

**Incorporating Risk in Operational Water Resources Management
Probabilistic Forecasting, Scenario Generation, and Optimal Control**

van der Heijden, Ties; Mendoza-Lugo, Miguel Angel; Palensky, Peter; van de Giesen, Nick; Abraham, Edo

DOI

[10.1029/2024WR037115](https://doi.org/10.1029/2024WR037115)

Publication date

2025

Document Version

Final published version

Published in

Water Resources Research

Citation (APA)

van der Heijden, T., Mendoza-Lugo, M. A., Palensky, P., van de Giesen, N., & Abraham, E. (2025). Incorporating Risk in Operational Water Resources Management: Probabilistic Forecasting, Scenario Generation, and Optimal Control. *Water Resources Research*, 61(3), Article e2024WR037115. <https://doi.org/10.1029/2024WR037115>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Water Resources Research®



RESEARCH ARTICLE

10.1029/2024WR037115

Key Points:

- A framework is proposed to estimate operational uncertainty, generate scenarios, and the application of risk-aware decision-making
- Taking operational uncertainty into account in control can lead to up to 44 p.p. lower energy cost compared to a deterministic forecast
- The use of Exceedance Risk constraints on water levels allows for a nuanced approach, leading to safe water levels and lower energy cost

Correspondence to:

T. van der Heijden,
T.J.T.vanderHeijden-2@tudelft.nl

Citation:

van der Heijden, T., Mendoza-Lugo, M. A., Palensky, P., van de Giesen, N., & Abraham, E. (2025). Incorporating risk in operational water resources management: Probabilistic forecasting, scenario generation, and optimal control. *Water Resources Research*, 61, e2024WR037115. <https://doi.org/10.1029/2024WR037115>

Received 12 JAN 2024

Accepted 8 FEB 2025

Author Contributions:

Conceptualization: Ties van der Heijden, Edo Abraham

Data curation: Ties van der Heijden

Formal analysis: Ties van der Heijden

Funding acquisition: Ties van der Heijden, Edo Abraham

Investigation: Ties van der Heijden

Methodology: Ties van der Heijden, Miguel Angel Mendoza-Lugo, Edo Abraham

Project administration: Ties van der Heijden, Edo Abraham

Software: Ties van der Heijden, Miguel Angel Mendoza-Lugo

Supervision: Peter Palensky, Nick van de Giesen, Edo Abraham

Validation: Ties van der Heijden

Visualization: Ties van der Heijden

Writing – original draft: Ties van der Heijden, Edo Abraham

© 2025. The Author(s).

This is an open access article under the terms of the [Creative Commons Attribution License](#), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

Incorporating Risk in Operational Water Resources Management: Probabilistic Forecasting, Scenario Generation, and Optimal Control

Ties van der Heijden^{1,2,3,4} , Miguel Angel Mendoza-Lugo⁵, Peter Palensky² , Nick van de Giesen¹ , and Edo Abraham¹

¹Department of Water Management, Faculty of Civil Engineering, Delft University of Technology, Delft, The Netherlands,

²Department of Electrical Sustainable Energy, Faculty of Electrical Engineering, Mathematics, and Computer Science, Delft University of Technology, Delft, The Netherlands, ³HKV Consultants, Delft, The Netherlands, ⁴Pythia Energy Intelligence, Delft, The Netherlands, ⁵Department of Hydraulic Engineering, Faculty of Civil Engineering, Delft University of Technology, Delft, The Netherlands

Abstract This study presents an innovative approach to risk-aware decision-making in water resource management. We focus on a case study in the Netherlands, where risk awareness is key to water system design and policy-making. Recognizing the limitations of deterministic methods in the face of weather, energy system, and market uncertainties, we propose a scalable stochastic Model Predictive Control (MPC) framework that integrates probabilistic forecasting, scenario generation, and stochastic optimal control. We utilize Combined Quantile Regression Deep Neural Networks and Non-parametric Bayesian Networks to generate probabilistic scenarios that capture realistic temporal dependencies. The energy distance metric is applied to optimize scenario selection and generate scenario trees, ensuring computational feasibility without compromising decision quality. A key feature of our approach is the introduction of Exceedance Risk (ER) constraints, inspired by Conditional-Value-at-Risk (CVaR), to enable more nuanced and risk-aware decision-making while maintaining computational efficiency. In this work, we enable the Noordzeekanaal–Amsterdam-Rijnkanaal (NZK-ARK) system to participate in Demand Response (DR) services by dynamically scheduling pumps to align with low hourly electricity prices on the Day Ahead and Intraday markets. Through historical simulations using real water system and electricity price data, we demonstrate that incorporating uncertainty can significantly reduce operational costs—by up to 44 percentage points compared to a deterministic approach—while maintaining safe water levels. The modular nature of the framework also makes it adaptable to a wide range of applications, including hydropower and battery storage systems.

Plain Language Summary This research introduces a new method for managing water resources in the Netherlands that considers uncertainties in weather, sea-levels, and energy markets. Traditional methods often ignore these uncertainties, which can lead to inefficiencies and higher costs. Our approach combines advanced forecasting techniques and decision-making strategies to better account for these unpredictable factors. Using cutting-edge computer models, we predict future conditions as multiple scenarios with assigned probabilities. These predictions guide decisions about when to operate water pumps, aiming to save energy costs while keeping water levels safe. By enabling the Noordzeekanaal–Amsterdam-Rijnkanaal (NZK-ARK) system to adjust pumping schedules in response to electricity price fluctuations, we show that this approach can achieve significant cost savings—up to 44 percent compared to traditional methods. Our method is flexible and can be applied to other systems like hydropower plants or battery storage facilities, demonstrating its potential for broader use. This study shows how combining risk-aware forecasting and operational strategies can improve the efficiency and safety of water management systems while supporting the energy transition.

1. Introduction

The energy transition represents a path toward sustainability and resilience in the face of global climate change. In Europe, this journey is characterized by the integration of intermittent Renewable Energy Sources (RES), the smooth operation of an integrated European energy market, and the transition from a centralized to a distributed grid infrastructure. These elements are key in achieving the European Union's goals for a climate-neutral future. However, overcoming challenges in infrastructural limitations, policy and market dynamics, and the variability of RES is essential to secure an independent and climate-neutral energy system.

Writing – review & editing: Ties van der Heijden, Miguel Angel Mendoza-Lugo, Peter Palensky, Nick van de Giesen

Flexibility in the energy sector is increasingly recognized as a cornerstone for successfully integrating RES into the power grid (TenneT & Gasunie, 2022). As countries move toward more sustainable energy solutions, the inherent intermittency of variable renewable resources such as wind and solar power presents new challenges. Their variability introduces the need for innovative approaches to maintain a stable and reliable energy supply. Demand Response (DR) is an approach to achieve demand-side flexibility, and it refers to the changes in electric usage by end-use customers from their normal consumption patterns. These signals are achieved through prices that reflect scarcity of supply and marginal cost of production. The variable supply of renewables leads to changing electricity prices over time, and induces lower electricity use at times of high market prices. By enabling dynamic adjustments of demand, DR contributes to the grid's stability, reduces the need for peak generation capacity, and optimizes energy costs for consumers. European spot markets are designed to incentivize consumers to adjust their energy usage in response to variable prices, giving DR its business case. By incorporating DR into energy systems, consumers can leverage market mechanisms to achieve a more resilient, cost-efficient, and sustainable energy system.

The increasing variability and uncertainty in water resources management due to climate change, socio-economic shifts, and evolving energy demands necessitate adaptive strategies that leverage real-time data and probabilistic forecasts. Across the globe, water systems are under increasing pressure to balance competing demands for irrigation, energy production, flood control, and ecological preservation (Verhagen et al., 2021). These challenges are further exacerbated by the need for greater integration with renewable energy systems, which introduce additional operational uncertainties. The ability to dynamically adapt to these uncertainties is essential for maintaining the reliability and sustainability of water resources under changing conditions.

Recent advancements in real-time monitoring technologies and forecasting methods have enabled a shift from traditional rule-based management approaches to more sophisticated adaptive frameworks (Castelletti et al., 2023). Probabilistic forecasts, in particular, have emerged as a critical tool for representing uncertainty and enabling more informed decision-making. By accounting for a range of possible future scenarios, these forecasts provide the basis for strategies that can adapt dynamically to changing conditions, ensuring resilience and efficiency in water resources management.

The Netherlands, facing unique geographical challenges, has developed significant expertise in water management. Much of its land is below sea level, requiring methods to manage precipitation and groundwater levels. In the past, pumping stations with rule-based or deterministic model predictive control (MPC) have been central to these management efforts. However, as climate change can lead to more unpredictable weather, the limits of deterministic approaches are being reached. This situation incentivized a move toward adaptive strategies, leveraging real-time data and probabilistic forecasts in the optimization process. These advances enable more effective adaptation to increasingly variable climatic conditions and evolving hydrological regimes within the catchments, improving overall water management resilience (Castelletti et al., 2008, 2023; Giuliani et al., 2021).

Stochastic MPC approaches can employ chance constraints to ensure that certain conditions are met with a predefined probability level (Mesbah, 2016). However, their inherent binary nature, indicating that they either meet the conditions or don't, overlooks the potential magnitude or likelihood of constraint violations, which is critical in applications such as water management. Additionally, while robust optimization techniques aim to handle uncertainty by considering the worst-case scenarios, they can also be overly conservative, leading to suboptimal or overly cautious operational decisions that may not capitalize on potential opportunities or efficiencies. In stochastic MPC, uncertainty can be addressed explicitly by specifying the probability distribution functions of uncertain variables, or implicitly through sets of scenarios derived from ensemble hydro-meteorological forecasts. Recent literature often favors these implicit approaches, not only because they can incorporate model uncertainty but also because they can organize large scenario ensembles into structured scenario trees. Such trees capture the temporal evolution and dependency of uncertainties while reducing computational complexity compared to using ensembles directly (Castelletti et al., 2023). Generally, explicit approaches do not account for autocorrelation in disturbances or forecast errors (Castelletti et al., 2008; Pianosi & Soncini-Sessa, 2009), which can relatively easily be implicitly taken into account in scenario sets.

