t , while v_t denotes the residuals of normalized inflows. Local inflow is still ω_t^C , and it is given by

$$\omega_t^C = \bar{\mu}_t + \bar{\sigma}_t v_t \quad (15)$$

$$v_t = \phi_9 v_{t-1} + \sigma_4 \epsilon_{4,t}, \quad (16)$$

where ϕ_9 is a measure for inflow deviation autocorrelation, and $\epsilon_{4,t}$ is normally distributed with zero mean and unit variance. The local hydrological state, h_t^{loc} , indicates the degree to which the plant has more or less resources available than normal. This captures whether the plant is located in an area with drought or excess of water recently. We define the local hydrological state as the exponentially smoothed inflow deviations from the mean:

$$h_t^{loc} = \phi_8 h_{t-1}^{loc} + (1 - \phi_8) \bar{\sigma}_t v_t, \quad (17)$$

where ϕ_8 is the smoothing coefficient. The lower ϕ_8 the higher weight is given to recent inflow deviations. A dry period for some time locally often indicates a dry period in the system, and vice versa. We propose a linear regression model with an AR-1 error structure for the system hydrological state, h_t^{sys} . The relevant local hydrological state at the producer's location, h_t^{loc} , is an explanatory variable for the hydrological state in a larger geographical area, h_t^{sys} . Our motivation for this model is the fact that the system hydrological state is an aggregation of local hydrological states. An unobservable factor with AR-1 structure, η_t , captures the rest of the system hydrological state:

$$h_t^{sys} = \phi_6 h_t^{loc} + \eta_t \quad (18)$$

$$\eta_t = \phi_7 \eta_{t-1} + \sigma_3 \epsilon_{3,t}. \quad (19)$$

The parameter ϕ_6 measures the effect of local hydrological state on the system state, ϕ_7 measures the autocorrelation of system state deviations that are not explained by the local hydro system that is being modelled, and $\epsilon_{3,t}$ is normally distributed with zero mean and unit variance. The system hydrological state enters the spot price model as an independent factor. In other respects, the price model is equivalent to an additive and discrete time version of the two-factor model in [40]:

$$\omega_t^O = \phi_1 \cos\left((t + \phi_2)\frac{2\pi}{52}\right) + \phi_3 h_t^{sys} + \chi_t + \xi_t \tag{20}$$

$$\chi_t = \phi_4 \chi_{t-1} + \sigma_1 \epsilon_{1,t} \tag{21}$$

$$\xi_t = \beta + \xi_{t-1} + \sigma_2 \epsilon_{2,t}. \tag{22}$$

The first term captures price seasonality. The parameter ϕ_1 determines the amplitude of seasonal variations, and ϕ_2 determines when the peak is. The next term is the hydrological system state, h_t^{sys} , which captures the effect of resource availability on price, determined by the parameter ϕ_3 . Note that the hydrological system state is again explained by another factor, namely, local inflow deviation from the local mean, and influences h_t^{sys} through (15)–(17). The equilibrium price level, ξ_t , captures long-term price behavior, and evolves as a discrete-time arithmetic Brownian motion (ABM) with drift β . Short-term deviations that are not explained by the hydrological state, χ_t , evolves as an AR-1 process, where ϕ_4 measures the autocorrelation in short-term price deviations that are not explained by inflow deviations. The parameters σ_1 and σ_2 is the standard deviation of short-term price deviations and the equilibrium price, respectively, and ϵ_1 and ϵ_2 are normally distributed with zero mean and unit variance. Overall, our model consists of four stochastic factors ($v_t, \eta_t, \chi_t, \xi_t$). Note that local hydrology h_t^{loc} is derived deterministically from v_t , and system hydrology h_t^{sys} is derived deterministically from h_t^{loc} and η_t .

4.1 Risk-neutral process

We develop a risk-neutral version of the stochastic model in the previous section. Futures contracts can then be used to calibrate the model. Given a risk-neutral price process, cashflows can be discounted at the risk-free rate [14]. The stochastic model in (15)–(22) are additive Gaussian. The transformation to an equivalent martingale measure for this type of underlying processes can be done by making assumptions regarding the risk premium [27] or convenience yield [2]. We follow the latter and obtain the following risk-neutral version:

$$\omega_t^O = \phi_1 \cos\left((t + \phi_2)\frac{2\pi}{52}\right) + \phi_3 h_t^{sys} + \chi_t + \xi_t \tag{23}$$

$$\omega_t^C = \bar{\mu}_t + \bar{\sigma}_t v_t \tag{24}$$

$$\chi_t = \phi_4^* \chi_{t-1} + \sigma_1 \epsilon_{1,t}^* \quad (25)$$

$$\xi_t = \phi_5^* \xi_{t-1} + \sigma_2 \epsilon_{2,t}^* \quad (26)$$

$$h_t^{sys} = \phi_6 h_t^{loc} + \eta_t \quad (27)$$

$$\eta_t = \phi_7^* \eta_{t-1} + \sigma_3 \epsilon_{3,t}^* \quad (28)$$

$$h_t^{loc} = \phi_8 h_{t-1}^{loc} + (1 - \phi_8)(\omega_t^C - \bar{\mu}_t) \quad (29)$$

$$v_t = \phi_9^* v_{t-1} + \sigma_4 \epsilon_{4,t}^* \quad (30)$$

$$\epsilon_{i,t}^* \sim \mathcal{N}(0, \Sigma) \quad (31)$$

Here, $\epsilon_{i,t}^*$ has mean zero under the equivalent martingale measure. See Appendix C for derivations. In the next section, we present how to calibrate this model using forward and hydrological data.

5 Medium-term hydropower scheduling: a case study

In this section, we present a case study of a hydropower plant located in the western part of Norway. Its operator provided inflow data and plant characteristics. [31] provides system price data, while we retrieved data for system resource availability, in terms of reservoir content, from [33]. NVE represents the resource availability data series as a deviation from an estimated normal situation at any time of the year. We first present a description of the interconnected system and the aggregation into an energy equivalent reservoir (EER). Then, we present the calibration procedure. Finally, we present the numerical results.

5.1 Energy equivalent reservoir and head variations

The watercourse we study consists of multiple interconnected reservoirs. We aggregate the reservoirs into one energy equivalent reservoir (EER). The composite representation of the interconnected multireservoir system was originally proposed by [3], and has been applied in practice [16, 28]. Note that SDDP efficiently handles the situation of multiple interconnected reservoirs, but since our focus is to study the effect of covariance on the optimal MDP policy, we instead study an EER for comparison and consistency between our analytical two-stage results and numerical multi-stage results. The aggregation simplifies the modeling of uncertainties and allows us to study general insights of the effect of covariance on the optimal MDP policy. For the same reasons, we assume a constant reservoir-specific volume-energy

transformation coefficient. In theory, this transformation depends on the discharge x_t , where the amount of discharge influences the energy efficiency. Piece-wise linear efficiency curves can approximate this behavior. Head variations, which depend on s_t , also influence the volume-energy transformation but require non-convex subproblems.

When aggregating reservoirs, the problem size is small enough to be solved with a standard stochastic dynamic programming approach (SDP). However, since SDDP is state-of-the art in hydropower planning, and the model easily can be extended to a multi-reservoir setting, we demonstrate our model using SDDP. The multi-reservoir setting could potentially lead to computational difficulties, as several inflow processes must be discretized. A way to avoid this is to reduce the number of inflow factors, which often is reasonable in a limited geographical area where inflow often is highly correlated.



5.2 Calibration of the stochastic model

We use electricity futures data with time to maturity 1–5 years from 2011 to 2018, local inflow data from 2009 to 2018, and data for the system hydrological state from the same period when calibrating our model. Pricing models are commonly calibrated using Kalman filtering and maximum likelihood estimation, where observed futures prices are assessed against model predictions. Joint calibration is unattainable because future contracts for inflow do not exist. We therefore apply a two-step process. First, we employ maximum likelihood estimation with historical data for inflow and system hydrological state. We then use (18) and (19) to predict future hydrological states and extract the effect on futures prices. Afterwards we apply standard Kalman filtering and maximum likelihood estimation. We must calibrate both the joint model and a benchmark model where ω_t^O and h_t^{sys} are independent. D outlines how we calibrate the benchmark model, while the following six steps explain our estimation approach for the joint model in detail.