Adapting to the unpredictability of weather and its impact on water levels requires sophisticated forecasting methods. The integration of probabilistic forecasts not only enhances prediction accuracy but also enhances dynamic responses such as Demand Response (DR) in the energy sector to achieve a necessary balance between supply and demand. DR refers to the changes in electric usage by end-use customers from their normal

consumption patterns. These signals are achieved through prices that reflect scarcity of supply and marginal cost of production. The variable supply of renewables leads to changing electricity prices over time, and induces lower electricity use at times of high market prices. In DR, inaccuracies in forecasting can significantly impact system efficiency, leading to both operational challenges and financial consequences for consumers and producers. With the increasing reliance on renewable energy sources, accurately forecasting energy production has become more complex. This unpredictability and complexity call for enhanced predictive models that adequately reflect the operational uncertainty given the latest information. Moreover, integrating stochastic optimization methods into energy management strategies can mitigate the risks associated with operational uncertainty. These methods enable the consideration of multiple possible future scenarios in the decision-making process, thus providing a more comprehensive strategy that accounts for a wide range of outcomes. By adopting a probabilistic approach to forecasting and optimization, stakeholders can better prepare for and adapt to the inherent volatility associated with Renewable Energy Sources and demand fluctuations and the operational risk involved in their processes (Pape, 2018; Sahu & McLaughlin, 2018; Sweeney et al., 2020), inspiring research in uncertainty estimation in forecasting and making it an increasingly present industry requirement (Bessa et al., 2017; Makridakis et al., 2022; van der Heijden, Palensky, & Abraham, 2021). Similarly, ensemble forecasts have also become a standard in hydro-meteorological forecasting due to their ability to capture forecast uncertainty (Castelletti et al., 2023; Zhao et al., 2021).

Central to this study are the uncertainties faced in operational water resources management, from variable discharges to dynamic seawater levels influenced by tides and winds. The inherent volatility of energy spot markets, combined with real-time supply demand dynamics, and the integration of renewable energy, exemplify the need for advanced forecasting methodologies. This motivated our exploration into probabilistic methodologies, such as the Combined Quantile Regression Deep Neural Network (van der Heijden, Palensky, & Abraham, 2021) (CQRDNN) and the Non-parametric Bayesian network (van der Heijden, Palensky, et al., 2022) (NPBN). Our framework integrates probabilistic forecasting, scenario generation, and reduction approaches. We further extend the scenario-based MPC approach into Tree-based MPC (TB-MPC) (Maestre et al., 2013; Raso et al., 2014) formulations for improved and sometimes necessary computational efficiency in large uncertainty representations.

Our approach emphasizes risk-based constraints, inspired by the Conditional-Value-at-Risk (CVaR) approach, also known as the mean excess loss (Rockafellar & Uryasev, 2000). CVaR represents the tightest convex formulation of a chance constraint (Venkatasubramanian et al., 2020), while its linear formulation leads to reduced computational complexity. Unlike traditional chance constraints or robust optimization methods, CVaR represents the tail distribution of potential outcomes, revealing worst-case scenarios and associated risks. We propose the use of Exceedance Risk (ER) constraints that are formulated similarly to CVaR, but are tailored for the application in operational water resources management. In the Dutch water sector, risk-awareness is key to system design and policy-making. Our work considers probabilistic risk of exceedances for water levels in real-time operational control. We formulate a stochastic MPC problem, where the states are stochastic variables driven by uncertain disturbances and parameters, and the objective and constraints are formulated as expectation of cost and probabilistic constraint violations, respectively. As such, the optimal control problem of the MPC requires solving stochastic optimization problems via a scenario-based approach. We also specifically consider CVaR-type stochastic problems to model water-level exceedance risks.

In response to the societal challenges posed by climate variability and the energy transition on water management systems, this manuscript introduces several novel methodologies to capture dependencies between weather, markets, and control decisions, and shows how to exploit these in control for our case study area, the drainage canal Noordzeekanaal–Amsterdam-Rijnkanaal (NZK-ARK). First, we expand the use of probabilistic forecasts to open canal systems and demand response applications, areas where their adoption has been limited. We propose a Combined Quantile Regression Deep Neural Network (CQRDNN) (Section 3.1), which leverages deep learning to provide probabilistic forecasts of water levels and energy demands. Unlike traditional models, CQRDNN incorporates a quantile-based approach that inherently manages forecast uncertainty, offering a range of potential outcomes crucial for effective decision-making under uncertainty.

Second, we introduce a Non-parametric Bayesian network (NPBN) (Section 3.3) for generating conditional scenarios that respect the temporal dynamics of water systems and energy markets. This method allows for a more realistic simulation of future states, enhancing the strategic planning capabilities of water resource managers. This

novel combination improves computational efficiency, which is critical for real-time operational adjustments while allowing for incorporating advanced risk-based constraints and a stochastic objective tailored for water management and DR. We further propose the energy distance (Ziel, 2020) for scenario subset selection and propose a scenario-tree generation method (Section 3.5) to further reduce the computational complexity of the optimization problem by applying TB-MPC.

Third, we integrate these ensembles into stochastic optimization frameworks, an approach not yet widely explored in water management applications. We introduce the ER measure (Section 4) as a novel reformulation of chance constraints. ER constraints are designed to efficiently manage operational risk on water level bound exceedance. Lastly, we demonstrate how this integrated framework enables energy flexibility services from safety-critical infrastructure, such as electric pumping stations, under operational uncertainty. These contributions significantly advance risk-aware decision-making for interconnected water and energy systems.

By examining the operational flexibility of electric pumping stations and their role in the broader energy system, this manuscript aims to contribute to a more sustainable and integrated approach to managing the interdependencies between water and energy in the Netherlands. The presented innovations collectively enhance the ability of water management systems to operate efficiently in the face of uncertainty, optimizing both water usage and energy expenditure. The integration of these methodologies not only addresses the immediate needs of operational control but also sets a foundation for future research and application in sectors grappling with similar variability and uncertainty, notably in the context of energy flexibility and DR.

2. Demand Response in the Netherlands: Case Study IJmuiden

In this section, we describe the case study for our proposed framework. First, we describe the water system; the NZK-ARK, located in IJmuiden, the Netherlands (Section 2.1). Then we describe the Dutch electricity markets that are considered for DR services which recent studies have shown to offer significant flexibility and cost reduction potential (van der Heijden, Lugt, et al., 2022) (Section 2.2).

2.1. The Noordzeekanaal–Amsterdam–Rijnkanaal

The NZK–ARK is an intricate open canal system. It's equipped with multiple undershot gates and has a pumping station at IJmuiden, which aids in consistently directing water into the North Sea, regardless of the sea's variable water levels. Water from four regional water boards flows into the NZK–ARK, helping redirect excess rainwater. This water is methodically channeled or pumped out to the North Sea. A detailed representation of this water system, focusing on the inflow and outflow mechanisms, is provided in Figure 1a. A simplified schematic of the system is presented in Figure 1c.

The operations at the IJmuiden pumping station highlight the complexities in decision-making caused by the presence of operational uncertainty. IJmuiden is not only Europe's largest pumping station, but also plays a key role in managing ship traffic, controlling saltwater intrusion, and working with four local water authorities. The complexity of its operations is further amplified with the possible addition of Demand Response (DR) and energy trading in the spot market. Previous work has shown that a control strategy that taps into varying electricity prices from the Day Ahead Market (DAM) and the Intraday Market (IDM) could prove more cost-effective than the current method, especially as the penetration of renewable energy grows (van der Heijden, Lugt, et al., 2022). However, that analysis was based on perfect foresight and did not consider operational uncertainty. The interaction between water management and energy usage in pumping stations is increasingly significant within the context of the water-energy nexus (Dang Doan et al., 2013; Horváth et al., 2022; Pour et al., 2022). Despite this significance, the opportunity to optimize pump schedules with regard for the energy markets is still largely untapped in practice.

Within the MPC's internal model, the canal is depicted as a linear reservoir with predefined surface area. This approach is consistent with the current control system of the NZK–ARK. Owing to the canal's considerable depth and width, water movement occurs at reduced speeds, ensuring that friction remains minimal (Goedbloed, 2006). The primary sources of water inflow come from the water board discharges and water from the Rhine directed through Maarssen, the outflow is directed toward the North Sea.

At IJmuiden, the gate operations are dictated by specific water level differentials. They can be activated when there's a 16 cm water level difference and deactivated at 12 cm. This modulation is necessary to account for the

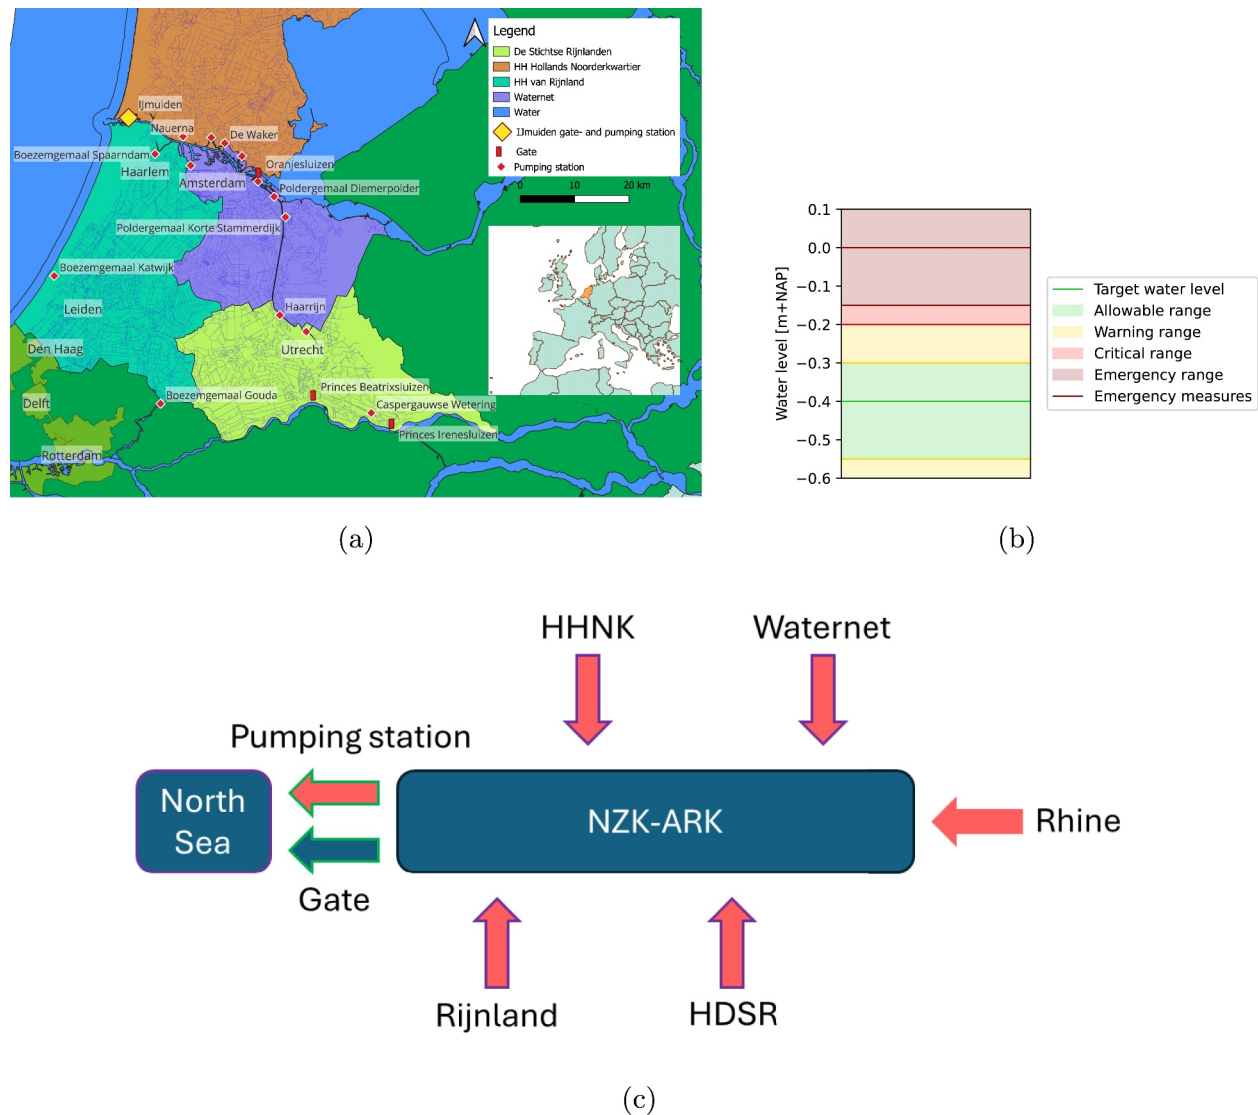


Figure 1. The Netherlands, Noordzeekanaal–Amsterdam–Rijnkanaal. (a) The water management area of different local water authorities are shown by color. Pumping stations owned by local water authorities are shown by red diamond, gate structures with red rectangles, and the IJmuiden gate- and pumping station with a yellow diamond. (b) The water level regime of the NZK-ARK showing target ranges and warning/emergency levels. The water levels are shown in m + NAP, which is the standard reference for water levels in the Netherlands, equivalent to Amsterdam Ordnance Datum. (c) NZK-ARK schematic for control, where the purple-outlined items indicate forecast variables, and the green-outlined items indicate control variables.

contrasting densities of salt and freshwater, along with inherent system friction (Janssen, 2017). With automated controls in place, these gates have a maximum discharge rate of $500 \text{ m}^3 \text{ s}^{-1}$, a precaution to protect the foundation of the gate complex. The IJmuiden station incorporates six distinct pumps, with a combined maximum power consumption nearing 6 MW. For modeling convenience, the MPC represents these multiple pumps as a singular entity, merging their unique attributes. Even though there are six individual pumps, an integrated approach is adopted, utilizing various Q–dH curves to represent the pumping station's overall discharge ability. Another composite representation is designed to represent the pumping station's energy use in a single equivalent PQH-curve, encapsulating pump energy consumption in relation to discharge levels and the pump head. A comprehensive description of these modeling techniques is available in van der Heijden, Lugt, et al. (2022).

At present, the IJmuiden pumping station operates using MPC, focusing on energy efficiency by drawing power from the Futures market with fixed-price contracts. It works with a 24-hr prediction, ensuring water levels stay




Figure 2. 2D-KDE of top: the Dutch IDM electricity prices over the DAM prices, bottom: the relative IDM prices over the DAM prices (2019–2022).

within set limits. To better fit the needs of energy trading on the DAM, we've expanded the prediction horizon to 48-hr. This change allows for a detailed look at the cost-saving potential of energy cost optimization over flexible energy markets for the NZK–ARK system, and aligns the control system with DAM rules. We've considered uncertainties in sea water levels, incoming discharge from local water authorities, and electricity prices. While sea water levels and discharge, which directly affect operational constraints, are considered jointly in detail, electricity prices are addressed separately to focus on minimizing expected costs. By treating price scenarios independently, the optimization can account for market variability without adding extra complexity or introducing additional decision points within the control horizon.

2.2. Energy Markets in the Netherlands

The energy trading landscape in the Netherlands is predominantly facilitated through the DAM and the IDM, the two main markets for trading renewable energy in Europe. These markets reward the exploitation of energy flexibility through DR services. On the DAM, energy transactions are made a day in advance, with consumers placing bids for the upcoming day in hourly blocks. By 12:00 CET each day, bids are collected on the Amsterdam Power Exchange (APX)-Endex, where the market establishes a price equilibrium where demand meets supply.

In contrast, the IDM offers a more dynamic trading environment, allowing participants to continuously buy or sell energy in quarterly, half-hourly, or hourly segments, sometimes up to just 5 min before the actual delivery time. Due to this flexible nature, individual buy and sell orders set varying contract prices. This study utilizes the volume-weighted price over the 3 hours leading up to delivery, termed the ID3-price. However, in real-world scenarios, every transaction could exhibit a distinct price. The IDM serves as a platform to adjust day-ahead plans and sidestep potential imbalances. Given its adaptability, many speculate that the IDM is poised to become the primary trading platform for renewable energy in the near future (De Jong et al., 2017), a hypothesis that is being confirmed by the yearly doubling of traded volume.

Both DAM and IDM incentivize participants to tailor their energy schedules in response to price signals that echo supply scarcity and production costs. By adhering to such price-sensitive DR strategies, there's potential for shifts in energy consumption patterns in correspondence with price changes (Jordehi, 2019). Figure 2 shows a 2-Dimensional Kernel Density Estimate (KDE) that sheds light on the evolving trends of DAM and IDM prices in the Netherlands. Interestingly, Dutch prices have demonstrated more pronounced fluctuations over the years, which might be attributed to various factors, including the maturity of the market, changing consumption patterns in the COVID pandemic, the lack of nuclear base-load in France, and the Russian invasion of Ukraine leading to a reduced gas and oil supply to Europe, and increased prices due to embargos. This caused spot market fluctuations to increase from around 0–75 [€/MWh] in 2019 to –200 – +600 [€/MWh] in 2022. Interestingly, the relative ID₃ price seems normally distributed around the DAM price and has stayed in a similar range since 2020.

While the Futures market provides long-term or base-load trading opportunities, its fixed pricing model doesn't facilitate the intermittent nature of renewable energy. As energy supply becomes more uncertain, it is expected that fixed-price contracts will carry higher risk premiums, giving flexibility a business case. Energy users might turn to the Futures market to hedge against risks, and producers find it useful to guarantee consistent sales. However, its inherent rigidity with predetermined prices makes it less suited for strategies that aim to exploit energy flexibility.

In the following sections, we extend the multi-market approach as defined in (van der Heijden, Lugt, et al., 2022) to incorporate operational uncertainty. The proposed strategy combines both DAM and IDM, optimizing the

strengths of each. Combining multiple markets promises to amplify the efficacy of market-based DR initiatives (Schwabeneder et al., 2021). With the rising prominence of renewable energy sources, the roles of DAM and IDM are set to further magnify in order to keep supply and demand in balance.

3. Modeling Operational Uncertainty

In this section, we describe the methods applied to model the operational uncertainty and make it suitable for scenario-based MPC. We first describe the CQRDNN to forecast quantiles of water levels, pumped discharge, and electricity prices (Section 3.1), and then we describe the search algorithm used for feature and hyperparameter optimization (Section 3.2). We continue with how an Non-parametric Bayesian Networks can be applied to sample time series from the regression quantiles with realistic temporal dependency (Section 3.3), and finish with scenario reduction techniques for optimal subset selection (Section 3.4) and scenario tree generation (Section 3.5).