Step 1: Computing smooth electricity forward curves and de-seasonalized prices. We decompose the forward curve into a seasonal component and an adjustment function that accounts for deviations from seasonality. The seasonal component has the same form as in the pricing model from the previous section, while the adjustment function is defined based on arbitrage arguments and the maximum smoothness criterion. See [5, 10, 18] for procedures for computing smooth synthetic forward curves. We estimate coefficients of the seasonal component by regression using front month forward contracts. Figure 3 illustrates examples of smooth synthetic forward curves computed using this procedure. Figure 4a displays synthetic forward curves for the entire electricity forward contract data set. After having estimated the seasonality in step 1, we de-seasonalize forward prices and continue the estimation procedure with forward prices adjusted for seasonality. Figure 4b outlines the de-seasonalized forward curves.

Step 2: Calibration of the local inflow model. We first normalize weekly inflow data based on weekly mean and standard deviation, similarly as in Gjelsvik et al. (2010) [20], before we fit an AR-1 process to account for serial correlations. Figure 5a displays

Fig. 3 Examples of smoothed futures curves by applying the method by Benth et al. [5]. Futures prices are obtained from [31]. The upper was observed on 3 January 2011 and the lower was observed on 27 May 2015




(a) Historical futures curves.




(b) Seasonality-adjusted futures curves.

Fig. 4 Smoothed futures curves for the entire data set



(a) Local inflow from an area in the Western part of Norway. The solid line is the weekly mean $\bar{\mu}_t$ and the dashed line is one standard deviation $\bar{\sigma}_t$.



(b) Smoothed local inflow (dashed line) and hydrological system state (solid line). NVE (2022) provides data for the system state.

Fig. 5 Hydrological data from 2009 to 2018

the local inflow. We assume that inflow risk is unrelated to the world economy and therefore diversifiable, commanding no risk premium.

Step 3: Constructing a time series of local hydrology. We use the exponentially weighted inflow series as the hydrological state

$$\begin{aligned}
 h_t^{loc} &= \phi_8(\omega_t^C - \bar{\mu}_t) + (1 - \phi_8)h_{t-1}^{loc} \\
 &= \phi_8(\omega_t^C - \bar{\mu}_t) + \phi_8(1 - \phi_8)(\omega_{t-1}^C - \bar{\mu}_{t-1}) + (1 - \phi_8)^2 h_{t-2}^{loc} \\
 &= \phi_8(\omega_t^C - \bar{\mu}_t) + \phi_8(1 - \phi_8)(\omega_{t-1}^C - \bar{\mu}_{t-1}) \\
 &\quad + \phi_8(1 - \phi_8)^2(\omega_{t-2}^C - \bar{\mu}_{t-2}) + (1 - \phi_8)^3 h_{t-3}^{loc} \\
 &\quad \vdots \\
 &= \phi_8 \sum_{k=0}^K (1 - \phi_8)^k (\omega_{t-k}^C - \bar{\mu}_{t-k}) + (1 - \phi_8)^K h_{t-K}^{loc},
 \end{aligned}
 \tag{32}$$


where K denotes the number of time periods. This time series captures the situation of a local producer having either less or more resources available than normal, depending on how inflow has deviated from its mean for a longer period. This may partly explain water supply in the system. Figure 5b displays the time series for the local hydrology in red together with system hydrology in black.

Step 4: Calibration of the local and system hydrology model. The system hydrology model is a linear regression model with an AR-1 error structure. It uses local hydrology as an explanatory variable, which again is driven by local inflow. We consider parameters of the inflow model fixed, and jointly estimate parameters of the local and system hydrology model. We enforce $\phi_8 = \phi_7$, which ensures that the local hydrology mean reverts at the same speed as the AR-1 process. We consider this reasonable because system hydrology is an aggregation all reservoirs in the system. This restriction is necessary to get reasonable hydrology predictions in the next step. When estimating the parameters, we iterate over ϕ_6 and ϕ_8 , or equivalently ϕ_7 , and study the likelihood surface of the linear regression model with AR-1 structure in (18) and (19). Figure 6 displays the likelihood surface of $\phi_8 = \phi_7$ and ϕ_6 with a mark at optimum.

Step 5: Estimating the effect of hydrology on prices and using the inflow-hydrology model to predict and extract the impact on futures prices. Given the model parameters $\theta = (\phi_1, \phi_2, \phi_3, \phi_4^*, \phi_5^*, \phi_6, \phi_7^*, \phi_8, \phi_9^*)$, futures prices at any given time t with maturity τ can be expressed by our model as follows:

$$\begin{aligned}
 \hat{F}_{t,\tau}(\chi_t, \xi_t, h_t^{loc}, \eta_t, \nu_t; \theta) &= \mathbb{E}^*(\omega_\tau^O) \\
 &= \phi_1 \cos\left((t + \phi_2)\frac{2\pi}{52}\right) + \phi_3 \hat{h}_{t,\tau}^{sys} + \phi_4^* \tau^{-t} \chi_t + \phi_5^* \tau^{-t} \xi_t
 \end{aligned}
 \tag{34}$$

Fig. 6 Likelihood surface for estimation of $\phi_8 = \phi_7$ and ϕ_6



where

$$\hat{h}_{t,\tau}^{sys} = \phi_6(\phi_8^{\tau-t} + h_t^{loc} + \phi_9^* \tau^{-t}(1 - \phi_8)\bar{\sigma}_\tau v_t) + \phi_7^* \tau^{-t} \eta_t. \tag{35}$$

We estimate ϕ_3 by regressing front month futures contracts on h_t^{sys} . Afterwards, we adjust the effect of hydrology by using model predictions and create a new data set of smooth synthetic forward curves from step 1, adjusted for hydrology. Figure 7a displays the hydrology predictions and Fig. 7b shows hydrology adjusted futures contracts, $F_{t,\tau} - \phi_3 \hat{h}_{t,\tau}^{sys}$. By comparing Figs. 4b and 7b, we observe less variability in futures contracts after accounting for hydrology. Forward curves adjusted for hydrology are then used in the next step.

Step 6: Calibrating the long- and short-term factors to futures residuals. In the final step, we run the Kalman filter together with maximum likelihood under the restriction of zero correlation between h_t^{sys} and χ_t . This is motivated by the fact that the price component that explains hydrology cannot be correlated by the price component explaining everything but hydrology. See [40] and [21] for more information on the calibration of geometric models of latent factors χ_t and ξ_t . We modify the likelihood function to calibrate our additive model. Appendix E provides nomenclature and parameter estimates.