3.1. Combined Quantile Regression Deep Neural Network

Operational uncertainty is characterized by both variability and unpredictability. We propose the use of the CQRDNN (van der Heijden, Palensky, & Abraham, 2021) for modeling operational uncertainty in incoming waterboard discharge, the water level of the North Sea, and the DAM electricity prices. The CQRDNN, a neural network architecture, is well-suited to capture non-linear data patterns and can be applied domain-free, leading to a manageable workflow complexity. The advantage of the CQRDNN is its ability to address the crossing quantile problem observed in ensemble models, where individual quantiles are represented by separate models. Ensemble models can display non-monotonicity, an issue where higher quantiles provide values lower than those of lower quantiles.

To address this problem, the CQRDNN employs a combined quantile loss function, which enables the training of multiple quantiles within a single DNN, benefiting the natural order of quantiles. The architecture of the CQRDNN is illustrated in Figure 3. In this work, we forecast 13 quantiles (0.01, 0.05, 0.1, 0.2, ...) to represent the operational uncertainty at a given time step.

Training the model involves the use of the pinball loss functions (Koenker & Bassett, 1978), sometimes referred to as the quantile regression error function.

$$L_{\tau} = \max(\tau \cdot e, (\tau - 1) \cdot e), \quad (1)$$

$$e = y - \hat{y}, \quad (2)$$

$$L_{CQ} = \frac{1}{N} \sum_{n=1}^N L_{\tau_n}, \quad (3)$$

where L denotes the loss, τ represents the quantiles, e signifies the quantile forecast error, with y being the observed value and \hat{y} the quantile forecast. L_{CQ} is the combined quantile loss function applied to the neural network to train all N quantiles simultaneously.

Forecasting performance of the models is evaluated based on the Mean Absolute Error (MAE) of the expected value of the forecast (50th percentile), the discrete approximation of the Continuous Ranked Probability Score (CRSP), and the Prediction Interval Confidence Percentage (PICP) of the 80% prediction interval, and the Prediction Interval Normalized Average Width (PINAW) of the 80% prediction interval. These metrics are defined as.

$$\text{MAE}(y, \hat{y}) = \frac{1}{N} \sum_{i=1}^N |y - \hat{y}|, \quad (4)$$

$$\text{CRPS}(y, \hat{y}) = \frac{1}{N} \sum_{\tau \in T} \max(\tau \cdot (\hat{y}_{\tau} - y), (\tau - 1) \cdot (\hat{y}_{\tau} - y)), \quad (5)$$

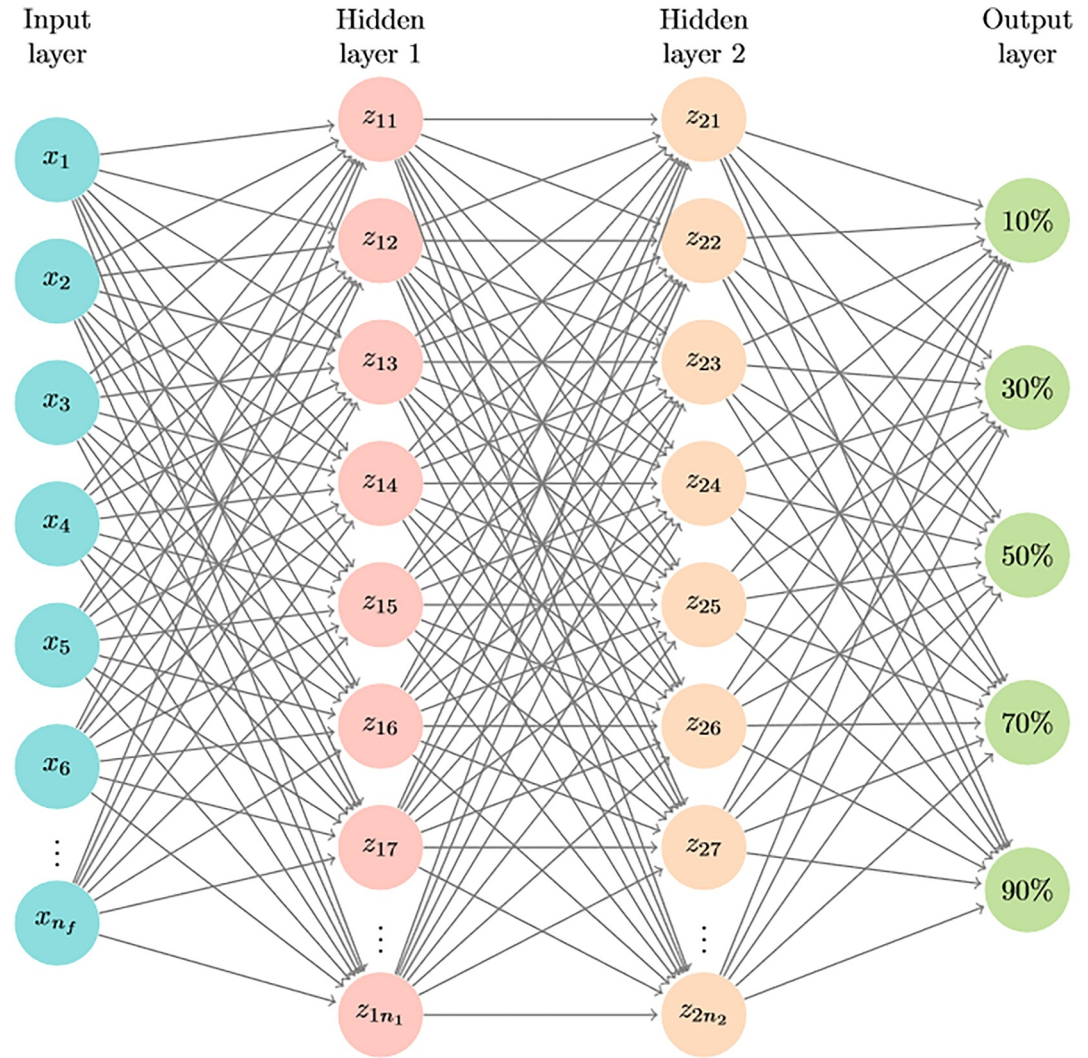


Figure 3. Structure of the CQRDNN showcasing separate quantile output nodes. Adopted from van der Heijden, Palensky, & Abraham (2021).

$$\text{PICP}_{\tau_0:\tau_1}(y, \hat{y}) = \frac{1}{N} \sum_{i=1}^N \mathbb{I}_{\hat{y}_{\tau_0} \leq y \leq \hat{y}_{\tau_1}}, \quad (6)$$

$$\text{PINAW}_{\tau_0:\tau_1}(y, \hat{y}) = \frac{\hat{y}_{\tau_1} - \hat{y}_{\tau_0}}{y_{\max} - y_{\min}}, \quad (7)$$

where y is the observation, \hat{y} the expected value of the quantile forecast, and \hat{y}_{τ} the τ -quantile. The MAE depicts an image of the error of the expected value of the quantile forecast, the CRPS (continuous ranked probability score) is a combined weighted metric for the quantile-errors, the PICP (prediction interval coverage percentage) shows the amount of observations that fall within a predefined prediction interval (PI) (e.g., the 80% PI should cover 80% of the observations), and the PINAW (prediction interval normalized average width) depicts the normalized width of a predefined PI. An ideal 80% PI should therefore have a low PINAW but a PICP of exactly 80%.

3.2. Combined Feature and Hyperparameter Optimization

For accurate and effective application of ML models it is necessary to optimize both the features applied in the model and the hyperparameters of the model. The Tree Parzen Estimator (TPE) algorithm, which is a variant of

Sequential Model Based Optimization, is capable of simultaneously optimizing both features and hyperparameters (Bergstra et al., 2013; Lago et al., 2018; van der Heijden, Lago, et al., 2021). It allows for the construction of a custom search space with a large number of dimensions which can be either discrete, continuous, or a combination.

The TPE employs Bayes' rule to create a surrogate model. The primary differentiator here is how the TPE segments the search space by evaluating the likelihood of the observed loss being above or below a threshold, typically denoted as y^* . In this work, we apply the TPE to optimize the features and hyperparameters of our CQRDNNs, where the loss being optimized for is the combined pinball loss function.

This surrogate model describes the probability of the loss being higher ($h(x)$) or lower ($l(x)$) than a specified threshold value (y^*) as a function of the search space instantiation:

$$p(y|x) = \frac{p(x|y) \cdot p(y)}{p(x)}, \quad (8)$$

In the above equation, y represents the model performance, and x denotes a search space instantiation.

Consequently, the model performance is estimated based on the features and hyperparameters, where $p(x|y)$ is defined as

$$p(x|y) = \begin{cases} l(x) & \text{if } y < y^* \\ h(x) & \text{if } y \geq y^*. \end{cases} \quad (9)$$

Within the TPE algorithm, samples are drawn from both $l(x)$ and $h(x)$, after which the ratio $\frac{l(x)}{h(x)}$ is evaluated for all samples. The next suggested candidate is the one with the highest expected improvement, that is, the candidate with the largest ratio between low and high probabilities in $l(x)$ and $h(x)$, respectively. In our case, we apply the combined pinball loss function (Equation 3).

To optimize the CQRDNN for forecasting pumped discharge by the four local water authorities, as well as the sea water level of the North Sea and Day Ahead Market (DAM) electricity prices, we apply the same hyperparameter search space for all modeling tasks while adjusting the feature search space. Details of the feature and hyperparameter search spaces are provided in Table C1. Model selection is based on the combined pinball loss (Equation 3) and an independent validation data set. All results presented in this manuscript are derived from independent test data sets. For the DAM, the model is retrained weekly to adapt to the rapidly changing market dynamics.

For pumped discharge modeling, data from early 2014 to January 2021 was used, with the period from 2019 to 2021 serving as the test set, 2016 as the validation set for feature and hyperparameter optimization, and the remaining years for training. The validation year was chosen based on preliminary analysis to represent an average year in terms of discharge rates. In essence, the CQRDNN is trained to model a rainfall-runoff system influenced by human interventions through pump control. To achieve this, we optimize a 'rolling window' approach by summing fluxes over a variable window length (rolling window size in Table C1) and selecting features such as precipitation and evaporation through a binary search space. For waterboards Rijnland and HHNK, we also include sea water level and wind as optional features, as these influence direct discharge to the North Sea. This approach works best for the waterboard Waternet, possibly due to the relatively low amount of storage present in the Waternet system. The NZK-ARK is used as a drainage canal (in Dutch: 'boezem'), where most waterboards have an intermediate drainage canal with higher storage capacity. This could make Waternet's system less influenced by human behavior. In the case of Rijnland, the main pumping stations are not variable speed pumps, leading to discrete pump modes that CQRDNN regression seems to have difficulties with. It could also indicate that pump scheduling can't be forecast based on the current feature data only, leading to high upper quantiles due to the inability to time the actual pumping.

The Day Ahead price forecasting model is based on the methodology presented in (van der Heijden, Palensky, & Abraham, 2021), and we had access to data from 2015 onward. The search space, however, was expanded to include renewable energy data. Due to the unstable markets during the time of simulation, the model was retrained on a weekly basis to be able to learn from the most recent data.




Figure 4. Non-parametric Bayesian Network structures for uncertainty modeling of panel (a) Day Ahead Market prices, incoming discharge, and the water level of the North Sea (adopted from van der Heijden, Palensky, et al. (2022)), and (b) Intraday market prices. N_i represents the i -th hour forecast lead time, DA_i and ID_i represent the DAM and IDM prices for the i -th hour of the day.

3.3. Scenario Generation

To assess decision-making processes in uncertain environments, generating operational scenarios is essential. The Non-Parametric Bayesian Network (NPBN) approach is useful here. It employs univariate parametric distributions based on quantiles from the CQRDNN to depict the marginal distributions of the variable at a timestep, while bi-variate copulae are used to model temporal dependencies as observed in the data. Particularly, in the NPBN, Gaussian copulas are used (In previous literature, Non-Parametric Bayesian Networks are sometimes referred to “Gaussian copula-based” Bayesian Networks (GCBN)) due to their ability to effectively handle a large number of variables by simplifying the process of joint distribution sampling (Mendoza-Lugo & Morales-Nápoles, 2024). This feature makes NPBN effective for scenario generation (van der Heijden, Palensky, et al., 2022) with scenarios that obey both the forecast marginal distribution and the observed auto-correlation. In this work, we apply NPBNs to sample scenarios containing multiple time steps of incoming waterboard discharge, the water level of the North Sea, and the DAM and IDM electricity prices.

In NPBNs, multivariate distributions are characterized by univariate marginals and a copula to represent the dependencies. The joint density of an NPBN with n variables is factorized as:

$$f_{1,\dots,n}(x_1, \dots, x_n) = f_1(x_1) \prod_{i=2}^n f_{i|Pa(i)}(x_i | x_{Pa(i)}), \quad (10)$$

here, $f_{1,\dots,n}$ denotes the joint density of the n variables, f_i represents the marginal distribution of each variable, and $f_{i|j}$ represents the conditional distributions. Each random variable x_i is associated with a node i , and the parent nodes of node i form the set $Pa(i) = i_1, \dots, i_{p(i)}$. The arcs in the NPBN are assigned one-parameter conditional copulas (Joe, 1997), parameterized by Spearman's rank correlations (Hanea et al., 2015). The arc from the parent node i_m to node i is assigned a conditional rank correlation.

In order to generate scenarios with realistic temporal dependency, we first fit parametric distributions, as available in the Scipy (Virtanen et al., 2020) python package, on the quantile forecasts, to be used in an NPBN with pre-defined structure as depicted in Figure 4a. For all variables that are forecast using the CQRDNN, we apply the same BN structure.

Forecasting electricity prices on the IDM requires renewable energy production data that is updated during the day. ENTSO-E's open database ENTSO-E (2018) is of insufficient quality to accurately model the IDM. Therefore, we model IDM price uncertainty using an NPBN that conditionalizes the IDM price to the DAM price and already observed IDM prices, as depicted in Figure 4b. We use the distributions of the relative ID_3 price (e.g., $\frac{ID_3}{DAM}$) to model and infer IDM prices.

3.4. Scenario Reduction

While generating a set of scenarios that fully describes the uncertainty space is invaluable, it is equally valuable to have a practical size, especially in simulation and optimization contexts where computational resources are constrained. Hence, it is beneficial to distil the original scenario set into an optimal subset with statistical properties that best approximates the properties of the original set. The criterion we use for this selection is the “minimal energy distance” between the reduced set and the original set of scenarios. The energy distance was applied since it tends to show better statistical approximation properties than the more widely used Wasserstein distance (Ziel, 2020).

The energy distance quantifies the dissimilarity between two distributions. This metric describes the distances between the elements of X and X^* , and then corrects these with the distances found between members within X and X^* themselves (Székely & Rizzo, 2013; Ziel, 2020).

Mathematically, the energy distance is formulated as.

$$E_p(X, X^*) = 2 \sum_{i \in I} \sum_{j \in I^*} p_i p_j^* d_{ij}^p - \sum_{i \in I} \sum_{j \in I} p_i p_j d_{ij}^p - \sum_{i \in I^*} \sum_{j \in I^*} p_i^* p_j^* d_{ij}^p, \quad (11)$$

$$d_{ij}^p = |\mathbf{x}_i - \mathbf{x}_j|^p, \quad (12)$$

where X denotes the original scenario set with index set I , and scenarios (time series set elements) \mathbf{x}_i for $i \in I$ with respective probabilities p_i . X^* denotes the considered subset with index set $I^* \subset I$, with corresponding scenarios \mathbf{x}^* , and probabilities p^* . The term d_{ij}^p denotes the p -norm distance between two scenarios, with $p = 1$ being our choice for this study. Since the primary set (X) remains unaltered when during scenario selection, the middle term is often neglected. This revised metric is termed the Energy Score. Probabilities p^* of X^* can be optimized for the minimal Energy Score relative to X using the quadratic program

$$\begin{aligned} p^* &= \arg \min_{p^*} E_p(X, X^*), \\ \text{s.t. } &\sum p^* = 1. \end{aligned} \quad (13)$$

To select an optimal subset X^* (i.e., the subset for which the chosen distance metric is minimal), we can compute the Energy distances for all potential subsets where $|I| = N$. This method works for small scenario sets but becomes impractical for larger sets. The forward selection algorithm (Growe-Kuska et al., 2003) was suggested to greedily append scenarios to the subset X^* until the required subset size is attained. This approach is detailed in Algorithm 1, accompanied by an illustrative example founded on the Bernoulli walk, presented in Figure A1.

To represent the uncertainty space sufficiently, we sample 1,000 time series from the BN for each variable. This number was selected a posteriori based on a trade-off between the optimization problem's computational feasibility and the uncertainty representation's completeness.

Algorithm 1. Forward Selection(X, I, N)

```

1  $X^* \leftarrow \emptyset$ 
2  $I^* \leftarrow \emptyset$ 
  while  $|I^*| < N$  do
    // Execute the forward selection algorithm, adding scenarios to the
    // subset until reaching N.
3    $i^* \leftarrow \arg \min_{i \in I \setminus I^*} \text{Distance}(X, X^* \cup \mathbf{x}_i)$  // Acquire the scenario index
    // yielding the smallest Energy Score.
4    $I^* \leftarrow I^* \cup i^*$  // Update the optimal subset's index set.
5    $X^* \leftarrow X^* \cup \mathbf{x}_{i^*}$  // Update the optimal subset.
  end
6  $p^* \leftarrow \text{Probabilities}(X^*)$  // Extract the scenario probabilities of the optimal
  // subset either through the optimal redistribution rule or using the
  // quadratic program defined in Equation (13).
7 return  $X^*, I^*, p^*$ 
```

3.5. Scenario Tree Reduction

Building on the concept of scenario reduction, scenario trees further simplify the task by offering a structured representation. The core principle here is to model the inherent growth of uncertainty over forecast horizons. To transition from scenario subsets to tree representations, we utilize a method inspired by Growe-Kuska et al. (2003). However, we search for the optimal tree shape and use the energy distance (Ziel, 2020) instead of the Wasserstein distance to calculate distances. The subsequent trees model the progression of uncertainty over time, as depicted in Figure A2.

The Backward Tree Reduction method depicted in Figure A2, while effective, has its limitations. It relies on predefined parameters like node locations (Φ) and the maximum children count per node (Θ), while fixing these parameters can yield suboptimal trees. To circumvent this limitation, a Genetic Algorithm (GA) is introduced to dynamically adjust these parameters in order to optimize tree structures for a given scenario set by solving the optimization problem

$$\begin{aligned} \Phi^*, \Theta^* = \arg \min_{\Phi, \Theta} \text{Distance}(\tilde{X}, X^*) \\ \text{s.t. } \sum_{x \in \tilde{X}} \frac{|x|}{N \cdot H} \end{aligned} \quad (14)$$

where we solve optimal split locations (Φ^*) and amount of nodes (Θ^*) by minimizing the distance between tree \tilde{X} and original subset X^* . We define tree complexity as the number of time-discretized variables necessary to describe the scenario set (i.e., the sum of the length of all sub-scenario time series in each node of the tree), and constrain the tree complexity to a fraction, β , of the original complexity defined by subset size N and time series length H . In this work, we apply a β value of 0.5 to constraint tree complexity reduction with a minimum of 50%, which was selected a posteriori based on the speed at which the GA could find feasible solutions.