5.3 Numerical results

We implement the Markov chain SDDP in Julia [6] using the SDDP package [12] and the Gurobi solver. We compute policies on the dependent and independent Markov chain based on data from a Norwegian hydropower producer. The procedure for estimating the parameters of the stochastic process is explained in Sect. 5.2, and parameter estimates are provided in Table 5 (Appendix E). When discretizing the stochastic processes we focus on preserving the mean and variance of prices and inflow, and the cross-moment. The discretization procedure is described in Appendix F. We use a 2% annual discount rate, reflecting the risk-free rate of return. Revenues are from 2 years of operations to ensure both policies have the same amount of inflow to allocate. However, policies in time periods close to the terminal stage

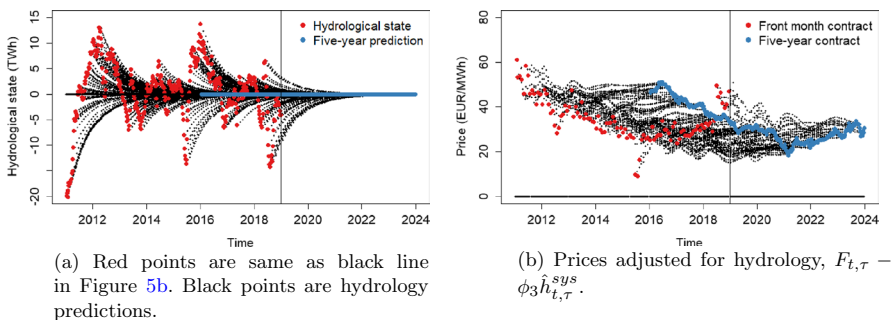


Fig. 7 Hydrology prediction and prices adjusted for hydrology

Fig. 8 Mean and 10th to 90th percentiles of reservoir trajectories for independent and dependent Markov chains

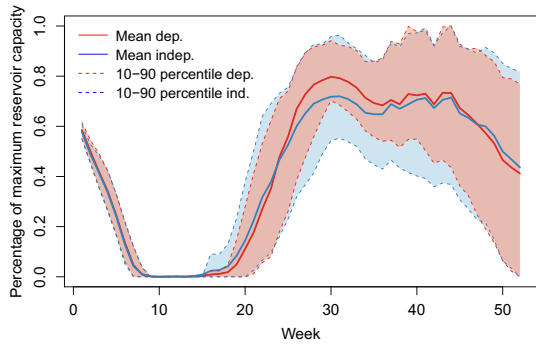
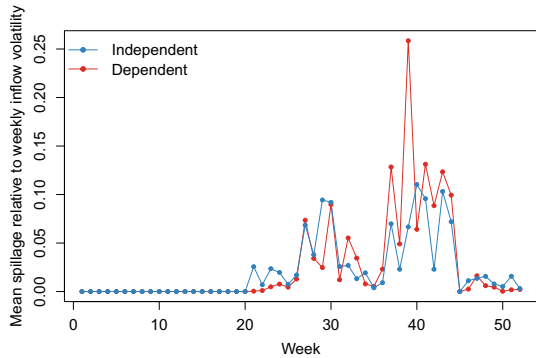


Fig. 9 Mean of spillage for independent and dependent Markov chains relative to weekly inflow volatility



experience end of horizon effects because they want to efficiently use all resources and thus drain the reservoir. This does not happen in practice, so the policy insights we present are only from the first year of operation where end-of-horizon effects have negligible impact.

Figure 8 displays the reservoir trajectories under optimal policies. While both policies drain their reservoirs to capitalize on high winter prices, a period with traditionally high demand and prices, their decisions start to diverge in the late spring. We observe that the reservoir mean is slightly higher during the autumn (from around week 25 to 40) for the policy computed on the dependent Markov chain. Prices usually increase in this period because we move toward higher winter prices. If prices correlate with inflows to the reservoir, the producer favors a slightly higher mean reservoir if it accounts for a negative correlation. This is because if a high-inflow state occurs in the next stage, the price tends to be low. From a current viewpoint, since reservoirs are limited, this makes current inflow less valuable than if prices and inflows were uncorrelated. As a result, the producer tolerates a higher amount of spillage in this period, as shown in Fig. 9. Figure 8 also shows that considering correlations reduces the reservoir trajectory variability. We observe from the 10th percentile that a policy based on the dependent Markov chain has fewer low-reservoir trajectories during the autumn, compared to the independent Markov chain. This is the period where the covariance is highest according to our estimated model parameters and where prices are expected to increase. The policies differ less

Fig. 10 Mean and 10th to 90th percentiles of generation decisions for independent and dependent Markov chains

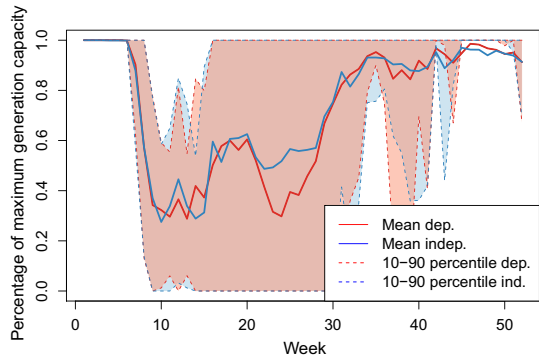


Table 4 Difference in expected revenue of policies relative to $V_1^p(s_0, \omega_1)$

Policy computation	Policy simulation	
	Independent MC	Dependent MC
Independent MC	2.55%	- 0.17%
Dependent MC	-	0 %

The first column tells which Markov chain that is used under policy computation, and the last two column headers tell which Markov chain that is used when simulating respective policies

in the other parts of the year. Figure 10, which displays the production, reinforces the observations of the reservoir trajectories. We observe that the producer generates less on average until around week 30.

Table 4 outlines the expected revenues for 2 years. The first column states which Markov chain is used when computing the policy, and the two rightmost columns state which Markov chain the policy is evaluated on. The numbers are percentage differences relative to the value of the optimal policy of the dependent Markov chain, $V_1^p(s_0, \omega_1)$. We observe that the optimal value of the independent Markov chain policy overestimates the value by 2.55%. However, when simulating the policy from the independent Markov chain, we find approximately 0.17% loss incurred by ignoring correlations under policy computation. The policy computed on the independent Markov chain has a slight disadvantage because the simulations are from the dependent Markov chain, which may have a slight discrepancy in the values of the independent Markov chain’s nodes despite similar first and second moments. Therefore, our results indicate that although there is a strong relationship between inflows to reservoirs and prices, the producer does not necessarily gain much by incorporating co-movements when establishing the seasonal operational policy.

Figures 11a, b show scenarios for the 100 lower and 100 upper revenues, respectively, out of 1000 scenarios. The black line is optimized revenue for the dependent Markov chain, and the red dashed line is revenue that the policy based on the independence assumption attains on the dependent Markov chain. The black solid line in Fig. 11a, b lies slightly above the dashed red line, which illustrates that the producer can slightly increase revenue in both low and high scenarios by accounting for

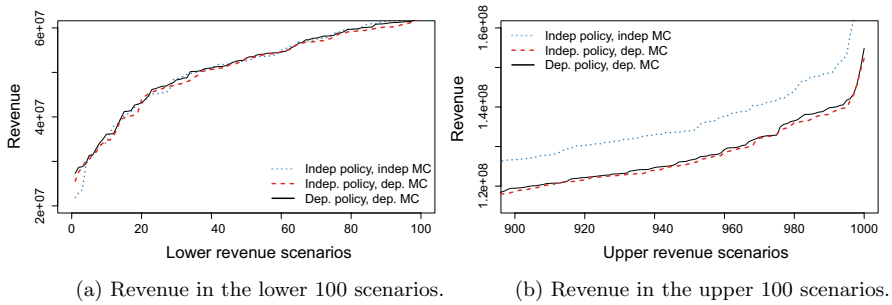



Fig. 11 Revenues in upper and lower 100 scenarios out of 1000

co-movements between prices and inflows when establishing the operational policy. Moreover, we observe from Fig. 11a, b that the blue dotted line, which is optimized revenue for the independent Markov chain, mostly lies below the black line in the lower 100 scenarios and above the black line in the upper 100 scenarios. An explanation for this is that the negative correlation between prices and inflow creates a natural hedge in the lower 100 scenarios. Revenues are therefore underestimated if correlations are ignored. In the upper 100 scenarios, however, revenues are overestimated because both high price and high inflow realizations are less likely to happen if co-movements are accounted for.

Figure 12 displays a scatter plot of the revenues from 1000 simulations obtained when the policy is optimized based on the independence assumption, but evaluated on the dependent Markov chain. The black line illustrates the one-to-one mapping of optimized revenue for the dependent Markov chain. We observe that scenarios with low revenue obtained by the independent policy correspond to low revenue scenarios when following the dependent policy. This further accentuates that the producer cannot gain much by accounting for the negative relationship between prices and inflows to reservoirs when establishing the operational policy.

To test how sensitive our numerical results are to parameter estimates of the stochastic process and plant characteristics, we conduct two additional numerical experiments. The analytical results from Sect. 2.1 find that a stronger degree of co-movements between prices and inflows to reservoirs can increase the opportunity value from accounting for these co-movements when computing the policy.

Fig. 12 Scatter plot of revenues. Red points show revenues when assuming zero co-movements when both optimizing and evaluating the policy. Red points show revenues when assuming zero co-movements when optimizing and evaluated on the dependent Markov chain (color figure online)



Moreover, a smaller reservoir capacity relative to inflow measures and more production capacity could also enhance revenue, as the opportunity value stems from spillage risk. Results from such case studies are provided in Appendix G. A summary of the results is that reservoir trajectories and policy values are highly case-specific, as expected. In line with intuition, both experiments find that the expected revenue decreases. Nonetheless, the reductions are a modest 0.24% for stronger degree of co-movements and 0.3% for smaller reservoir capacity and increased production capacity. The revenue variances have a similar behavior as the base case. The moderate changes reinforce the evidence for just modest relative gains from considering co-movements in price and inflow when establishing an operational policy.

Our findings diverge from [32], who report 2.5% and 3.1% reduced expected revenue on a case study of another Norwegian hydropower plant. We apply our model to different case studies, and we have different assumptions regarding the evolution of uncertainties and co-movements, which may explain the differences in results. We also note that there is no clear relationship between reservoir trajectories and correlation coefficient in their results. This is different from our findings, which indicate that the producer tolerates slightly more spillage if correlations are taken into account and does therefore prefer a slightly higher average reservoir trajectory.

6 Conclusions

Hydro-dominated systems in liberalized electricity markets have a clear relationship between prices and inflows to the system's reservoirs. Full reservoirs mean abundant supply and even risk of spillage, which depress prices. Low reservoirs, on the other hand, may indicate a drought and thus high prices. Nevertheless, common industry practice is to consider price and inflow as independent stochastic processes when optimizing release decisions. This paper investigate the impact of this assumption.

In order to study the co-movements between electricity prices and inflows to reservoirs in hydro-dominated electricity markets, we develop a novel stochastic model that captures the relationship between system prices and supply. Hydrological states at the producer's location and in the system as a whole represent the latter. The local state indicates the resource availability of the producer. This state partly explains the supply in a larger geographical area, which again partly explains prices. We provide general insights about the effect of such co-movements on the optimal policy of a production scheduling problem in a two-stage setting. We apply the stochastic model in a multistage setting by studying the medium-term hydropower scheduling problem on industry data. Policies that consider co-movements prefer a slightly higher average reservoir trajectory during periods of high correlations. Moreover, the reservoir trajectory variance reduces, and the producer tolerates a slightly higher spillage probability when prices are expected to increase in the near future. From case studies using industry data, we find that a policy that ignores correlations only incurs 0.17% to 0.30% loss in expected revenues for different reservoir and production capacities and inflow behavior. This is valuable information for reservoir operators, who receive

evidence indicating that the theoretical differences only amount to modest economic losses in practice.

Proof of Propositions

This section presents the proofs of propositions.

Proof of Proposition 2.1

Proof The proof requires finiteness. No production results in zero value, hence $V_t(s_{t-1}, \omega_t) \geq 0$, which implies $\alpha_{t+1}(s_t, \omega_t) \geq 0$. An upper bound is provided by maximum production at every stage:

$$\alpha_{t+1}(s_t, \omega_t) = \delta \mathbb{E}(V_{t+1}(s_t, \omega_{t+1}) | \omega_t) \tag{36}$$

$$\leq \delta G \left(\sum_{i=t+1}^T \mathbb{E}(\omega_i^O | \omega_i^O) \right) \tag{37}$$

Furthermore,

$$V_t \leq \omega_t^O G + \alpha_{t+1}(s_t, \omega_t). \tag{38}$$

Hence, the value function and continuation function are finite. We proceed by showing that the terminal value is concave in s_t for a fixed ω_t at stage T . For a given ω_T , we have

$$V_T(s_{T-1}, \omega_T) = \max_{x_T \in \mathcal{X}(s_{T-1}, \omega_T^C)} \omega_T^O x_T \tag{39}$$

$$\alpha_{T+1}(s_T, \omega_T) = 0 \tag{40}$$

This is a linear program where $R - s_{T-1} - \omega_T^C$ is the upper bound of the feasible set $\mathcal{X}(s_{T-1}, \omega_T^C)$. From standard linear programming results, $V_T(s_{T-1}, \omega_T)$ is piecewise linear concave in s_{T-1} . The continuation function is zero and is therefore also piecewise linear concave.

By finiteness of the continuation function and the induction hypothesis, it is easy to verify that $\alpha_{t+1}(\cdot, \omega_t)$ is piecewise linear concave in s_t . Hence, for a feasible stage t action set $\mathcal{X}(s_{t-1}, \omega_t^C)$ which is bounded by $R - s_t - \omega_t^C$, a linear reward function $\omega_t^O x_t$, and a concave continuation function, it follows that the value function $V_t(s_{t-1}, \omega_t)$ is piecewise linear concave. This implies that $\frac{\partial \alpha_{t+1}(\cdot, \omega_t)}{\partial s_t}$ is decreasing in s_t .

□

Proof of Proposition 2.2

Proof The Lagrangian is given by

$$\begin{aligned}
 L(x_t, v_t, s_t) = & \omega_t^O x_t + \gamma \alpha_{t+1}(s_t, \omega_t) + z_1(s_{t-1} + \omega_t^C - s_t - x_t - v_t) \\
 & + z_2(R - s_{t-1} - \omega_t^C + x_t + v_t) \\
 & + z_3(s_{t-1} + \omega_t^C - x_t) + z_4(G - x_t) + z_5 x_t + z_6 v_t.
 \end{aligned}$$

This yields optimality conditions

$$\begin{aligned}
 \frac{\partial L(x_t, v_t, s_t)}{\partial x_t} &= \omega_t^O - z_1 + z_2 - z_3 - z_4 + z_5 = 0 \\
 \frac{\partial L(x_t, v_t, s_t)}{\partial v_t} &= -z_1 + z_2 + z_6 = 0 \\
 \frac{\partial L(x_t, v_t, s_t)}{\partial s_t} &= \frac{\partial \alpha_{t+1}(s_t, \omega_t)}{\partial s_t} - z_1 = 0,
 \end{aligned}$$

where we have set $\gamma = 1$ for ease of notation. Rearranging,

$$\frac{\partial \alpha_{t+1}(s_t, \omega_t)}{\partial s_t} = \omega_t^O + z_2 - z_3 - z_4 + z_5 \tag{41}$$

$$\frac{\partial \alpha_{t+1}(s_t, \omega_t)}{\partial s_t} = z_6 + z_2 \tag{42}$$

In the trivial case, if $\omega_t^O \leq 0$ the optimal decision is $x_t^*(s_t) = 0$. Moreover, the marginal water value is always non-negative, since if the price is negative, the producer can spill the water at zero cost. Therefore, this trivial situation is covered by the upper case in Proposition 2.2. Below, we prove each of the two cases, assuming $\omega_t^O > 0$.

Case 1: $\frac{\partial \alpha_{t+1}(s_t, \omega_t)}{\partial s_t} > \omega_t^O$.

Since $\omega_t^O > 0$, we can in this case safely assume $s_{t-1} + \omega_t^C - x_t < R$, since the marginal water value cannot be positive if the reservoir plus incoming inflow is larger than the reservoir plus generation capacity. From (41) we get

$$z_2 + z_5 > z_3 + z_4. \tag{43}$$

Because $s_{t-1} + \omega_t^C < R$, this constraint is not binding and thus $z_2 = 0$. This means that z_5 must be positive for the condition to hold. Hence, in this case it is optimal to produce nothing, i.e., $x_t^* = 0$. This proves the first case.

Case 2: $\frac{\partial \alpha_{t+1}(s_t, \omega_t)}{\partial s_t} < \omega_t^O$.