4. Risk-Aware Optimal Control

In this section, we describe the control model and methods that we propose to optimize the gate and pump schedule for IJmuiden with regard for uncertainty. We start with an introduction to stochastic MPC and our justification for the use of ER constraints (Section 4.1), continued by the water system constraints (Section 4.2) and the multi-market trading objective (Section 4.3). In our internal model, we have the pump- and gate discharge from IJmuiden (q_p, q_g) as decision variables, the water-level of the North Sea (h_{nzk}) as system state, two binary auxiliary variables (z_p, z_g) to indicate pump- and gate discharge possibility, and an auxiliary variable is used as slack variable for water level constraint violations (z_{wl}). Similarly, an auxiliary variable (H) is used to transform the quadratic-linear term in the energy consumption (Equation 24) of the pumping station to a bilinear term. A detailed overview of the control variables can be found in Appendix B1, and the full MPC control problem is formulated in Appendix B2.

4.1. Conditional-Value-at-Risk Constraints

Stochastic MPC is a powerful extension to Deterministic MPC for control problems in the presence of uncertainties (Mesbah, 2016). It leverages mathematical optimization to make optimal control decisions based on probabilistic information about system dynamics and disturbances. The key idea behind stochastic MPC is to formulate the control problem as a stochastic optimization problem, where the objective is to find a control policy that minimizes an expected cost while considering the constraints and uncertainties in the system. In addition to handling uncertainties in a single time step, stochastic MPC can be extended to multi-stage scenarios, where the uncertainty evolves over the prediction horizon. This extension enables the controller to make decisions that take into account the changing nature of uncertainties over time. By making use of real-time updated information and probabilistic forecasts in the optimization, stochastic MPC can help water systems cope with changes in the climate or catchment, and mitigate impacts of extreme hydrological events (Castelletti et al., 2008).

Robust optimization techniques play a significant role in stochastic MPC by addressing uncertainties and maintaining system performance under worst-case scenarios. Scenario-robust optimization aims to find control

policies that can withstand uncertainties by optimizing over a range of possible scenarios or uncertainty realizations. By considering the most unfavorable outcomes in the scenario set, scenario-robust optimization provides assurances regarding system stability and performance, even when uncertainties deviate from their nominal distributions. In stochastic MPC, scenario-robust optimization techniques are integrated into the optimization framework, enabling controllers to account for uncertainties while minimizing the potential impact of worst-case scenarios on system behavior. By explicitly considering worst-case scenarios and their associated constraints or objectives, scenario-robust optimization provides a reliable control framework. This approach is particularly valuable in safety-critical systems or applications where deviations from nominal conditions can have severe consequences, such as the NZK-ARK where constraint violation can lead to flooding.

Robustness in MPC can extend to both constraints and objectives. In scenario-robust constraint formulation, constraint violation will not be tolerated in the whole range of uncertainty realizations. Scenario-robust optimization ensures that the system remains within predefined safety bounds under all considered scenarios, providing a reliable control strategy. Although in most cases, subjecting to constraints even in the most unlikely considered scenarios can be considered conservative. The authors hypothesize that this robustness in control has a price when put in context of DR.

We propose the use of CVaR-inspired constraints for managing risk by controlled relaxing of robustness in decision-making, while minimizing the expected value of the energy cost by trading over multiple markets. Unlike traditional measures like chance constraints that focus solely on a probability threshold or variance constraints that constrain a variance around the expectation, CVaR considers the size of the tail of the distribution beyond the threshold. By capturing the tail of the risk, CVaR provides a more comprehensive measure of the risk associated with the system's outcomes while being the tightest convex approximation of the chance constraint (Venkatasubramanian et al., 2020). The constraint can be formulated as a linear constraint with a slack variable, while being a coherent risk measure (Conejo et al., 2010). Consider the optimization problem.

$$\min_x \mathbb{E}[C(x, \mathcal{W})] \quad (15)$$

$$\text{s.t. } \text{CVaR}_\alpha[C(x, \mathcal{W})], \quad (16)$$

where x represents the control variable, and \mathcal{W} represents the set of possible uncertainty realizations. The cost function $C(x, \mathcal{W})$ captures the system's cost under control variable x and uncertainty set \mathcal{W} . The CVaR_α constraint ensures that the Conditional Value at Risk of the cost function remains below a certain threshold γ with confidence level α . This constraint can be discretely approximated by introducing a slack variable and two linear constraints (Krokhmal et al., 2003)

$$\zeta + (1 - \alpha)^{-1} \sum_{w \in \mathcal{W}} p[w] \cdot z[w] \leq \gamma, \quad (17)$$

$$z[w] \geq C[x, w] - \zeta \quad \forall \quad w \in \mathcal{W}, \quad (18)$$

where ζ represents the VaR, α the confidence level, w an uncertainty realization in set \mathcal{W} with probability set p , γ the CVaR upper bound, C the cost function, and z the newly introduced slack variable representing the cost exceeding the VaR.

Operational water resources management, however, generally requires the risk measure to reflect the risk of water level bound exceedance and, therefore, the tail should be calculated with respect to a fixed threshold. As a fixed threshold would not always equal the VaR, we call this reformulation Exceedance Risk (ER) constraints. This constraint imposes a limit on the risk associated with the constraint violation (i.e., violation of the upper bound on the water level in the canal), ensuring that the system operates within acceptable bounds with confidence level α . By incorporating ER_α constraints, the optimization algorithm ensures that the system satisfies the constraint with predefined statistical confidence. As γ approaches the threshold and α approaches 1, the constraint converges toward a scenario-robust constraint where no constraint violation is allowed in any of the considered uncertainty realizations.

4.2. Water System Constraints

In this section, we describe the constraint formulation of the stochastic MPC based on the already published deterministic model (van der Heijden, Lugt, et al., 2022). The constraints describe the water balance on the linear reservoir model (Equation 19a), and constraints on gate discharge formulated as big-M constraints (Equations 19c, 19d) using binary indicator variable z_g , which is 1 when the gate discharge is actuated, and gate discharge curve (Equation 19e). The physical constraints for pump discharge are similarly formulated with big-M constraints (Equations 19f, 19g) using binary indicator variable z_p and the pump discharge curve (Equation 19h), the lower bound constraint on the water level (Equation 19i), and a quadratic constraint on a new variable to make the pump energy use in the objective bi-linear (Equation 19j).

$$h_{nzk}[t, w] = h_{nzk}[t - 1, w] + (q_{in}[t - 1, w] - q_g[t - 1, w] - q_p[t - 1, w]) \cdot \frac{\Delta t}{A_{nzk}} \quad \forall \quad t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (19a)$$

$$dh[t, w] := h_{nzk}[t, w] - h_{NS}[t, w], \quad (19b)$$

$$dh[t, w] - dh_g^- + (1 - z_g[t, w]) \cdot M_g \geq 0 \quad \forall \quad t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (19c)$$

$$dh[t, w] - dh_g^- - z_g[t, w] \cdot M_g \leq 0 \quad \forall \quad t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (19d)$$

$$q_g[t, w] \leq (a_g \cdot dh[t, w] + b_g) \cdot z_g[t, w] \quad \forall \quad t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (19e)$$

$$-dh[t, w] - dh_p^- + (1 - z_p[t, w]) \cdot M_p \geq 0 \quad \forall \quad t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (19f)$$

$$-dh[t, w] - dh_p^- - z_p[t, w] \cdot M_p \leq 0 \quad \forall \quad t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (19g)$$

$$q_p[t, w] \leq z_p[t, w] \cdot \sum_{i=1}^6 (a_p[i] \cdot -dh[t, w] + b_p[i]) \quad \forall \quad t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (19h)$$

$$h_{nzk}[t, w] \geq h^- \quad \forall \quad t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (19i)$$

$$H[t, w] = (h_{NS}[t - 1, w] - h_{nzk}[t - 1, w])^2 \quad \forall \quad t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (19j)$$

with time steps $t \in [0, T]$, prediction horizon T , sea water level uncertainty realization set \mathcal{W}_h , incoming discharge uncertainty realization set \mathcal{W}_q , their exhaustively combined set $\mathcal{W}_{q \times h} = (\mathcal{W}_q \times \mathcal{W}_h)$, the water level of the NZK-ARK $h_{nzk,t}^w$ with probabilities $p_q \cdot p_h$, the outgoing gate discharge scenarios $q_{g,t}^w$ with probabilities $p_q \cdot p_h$, the outgoing pump discharge scenarios $q_{p,t}^w$ with probabilities $p_q \cdot p_h$, incoming discharge scenarios $q_{in,t}^w \in \mathcal{W}_q$ with probabilities p_q , and the North Sea water level scenarios $h_t^w \in \mathcal{W}_h$ with probabilities p_h .