From (41) we get

$$z_2 + z_5 < z_3 + z_4 \tag{44}$$

Either z_3 or z_4 must be positive, which means $z_5 = 0$ since both the maximum and minimum generation constraint cannot be binding at the same time. Since either z_3 or z_4 is positive, the optimal generation quantity is $x_t^* = \min\{G, s_t + \omega_t^C\}$ if $\frac{\partial \alpha(s_{t+1}, \omega_t)}{\partial s_t} < \omega_t^O$ for all s_t . However, since $\frac{\partial \alpha_{t+1}(s_t, \omega_t)}{\partial s_t}$ is decreasing in s_t , see Proposition 2.1, it is optimal to produce until the price and derivative of the continuation function, or the marginal water value, are equal, i.e. $\frac{\partial \alpha_{t+1}(s_t, \omega_t)}{\partial s_t} = \omega_t^O$. For an illustration, see Fig. 1b. This means that in a situation where $K = s_{t-1} + \omega_t^C - \arg \min_{s_t} \left| \omega_t^O - \frac{\partial \alpha_{t+1}(s_t, \omega_t)}{\partial s_t} \right| \in (0, \min\{G, s_t + \omega_t^C\})$, the optimal generation quantity is K . Therefore, in total, the optimal generation quantity is $x_t^* = \min\{G, K, s_t + \omega_t^C\}$. This proves the second case. \square

Proof of Proposition 2.3

Proof In stage 2 it is optimal to produce at maximum, and since $G = R$, the optimal production quantity is $x_2^* = \min\{s_1 + \omega_2^C, R\}$. Inserting this into the expression for $\alpha_2(s_1, \omega_1)$ gives

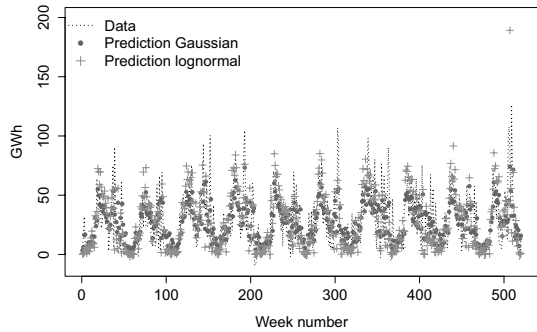
$$\begin{aligned}
 \alpha_2^p &= \delta \mathbb{E}(\omega_2^O \min\{s_1 + \omega_2^C, R\}) \\
 &= \mathbb{E}\left(\omega_2^O(s_1 + \omega_2^C) \mathbb{1}_{\{\omega_2^C \leq R - s_1\}} + R \mathbb{1}_{\{\omega_2^C > R - s_1\}}\right) \\
 &= (\mathbb{E}(\omega_2^O \omega_2^C) + \mathbb{E}(\omega_2^O) s_1) P(\omega_2^C \leq R - s_1) + R \mathbb{E}(\omega_2^O) P(\omega_2^C > R - s_1) \\
 &= (\mathbb{E}(\omega_2^O) \mathbb{E}(\omega_2^C) + \text{Cov}(\omega_2^O, \omega_2^C) + \mathbb{E}(\omega_2^O) s_1) P(\omega_2^C \leq R - s_1) + R \mathbb{E}(\omega_2^O) P(\omega_2^C > R - s_1) \\
 &= \alpha_2^0 + \text{Cov}(\omega_2^O, \omega_2^C) P(\omega_2^C \leq R - s_1)
 \end{aligned}
 \tag{45}$$

The term $\text{Cov}(\omega_2^O, \omega_2^C)$ is negative and $P(\omega_2^C \leq R - s_1)$ is positive and greater than zero if $s_1 < R$, hence $\alpha_2^p < \alpha_2^0$ when $s_1 < R$. Note that $P(\omega_2^C \leq R - s_1)$ is monotonically increasing in s_1 . Hence, $\alpha_2^0 - \alpha_2^p$ is a decreasing function from $\text{Cov}(\omega_2^O, \omega_2^C) P(\omega_2^C \leq R)$ at $s_1 = 0$ to 0 at $s_1 = R$. For an illustration, see Fig. 1a. This means that $\frac{\partial \alpha_2^p(s_1, \omega_1)}{\partial s_1} > \frac{\partial \alpha_2^0(s_1, \omega_1)}{\partial s_1}$. \square

Assessment of the inflow model

See Fig. 13.

Fig. 13 Inflow data, and one-step ahead predictions under the Gaussian model and the lognormal counterpart



Derivation of the risk-neutral process

We first derive the continuous-time risk-neutral version and then derive the discretized version. The general dynamic of prices is

$$dS_t = \alpha(S_t, t)dt + \sigma(S_t, t)dz, \tag{46}$$

and the following deterministic convenience yield process is

$$dC_t = \delta(S_t, t)dt. \tag{47}$$

The discounted gains process of owning a derivative or asset on S_t must be a martingale and is given by

$$dG_t = d(e^{-rt}S_t) + e^{-rt}dC_t \tag{48}$$

$$= -re^{-rt}S_tdt + e^{-rt}dS_t + e^{-rt}dC_t \tag{49}$$

$$= e^{-rt}((-rS_t + \alpha(S_t, t) + \delta(S_t, t))dt + \sigma(S_t, t)dz) \tag{50}$$

$$= e^{-rt}\sigma(S_t, t)(\lambda(S_t, t)dt + dz) \tag{51}$$

$$= e^{-rt}\sigma(S_t, t)dz^*, \tag{52}$$

where r denotes the risk-free rate, and where the risk premium is

$$\lambda(S_t, t) = \frac{\alpha(S_t, t) + \delta(S_t, t) - rS_t}{\sigma(S_t, t)}. \tag{53}$$

By Girsanov's theorem, the new process dz^* is a standard Brownian motion [35]. The risk-neutral dynamics can now be written as

$$dS_t = (rS_t - \delta(S_t, t))dt + \sigma(S_t, t)dz^*, \tag{54}$$

by rearranging (49) and (52).

We now move to our specific stochastic model for prices, and start with the equilibrium price level, ξ_t . If we let $\alpha(\xi_t, t) = \alpha$ and $\sigma(\xi_t, t) = \sigma$, the underlying is an arithmetic Brownian motion (ABM), which can be written in discrete time as in (22). Following Alexander et al. (2012) [2], we impose the long-term convenience yield $\delta(\xi_t, t) = \delta\xi_t$. Note that since ξ_t can be positive and negative, convenience yield is negative if ξ_t is negative. The risk premium is then given by $\sigma_2\lambda_\xi(\xi_t, t) = \alpha - (r - \delta_\xi)\xi_t$, which inserted into (54) leads to the following risk-neutral dynamics:

$$d\xi_t = (r - \delta_\xi)\xi_t dt + \sigma_\xi dz^* \tag{55}$$

This process has geometric expected growth and additive noise. It can be discounted by the risk-free rate for getting the value today of an asset on ξ_t . Rearranging, a discrete-time version risk-neutral process can be written as

$$\xi_t = \phi_5^* \xi_{t-1} + \sigma_2 \epsilon_2^* \tag{56}$$

where $\phi_5^* = e^{(r-\delta_\xi)\Delta t}$ and $\sigma_2^2 = \frac{1}{2(r-\delta_\xi)}\sigma_\xi^2(1 - e^{(r-\delta_\xi)\Delta t})$.

For the short-term deviations, χ_t , we let $\alpha(\chi_t, t) = -\kappa\chi_t$ and $\sigma(\chi_t, t) = \sigma_\chi$. The underlying is then an additive Ornstein-Uhlenbeck process, which can be seen by inserting into (46). This can be written as an AR-1 process in discrete time, as in (21). Imposing the short-term convenience yield $\delta(\chi_t, t) = \delta_\chi\chi_t$ entails short-term risk premium $\sigma_\chi\lambda_\chi(\chi_t, t) = -\kappa_\chi\chi_t - (r - \delta_\chi)\chi_t$. Inserting convenience yield into (54) gives

$$d\chi_t = (r - \delta_\chi)\chi_t dt + \sigma_\chi dz^* \tag{57}$$

Rearranging, a discrete-time version risk-neutral process for the short-term deviations can be written as

$$\chi_t = \phi_4^* \chi_{t-1} + \sigma_1 \epsilon_1^* \tag{58}$$

where $\phi_4^* = e^{(r-\delta_\chi)\Delta t}$ and $\sigma_2^2 = \frac{1}{2(r-\delta_\chi)}\sigma_\chi^2(1 - e^{(r-\delta_\chi)\Delta t})$. For the stochastic processes v_t and η_t in (16) and (19) the derivations are similar to the derivations for χ_t .

Calibration with ω_t^O and h_t^{sys} being independent

For the case study in Sect. 5, we need to benchmark our stochastic model against a model where price is not affected by inflows to reservoirs. This can be obtained by setting the coefficient of h_t^{sys} to zero and let the short-term price factor χ_t capture all price deviations from the long-term price level, including price deviations related to hydrology. We then need to re-estimate parameters of the short-term factor, such that mean-reversion speed and variance are approximately equivalent to the model where hydrology enters as an explicit factor. Therefore, we let

$$\phi_3^l = 0 \tag{59}$$

$$\phi_6^I = 0 \quad (60)$$

$$\phi_4^{*,I} = \frac{1}{2}(\phi_6 + \phi_4^*) \quad (61)$$

$$\sigma_4^{*,I} = \sqrt{\frac{1 + \phi_4^{*,I\ 2}}{1 + \phi_4^{*2}} \sigma_1^2 + \frac{1 + \phi_4^{*,I\ 2}}{1 + \phi_6^{*2}} \sigma_3^2}, \quad (62)$$

where σ_3 is re-estimated to system hydrological data when ϕ_6 is set to zero, and superscript I indicates independent. By defining parameters this way we can closely resemble conditional means and long-term variance in our joint model and a model with no connection between prices and inflows to reservoirs.

Nomenclature and parameter estimates

Table 5 provides the parameters used in the case study in Sect. 5. The last columns provide parameter estimates from the joint model (J) and the independent model (I). Section 5.2 explains the procedure for calibrating the joint model, while Appendix D presents a possible procedure for calibrating the independent model.

Discretization of the stochastic model

We consider a producer planning its hydropower operation for the next 2 years. It thus considers 104 weekly stages. We construct Markov chains according to Algorithm 1 from 50,000 weekly Monte Carlo simulations of prices and inflows to reservoirs from expressions (23) to (31). Since inflow is normally distributed in our model, it can simulate negative inflow values that we set to zero. Negative inflow is a challenge in SDDP and could alternatively be addressed by penalization [20]. Each week has a price and inflow pair, and we use the k -means method to cluster. Transition probabilities are estimated as the number of scenarios that travel from one cluster to the next, divided by the total number of scenarios moving through the previous cluster. A joint clustering on price and inflow pairs created a Markov chain that did not capture the variance, mean, and covariance of the underlying stochastic model. Instead, we only cluster on prices and map the corresponding inflows to each price cluster. These inflows must transform into a representative inflow to the price cluster. Importantly, they must retain the properties of the underlying stochastic model. The average value of the inflows manages to capture the covariance and mean, but not the variance. On the other hand, a random sample of the inflows captures the variance but not the covariance and mean values. A compromise is the mean of a limited number of samples, which manages to provide representative covariances, variances, and means (experiments on our application find five samples to be representative). The variance has

Table 5 Parameters for case study

	Description	(J)	(I)
<i>State variables</i>			
r_t [MWh]	Reservoir volume		
x_t [€/MWh]	Price deviation from long-term price level		
y_t [€/MWh]	Long-term price level		
η_t [TWh]	Unexplained system deviation state		
v_t [MWh/week]	De-seasonalized and normalized inflow		
h_t^{loc} [MWh/week]	Smoothed local inflow deviations (standardized)		
<i>Decision variables</i>			
g_t [MWh/week]	Production		
w_t [MWh/week]	Spillage		
<i>Parameters</i>			
γ	Discount factor	0.02	
\underline{R} [MWh]	Minimum reservoir volume	0	
\overline{R} [MWh]	Maximum reservoir volume	334,989	
\underline{G} [MWh/week]	Minimum production	0	
\overline{G} [MWh/week]	Maximum production	39,439	
ϕ_1 [€/MWh]	Price seasonality parameter, amplitude	4.421	
ϕ_2 [€/MWh]	Price seasonality parameter, shift	2.792	
ϕ_3 [€/MWh TWh]	Effect from h_t^{sys} on p_t	- 0.820	0
ϕ_4^* [1]	Short-term price AR-1-coef	0.978	0.977
ϕ_5^* [1]	Long-term drift	0.999	
ϕ_6 [1]	Effect on from h_t^{loc} on h_t^{sys}	0.002	0
ϕ_7^* [1]	Unexplained system deviation AR-1-coef	0.978	0.978
ϕ_8 [1]	Smoothing coefficient for local inflow	0.978	
ϕ_9^* [1]	Local inflow deviations AR-1-coef	0.356	
σ_1 [€/MWh]	Volatility short term price	1.872	2.292
σ_2 [€/MWh]	Volatility long-term price	0.554	
σ_3 [TWh]	Standard dev. of unexplained system deviation	0.940	1.323
σ_4 [MWh]	Standard deviation inflow	0.905	
$\bar{\mu}_t, t = 1, \dots, 52$ [MWh]	Weekly historical local inflow mean		
$\bar{\sigma}_t, t = 1, \dots, 52$ [MWh]	Weekly historical local inflow standard deviation		
<i>Values derived from states</i>			
ω_t^C [MWh]	Local inflow		
ω_t^O [€/MWh]	Spot price		
h_t^{sys} [TWh]	System reservoir volume, deviation from the mean		

the same shape as the continuous process' variance, but has a lower magnitude that decreases with the number of samples. However, an appropriate factor can scale the variance to the magnitude of the underlying stochastic model. Before scaling, we need to subtract the mean of all the associated inflows in a cluster and add it back afterwards, in order to not scale the mean as well. Algorithm 1 shows

pseudocode for this operation. We use a Markov chain with five price and inflow pair nodes. Simulations verify that it is representative of the continuous price and inflow processes.

Algorithm 1 Pseudocode to build Markov chains

- 1: Initialize parameters in Table 5 and decide inflow sample size N_{sample} and scaler κ .
 - 2: Simulate random variables χ , ξ , η , and ν using equations (25), (26), (28), and (30), drawing random samples according to (31).
 - 3: Calculate h^{*y^s} according to (27).
 - 4: Calculate prices ω_t^C according to (24).
 - 5: Calculate inflows ω_t^O according to (23) and set negative values to 0.
 - 6: Standardize prices ω_t^C .
 - 7: Use k -means algorithm to get price centroids.
 - 8: Identify which price (and inflow) scenario belongs to which centroid.
 - 9: **for** t in all stages **do**
 - 10: **if** first stage **then**
 - 11: Calculate transition probability at first stage as number of scenarios in a centroid divided by the total number of scenarios.
 - 12: **else**
 - 13: Count what centroid each scenario comes from and to what scenario it moves to. Calculate transition probabilities by dividing by number of scenarios that came from each centroid.
 - 14: Convert prices to original values by removing standardization.
 - 15: **for** c in all centroids **do**
 - 16: Make list I_c of all inflows that belong to centroid c .
 - 17: Calculate mean μ_{I_c} of I_c .
 - 18: Make new list \hat{I}_c by subtracting mean μ_{I_c} from inflows in I_c .
 - 19: Calculate the mean $\mu_{\hat{I}_c}$ from a sample of size N_{sample} from \hat{I}_c .
 - 20: Multiply $\mu_{\hat{I}_c}$ with scaler κ and then add μ_{I_c} to get the centroid's inflow. Set value to 0 if negative.
 - 21: **end for**
 - 22: **end if**
 - 23: **end for**
 - 24: Return transition probabilities and centroids with price and inflow pairs.
-

Additional case studies

This section contains results for different parameters to investigate the sensitivity of the findings. We first change parameters of the stochastic process to test stronger co-movements and higher seasonal price variations. Afterwards, we test different plant characteristics by reducing reservoir capacity and increasing generation capacity.