The upper bound constraint on the water level is formulated as a ER_α constraint (Equations 20a and 20b) or as scenario-robust constraint (Equation 21). The ER_α constraints are formulated to represent the expectation of the water level exceeding the given bound; we formulate linear ER constraints inspired by the CVaR formulation proposed in Krokmal et al. (2003), described in Equations 17 and 18. By introducing slack variable $z_{wl} \in [0, \infty)$ for all scenarios and time steps, the constraint can be efficiently formulated as.

$$h^+ + (1 - \alpha)^{-1} \cdot \sum_{w \in \mathcal{W}^*} z_{wl}[t, w] \cdot p[w] \leq \gamma \quad \forall \quad t \in T, \quad (20a)$$

$$z_{wl}[t, w] \geq h_{nzk}[t, w] - h^+ \quad \forall \quad t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (20b)$$

where h^+ is the water level threshold after which we consider violations for which we use the currently applied target water level ($-0.4 \text{ m} + \text{NAP}$), z_{wl} is the slack variable measuring the upper bound violation of the water level in the NZK for each uncertainty realization w and timestep $t \in T$, α the confidence level, and γ the upper bound on the acceptable ER_α expressed in water level of the NZK ($\text{m} + \text{NAP}$) (i.e., the expectation of the scenarios that violate h^+ with confidence level α).

In this formulation, the slack variable z_{wl} represents the water level exceedance of a scenario and timestep with respect to h^+ . The sum of the probability-weighted exceedances over all scenarios is the expected water level violation at each timestep, making the formulation a discrete approximation of the conditional expectation of water level bound violation, given that the water level in that scenario is higher than h^+ .

As α approaches 1, the ER constraint becomes stricter, and the allowable exceedance above h^+ approaches zero. This means the expectation of violations becomes tighter and increasingly constrained. The constraints in Equation 20 become scenario-robust (i.e., allow no constraint violation in any of the considered uncertainty realizations) and are reformulated as

$$h_{nzk}[t, w] \leq h^+ \quad \forall \quad t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (21)$$

where h^+ is applied as upper bound constraint on the water level over all timesteps, taking all uncertainty realizations into considerations. The continuous decision-variables in this problem are water level h_{nzk} , gate discharge q_g , pumped discharge q_p , slack variable describing water level violations z_{wl} and the quadratic pump head H . Two binary variables, z_g and z_p are present to describe pump- and gate opportunities due to the water level variations between the NZK-ARK and the North Sea.

4.3. Multi-Market Trading Objective

To create a cost-effective pump schedule, we minimize the expected energy cost over all possible control actions under the combined uncertainty of waterboard discharge, North Sea water level, and possible Intraday and Day Ahead market prices.

$$J_{IDM} := \min_{q_g, q_p} \mathbb{E}[\Delta E_{t, w_1} c^{ID}[t, w_2]], \quad \forall t \in T_{IDM}, w_1 \in \mathcal{W}_{q \times h}, w_2 \in \mathcal{W}_{idm|dam}, \quad (22)$$

$$J_{DAM} := \min_{q_g, q_p} \mathbb{E}[E_{t, w_1} c^{DA}[t, w_3]], \quad \forall t \in T_{DAM}, w_1 \in \mathcal{W}_{q \times h}, w_3 \in \mathcal{W}_{dam}, \quad (23)$$

where.

$$E_{t, w} := a_p q_p[t, w]^2 + b_p H[t, w] q_p[t, w] + c_p q_p[t, w] - dh[t, w] \Delta t, \quad (24)$$

$$\Delta E_{t, w} := E_{t, w} - E_t^{\text{bid}}, \quad (25)$$

$$E_t^{\text{bid}} := \mathbb{E}_w[E_{t, w}], \quad (26)$$

with the set of possible DAM price realizations \mathcal{W}_{dam} and the set of possible intraday price realizations that is conditioned on the observed DAM prices of that day $\mathcal{W}_{idm|dam}$. The variable $\Delta E_{t, w_1}$ represents the difference between the actual energy used at the time of a given scenario E_{t, w_1} and the energy bought on the Day Ahead Market E_t^{bid} , where the bid is fixed to the expected energy use for a given hour. It is used with the intraday price c_{t, w_2}^{ID} for intraday trading via the expected cost in J_{IDM} . The energy consumption E_{t, w_1} and the Day Ahead price c_{t, w_2}^{DA} are used for day-ahead trading in J_{DAM} .

In our closed-loop simulation framework, we solve a single optimization problem over a rolling 48-hr control horizon at each decision step. The objective incorporates both J_{DAM} and J_{IDM} , reflecting the need to continuously prepare for future Day Ahead bids while managing ongoing Intraday adjustments.

- *Continuous DAM preparation:* At any given simulation time, we consider a 48-hr horizon that extends beyond the next Day Ahead bidding moment. The J_{DAM} objective is thus always active, guiding the selection of a baseline energy purchase schedule for the upcoming day. This ensures that we continuously prepare for the next DAM bid, even if it occurs several hours into the future.
- *At the DAM bidding hour (e.g., 11:00):* When the DAM bidding moment arrives, we finalize the Day Ahead bid based on the decisions formed in the previous optimization step. This action fixes E_t^{bid} for the next 24-hr

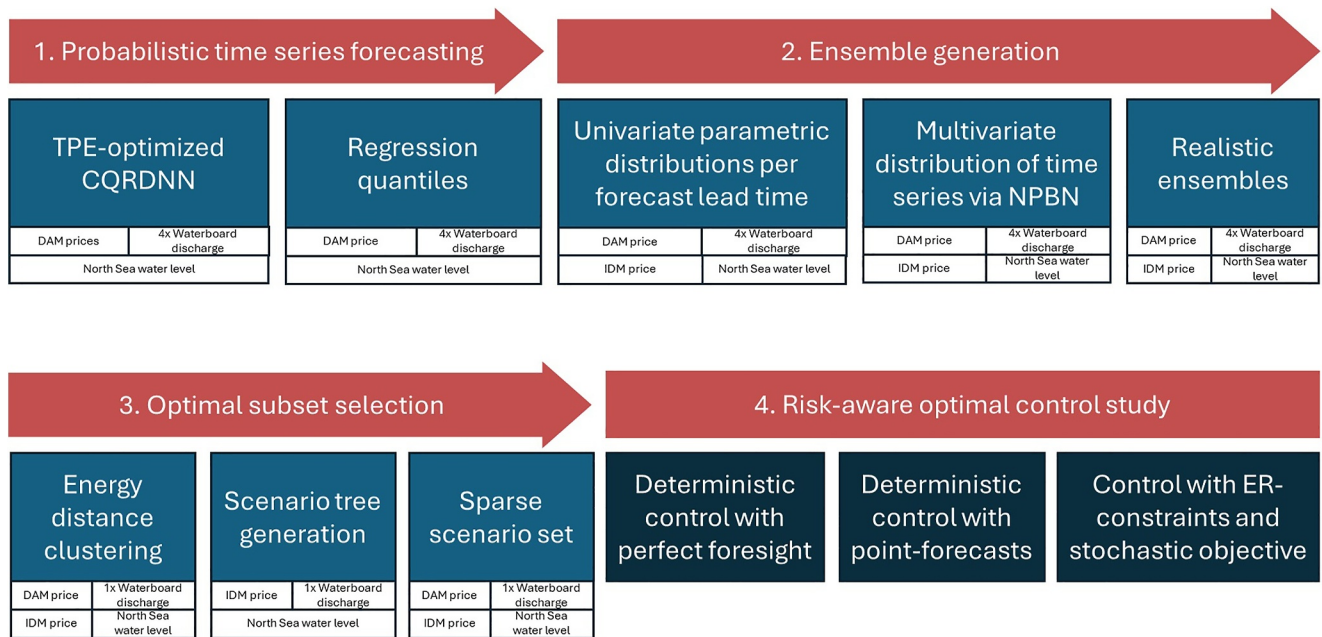


Figure 5. The applied workflow in this manuscript, describing what steps were taking to model operational uncertainty for each variable.

period (from 00:00 to 23:00 of the following day) by bidding the expected energy use over all scenarios. With this bid established, we have a clear reference against which to measure Intraday deviations.

- *Intraday adjustments after the DAM bid:* Once the DAM bid is set at 11:00 and E_t^{bid} is fixed, J_{IDM} becomes relevant for the next day's hours to correct deviations from the fixed DAM schedule. As time progresses and the horizon rolls forward, the DAM bid that was just made remains fixed, while the optimization problem continues to consider both J_{DAM} for future bids and J_{IDM} for current-day Intraday trading. This leads to a shrinking Intraday adjustment window until the next DAM bidding opportunity arrives, at which point a new baseline can be established and the cycle repeats.

5. Results and Discussion

In this section, we describe the results of the proposed framework applied to the case study area; the NZK-ARK. We simulate historical multi-market participation with the real water system and electricity price data, where the optimal control of the NZK-ARK system is simulated under uncertainty in a receding horizon fashion over multiple years and months. We first describe the probabilistic forecasting of incoming discharge from four local water authorities into the canal, the water level of the North Sea, and the DAM prices using the methodology described in Sections 3.1 and 3.2. We then describe how scenarios are generated and reduced to optimal subsets using the method described in Sections 3.3, 3.4 and 3.5. Finally, we perform closed-loop simulations for several months in 2019, 2020, and 2021 in Section 5.3. We analyze the difference between the deterministic (perfect forecasts and deterministic 'point' forecast), and stochastic MPC. We vary the stochastic MPC with different numbers of scenarios and risk-acceptance settings, including ER and scenario-robust constraints. Figure 5 depicts the workflow that is applied in this manuscript. It describes the steps that are taken to model operational uncertainty in each source, and what study is performed to evaluate performance.

5.1. Probabilistic Forecasting and Scenario Generation

Figure 6a shows example operational forecasts for the considered sources of uncertainty.

The model forecasting the water level of the North Sea seems to perform well in stable conditions. A model was trained with perfect foresight for the wind forecast since historic meteorological forecasts are not easily accessible, and fully optimizing for the best model is out of the scope of this work. For the North Sea water level model, we used 2016 and 2017 data as training set, 2018 data as validation set, and 2019–2021 as test set.