Extreme effect from hydrology on prices and higher seasonal price variations

To study the effect when seasonal price variations are higher and co-movements are stronger, we increase ϕ_1 to -1.5 and ϕ_3 to 10 . This also introduces a higher price variance. All other model parameters remain as they were. Figure 14a displays the reservoir trajectories under optimal policies. The behavior of the mean reservoir is similar to the base case. It shows that the producer favors a higher mean reservoir from around week 25 and is therefore more willing to risk spillage in this period. The 10th to 90th percentiles are narrower throughout the year in both the dependent and independent Markov chain compared to the base case. This is a result of increased seasonal price variations. Figure 14b shows a scatter plot of revenue. As in

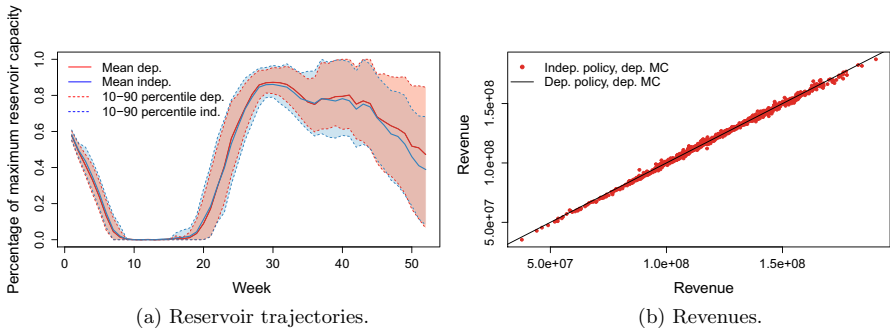


Fig. 14 Revenues and reservoir trajectories for the extreme case

Table 6 Difference in expected revenue of policies relative to $V_1^p(s_0, \omega_1)$ for the extreme case

Policy computation	Policy simulation	
	Independent MC	Dependent MC
Independent MC	0.68%	-0.24%
Dependent MC	-	0%

The first column tells which Markov chain that is used under policy computation, and the last two column headers tell which Markov chain that is used when simulating respective policies

the base case, the producer only slightly increases revenue in most scenarios when accounting for co-movements when establishing the operational policy. Table 6 outlines the percentage expected opportunity loss in the extreme case.

Changing plant characteristics: higher production capacity and lower reservoir capacity

The next case study keeps the stochastic model parameter estimates as the base case but changes the plant characteristics to $R = G = 106,382$ MWh. Both the generation capacity and reservoir capacity are equal to the average inflow of four weeks. The reservoir capacity is 32% and generation capacity is 270% of that of the base

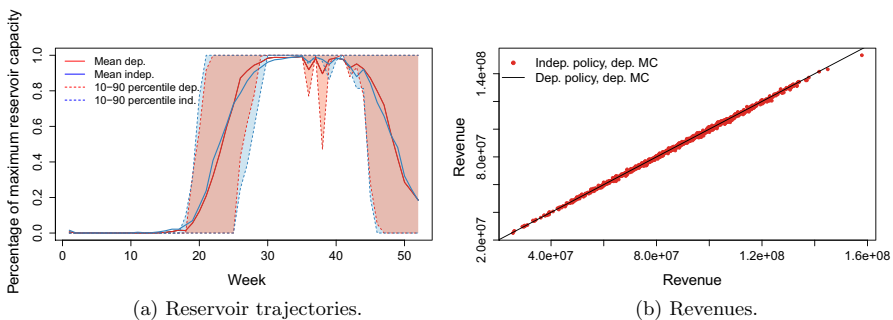


Fig. 15 Revenues and reservoir trajectories for case with different plant characteristics

Table 7 Difference in expected revenue of policies relative to $V_1^p(s_0, \omega_1)$ with different plant characteristics $R = G = 106,382$ MWh

Policy computation	Policy simulation	
	Independent MC	Dependent MC
Independent MC	2.94%	- 0.30%
Dependent MC	-	0 %

The first column tells which Markov chain that is used under policy computation, and the last two column headers tell which Markov chain that is used when simulating respective policies

case. Figures 15a, b show reservoir trajectories and revenues. The producer drains the reservoir already in the first week, because it can now produce the entire reservoir at any point in time, and the price seasonality trends downwards in January. Figure 15b demonstrates that the sample path revenues behave as in the other cases. Table 7 shows that the revenue improves by 0.30% in expectation.

Acknowledgements The authors would like to thank Selvaprabu Nadarajah for valuable discussions and inputs. We thank Eviny for case data. The first and third author gratefully acknowledge support from the research center HydroCen, RCN No. 257588.

Funding Open access funding provided by NTNU Norwegian University of Science and Technology (incl St. Olavs Hospital - Trondheim University Hospital).

Data availability Price data is available from commercial database providers; we are grateful to Montel for providing time series of futures price data. The system balance data is published by NVE (www.nve.no). Inflow and power plant data was graciously provided by Eviny, a hydropower producer.

Declarations

Conflict of interest Author Fleten is a guest editor and an associate editor with this journal. He has not been handling this paper in any editorial role.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Agrawal, S., Ding, Y., Saberi, A., et al.: Price of correlations in stochastic optimization. *Oper. Res.* **60**(1), 150–162 (2012). <https://doi.org/10.1287/opre.1110.1011>
2. Alexander, D.R., Mo, M., Stent, A.F.: Arithmetic Brownian motion and real options. *Eur. J. Oper. Res.* **219**(1), 114–122 (2012). <https://doi.org/10.1016/j.ejor.2011.12.023>

3. Arvanitidis, N.V., Rosing, J.: Composite representation of a multireservoir hydroelectric power system. *IEEE Trans. Power Appar. Syst.* **PAS 89**(2), 319–326 (1970). <https://doi.org/10.1109/TPAS.1970.292595>
4. Ávila, L., Mine, M.R., Kaviski, E., et al.: Evaluation of hydro-wind complementarity in the medium-term planning of electrical power systems by joint simulation of periodic streamflow and wind speed time series: A Brazilian case study. *Renew. Energy* **167**, 685–699 (2021). <https://doi.org/10.1016/j.renene.2020.11.141>
5. Benth, F.E., Koekkebakker, S., Ollmar, F.: Extracting and applying smooth forward curves from average-based commodity contracts with seasonal variation. *J. Deriv.* **15**(1), 52–66 (2007). <https://doi.org/10.3905/jod.2007.694791>
6. Bezanson, J., Edelman, A., Karpinski, S., et al.: Julia: a fresh approach to numerical computing. *SIAM Rev.* **59**(1), 65–98 (2017). <https://doi.org/10.1137/141000671>
7. Chopra, V.K., Ziemba, W.T.: The effect of errors in means, variances, and covariances on optimal portfolio choice. *J. Portfolio Manag.* **19**(2), 6–11 (1993). <https://doi.org/10.3905/jpm.1993.409440>
8. De Ladurantaye, D., Gendreau, M., Potvin, J.Y.: Optimizing profits from hydroelectricity production. *Comput. Oper. Res.* **36**(2), 499–529 (2009). <https://doi.org/10.1016/j.cor.2007.10.012>
9. Denault, M., Dupuis, D., Couture-Cardinal, S.: Complementarity of hydro and wind power: improving the risk profile of energy inflows. *Energy Policy* **37**(12), 5376–5384 (2009). <https://doi.org/10.1016/j.enpol.2009.07.064>
10. Dietze, M., Chavarry, I., Freire, A.C., et al.: A novel semiparametric structural model for electricity forward curves. *IEEE Trans. Power Syst.* **38**(4), 3268–3278 (2022). <https://doi.org/10.1109/TPWRS.2022.3197982>
11. Downward, A., Dowson, O., Baucke, R.: Stochastic dual dynamic programming with stagewise-dependent objective uncertainty. *Oper. Res. Lett.* **48**(1), 33–39 (2020). <https://doi.org/10.1016/j.orl.2019.11.002>
12. Dowson, O., Kapelevich, L.: SDDP.jl: a Julia package for stochastic dual dynamic programming. *Inf. J. Comput.* **33**(1), 27–33 (2021). <https://doi.org/10.1287/ijoc.2020.0987>
13. Dowson, O.: Applying stochastic optimisation to the New Zealand dairy industry. PhD thesis, University of Auckland (2018)
14. Duffie, D.: *Dynamic Asset Pricing Theory*. Princeton University Press, Princeton (2010)
15. Feng, Y., Ryan, S.M.: Scenario construction and reduction applied to stochastic power generation expansion planning. *Comput. Oper. Res.* **40**(1), 9–23 (2013)
16. Flatabø, N., Haugstad, A., Mo, B., et al.: Short-term and medium-term generation scheduling in the Norwegian hydro system under a competitive power market structure. In: *EPSOM'98 (International Conference on Electrical Power System Operation and Management)*, Switzerland (1998)
17. Fleten, S.E., Kristoffersen, T.K.: Short-term hydropower production planning by stochastic programming. *Comput. Oper. Res.* **35**(8), 2656–2671 (2008)
18. Fleten, S.E., Lemming, J.: Constructing forward price curves in electricity markets. *Energy Econ.* **25**(5), 409–424 (2003)
19. Girardeau, P., Leclere, V., Philpott, A.B.: On the convergence of decomposition methods for multi-stage stochastic convex programs. *Math. Oper. Res.* **40**(1), 130–145 (2015). <https://doi.org/10.1287/moor.2014.0664>
20. Gjelsvik, A., Mo, B., Haugstad, A.: Long- and medium-term operations planning and stochastic modelling in hydro-dominated power systems based on stochastic dual dynamic programming. In: *Handbook of Power Systems*. Springer, Berlin, pp. 33–55. https://doi.org/10.1007/978-3-642-02493-1_2 (2010)
21. Goodwin, D.: Schwartz-Smith 2-factor model—parameter estimation (2020). <https://www.mathworks.com/matlabcentral/fileexchange/43352-schwartz-smith-2-factor-model-parameter-estimation>. Accessed 23 June 2021
22. Labadie, J.W.: Optimal operation of multireservoir systems: State-of-the-art review. *J. Water Resour. Plan. Manag.* **130**(2), 93–111 (2004)
23. Lima, R.M., Conejo, A.J., Langodan, S., et al.: Risk-averse formulations and methods for a virtual power plant. *Comput. Oper. Res.* **96**, 350–373 (2018)
24. Lium, A.G., Crainic, T.G., Wallace, S.W.: Correlations in stochastic programming: a case from stochastic service network design. *Asia-Pac. J. Oper. Res.* **24**(02), 161–179 (2007)
25. Löhdorf, N., Shapiro, A.: Modeling time-dependent randomness in stochastic dual dynamic programming. *Eur. J. Oper. Res.* **273**(2), 650–661 (2019). <https://doi.org/10.1016/j.ejor.2018.08.001>

26. Löhndorf, N., Wozabal, D., Minner, S.: Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. *Oper. Res.* **61**(4), 810–823 (2013). <https://doi.org/10.1287/opre.2013.1182>
27. Lucia, J.J., Schwartz, E.S.: Electricity prices and power derivatives: evidence from the Nordic power exchange. *Rev. Deriv. Res.* **5**(1), 5–50 (2002). <https://doi.org/10.1023/A:1013846631785>
28. Maceira, M.E.P., Duarte, V., Penna, D., et al.: Ten years of application of stochastic dual dynamic programming in official and agent studies in Brazil: description of the NEWAVE program. In: 16th PSCC, Glasgow, Scotland, pp. 14–18 (2008)
29. Maceira, M.E., Melo A., Pessanha J.F, et al.: An approach to representing wind uncertainties in the long-term operation planning of systems with hydropower predominance. Tech. rep., EGU General Assembly, Vienna, Austria, 23–27 May 2022, EGU22-8993 (2022). <https://doi.org/10.5194/egusphere-egu22-8993>
30. Mo, B., Gjelsvik, A., Grundt, A.: Integrated risk management of hydro power scheduling and contract management. *IEEE Trans. Power Syst.* **16**(2), 216–221 (2001). <https://doi.org/10.1109/59.918289>
31. Montel Market data Nordic power (2021). <https://www.montelnews.com/en/>. Accessed 23 June 2021
32. Nersten, S., Dimoski, J., Fleten, S.E., et al.: Hydropower reservoir management using multi-factor price model and correlation between price and local inflow. In: 2018 IAEE International Conference (2018). <http://www.iaee.org/proceedings/article/15183>
33. NVE.: Norwegian Water Resources and Energy Directorate (2022). <https://www.nve.no/energi/analyse-og-statistikk/hydrologiske-data-til-kraftsituasjonsrapporten/>. Accessed 20 Feb 2022
34. Obersteiner, C., Sagan, M.: Parameters influencing the market value of wind power—a model-based analysis of the Central European power market. *Eur. Trans. Electr. Power* **21**(6), 1856–1868 (2011). <https://doi.org/10.1002/etep.430>
35. Øksendal, B.: Stochastic differential equations. In: *Stochastic Differential Equations: An Introduction with Applications*. Springer, Berlin, pp. 65–84 (2003). https://doi.org/10.1007/978-3-642-14394-6_5
36. Passos, A.C., Street, A., Barroso, L.A.: A dynamic real option-based investment model for renewable energy portfolios. *IEEE Trans. Power Syst.* **32**(2), 883–895 (2016)
37. Pereira, M.V., Pinto, L.M.: Multi-stage stochastic optimization applied to energy planning. *Math. Program.* **52**(1), 359–375 (1991). <https://doi.org/10.1007/BF01582895>
38. Pérez-Díaz, J.I., Guisández, I., Chazarra, M., et al.: Medium-term scheduling of a hydropower plant participating as a price-maker in the automatic frequency restoration reserve market. *Electr. Power Syst. Res.* **185**(106), 399 (2020). <https://doi.org/10.1016/j.epsr.2020.106399>
39. Philpott, A., Guan, Z.: On the convergence of stochastic dual dynamic programming and related methods. *Oper. Res. Lett.* **36**(4), 450–455 (2008). <https://doi.org/10.1016/j.orl.2008.01.013>
40. Schwartz, E., Smith, J.E.: Short-term variations and long-term dynamics in commodity prices. *Manage. Sci.* **46**(7), 893–911 (2000). <https://doi.org/10.1287/mnsc.46.7.893.12034>
41. Shapiro, A., Tekaya, W., da Costa, J.P., et al.: Risk neutral and risk averse stochastic dual dynamic programming method. *Eur. J. Oper. Res.* **224**(2), 375–391 (2013). <https://doi.org/10.1016/j.ejor.2012.08.022>
42. Stedinger, J.R.: Fitting log normal distributions to hydrologic data. *Water Resour. Res.* **16**(3), 481–490 (1980). <https://doi.org/10.1029/WR016i003p00481>
43. Suomalainen, K., Pritchard, G., Sharp, B., et al.: Correlation analysis on wind and hydro resources with electricity demand and prices in New Zealand. *Appl. Energy* **137**, 445–462 (2015). <https://doi.org/10.1016/j.apenergy.2014.10.015>
44. Wallace, S.W., Fleten, S.E.: Stochastic programming models in energy. In: Ruszczyński, A., Shapiro, A., (eds) *Stochastic Programming*. Vol. 10 of *Handbooks in Operations Research and Management Science*. Elsevier Science, pp 637–677 (2003). [https://doi.org/10.1016/S0927-0507\(03\)10010-2](https://doi.org/10.1016/S0927-0507(03)10010-2)
45. Wolfgang, O., Haugstad, A., Mo, B., et al.: Hydro reservoir handling in Norway before and after deregulation. *Energy* **34**(10), 1642–1651 (2009). <https://doi.org/10.1016/j.energy.2009.07.025>