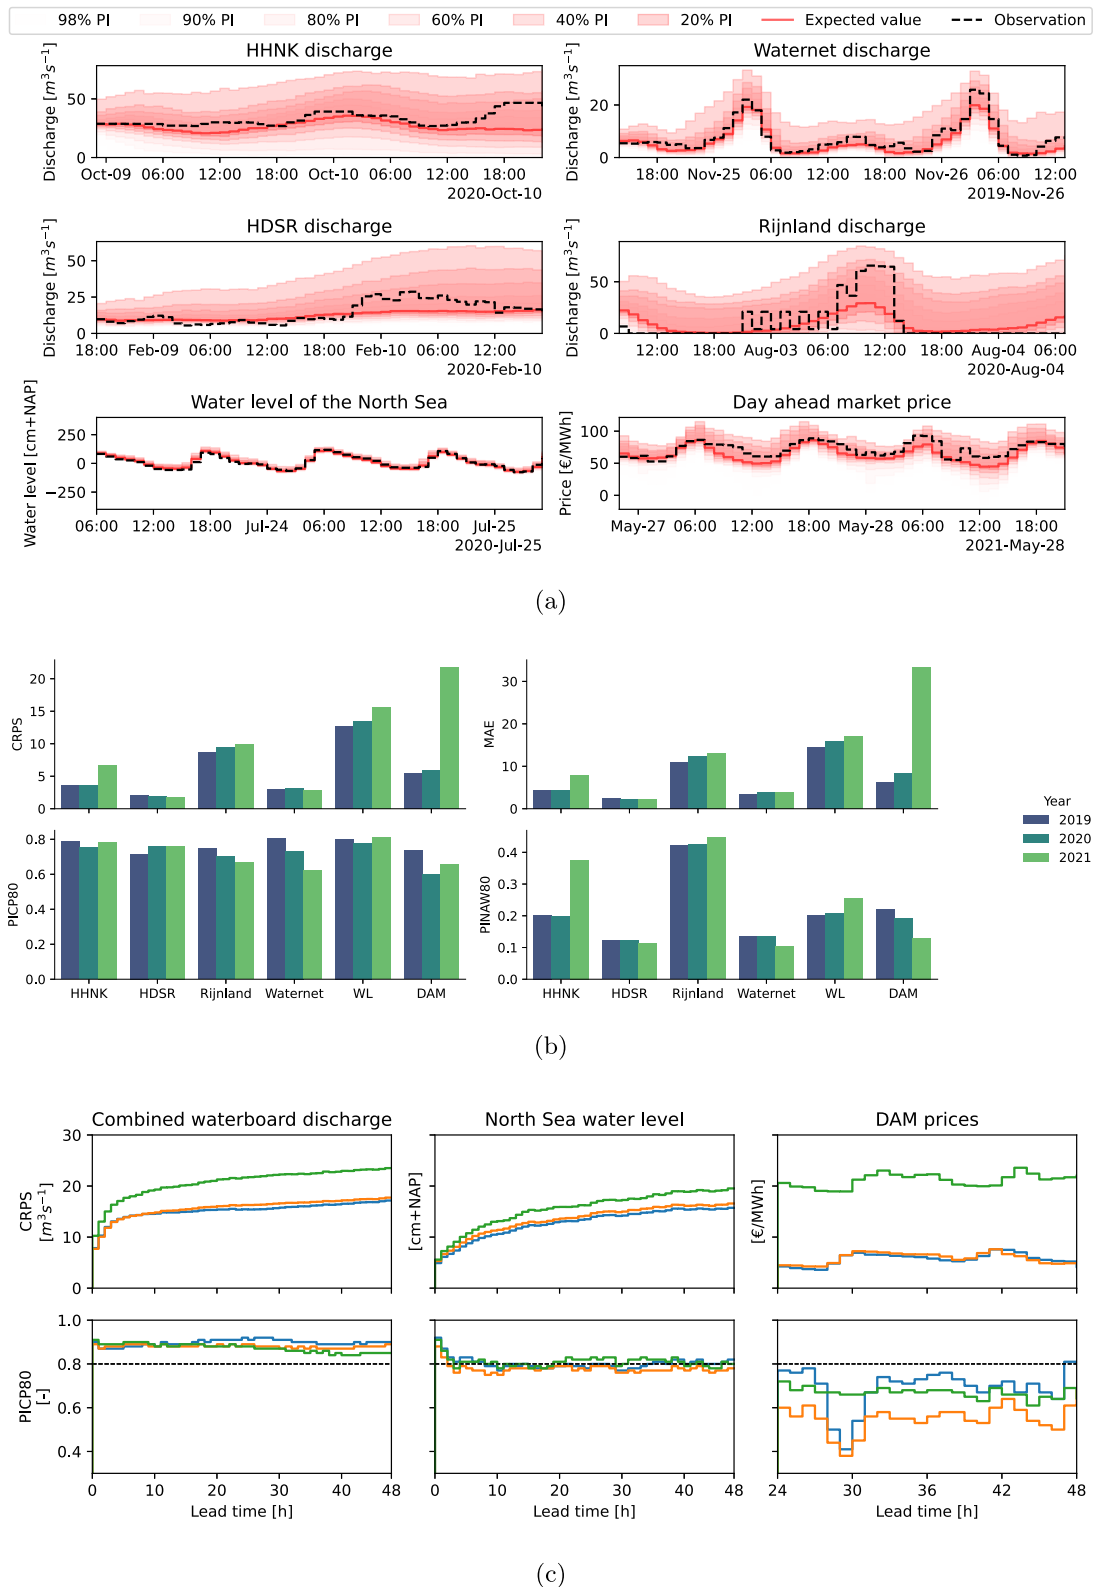


Figure 6. Operational forecast examples (a) discharge from the four local water authorities (HHNK, Rijnland, HDSR, and Waternet), the water level of the North Sea, and the DAM electricity price. Performance metrics of the forecasting models (b), where the metrics are averaged over the 48 hr forecast lead times for all years in the receding horizon simulation. Performance metrics over the lead time (c) per year in the test set. The water levels are shown in m + NAP, which is the standard reference for water levels in the Netherlands, equivalent to Amsterdam Ordnance Datum.




Figure 7. The energy-distance optimal scenario subset of size 10, and the constructed scenario fan and scenario tree with 50% complexity reduction for North Sea water level at a single simulation timestep.

Figure 6b shows some averaged performance metrics over the lead time of the models. The CRPS shows the Continuous Ranked Probability Score, which is a combined score for all quantiles in the forecast and is expressed in the unit of the forecast variable. The MAE is the Mean Absolute Error calculated with the 50th percentile of the forecasts. The PICp80 stands for the Prediction Interval Coverage Percentage at the 80% confidence interval, representing how many observations were between the 90th and 10th percentile. The PINAW80 stands for Prediction Interval Normalized Average Width at the 80% confidence interval. Some models show some performance decrease over the length of the test sets, which would be especially visible in the decrease of PICP and therefore the validity of confidence intervals.

The models could be optimized further in order to improve performance but seem to perform well on average. Every model can be optimized for performance, especially if performance evaluation is purely metric-based. We test model goodness based on operational performance, to be measured in constraint violation and energy cost - compared to perfect foresight. In this work, we focus on the probabilistic control framework and closed-loop simulation testing; the aim is not the exhaustive optimization of each single model based on classical time-series forecast metrics.

All models show some overconfidence in the quantiles, having PICPs lower than 0.8, indicating too narrow prediction intervals. The waterboard discharge model performance seem relatively stable, with a notable decrease in PICP for Waternet. This can be explained by the fact that Waternet decides on their pumping schedule informed by DAM prices, leading to more unpredictable behavior as prices become more volatile. The water level forecast of the North Sea shows an increasing CRPS and MAE over the years, possibly indicating a lack of training data or an in-stationary environment. A severe decrease in accuracy can be seen for the DAM model in 2020 and 2021, when the energy crisis was in effect. The MAE and CRPS increase, but the PICP shows the quantiles are less valid. A high observed price spread can explain the decrease in PINAW. Every model can be optimized for performance, especially if performance evaluation is purely metric-based. We test model goodness based on operational performance, measured in constraint violation and energy cost - compared to perfect foresight. In this work, we focus on the probabilistic control framework and closed-loop simulation testing; the aim is not the exhaustive optimization of each single model based on classical time-series forecast metrics.

On each simulation time step, we apply the CQRDNN models to forecast quantiles, fit a distribution, and sample scenarios with the NPNB approach (Section 3.3). After that, we apply Algorithm 1 to select an optimal subset with new weights. To transform the scenario set into a fan, we cluster the root node to a single scenario. We select a root node size of 3 hours, which corresponds with the time when the IDM has the most liquidity. To reduce computational complexity, we also consider scenario trees. A scenario tree can significantly reduce the amount of decision variables necessary to describe the optimization problem. We apply a GA as described in Section 3.5 to reduce the selected subset to a scenario tree and constrain a minimum 50% reduction in tree complexity. Figure 7 shows an example reduced scenario subset, a scenario fan, and a scenario tree, for the water level of the North Sea.

5.2. Stochastic MPC

To investigate the effect of the variations in risk-acceptance on the water level, we explore the output of the MPC's internal model. Within the internal MPC model, there is a pump discharge, gate discharge, and water-level decision variable for each timestep and all scenario combinations. However, since we need a single implementable output for control, a root node of 3h is maintained where the values of the energy-distance optimal scenario are used, leading to a deterministic pump- and gate schedule for the coming 3 hours.

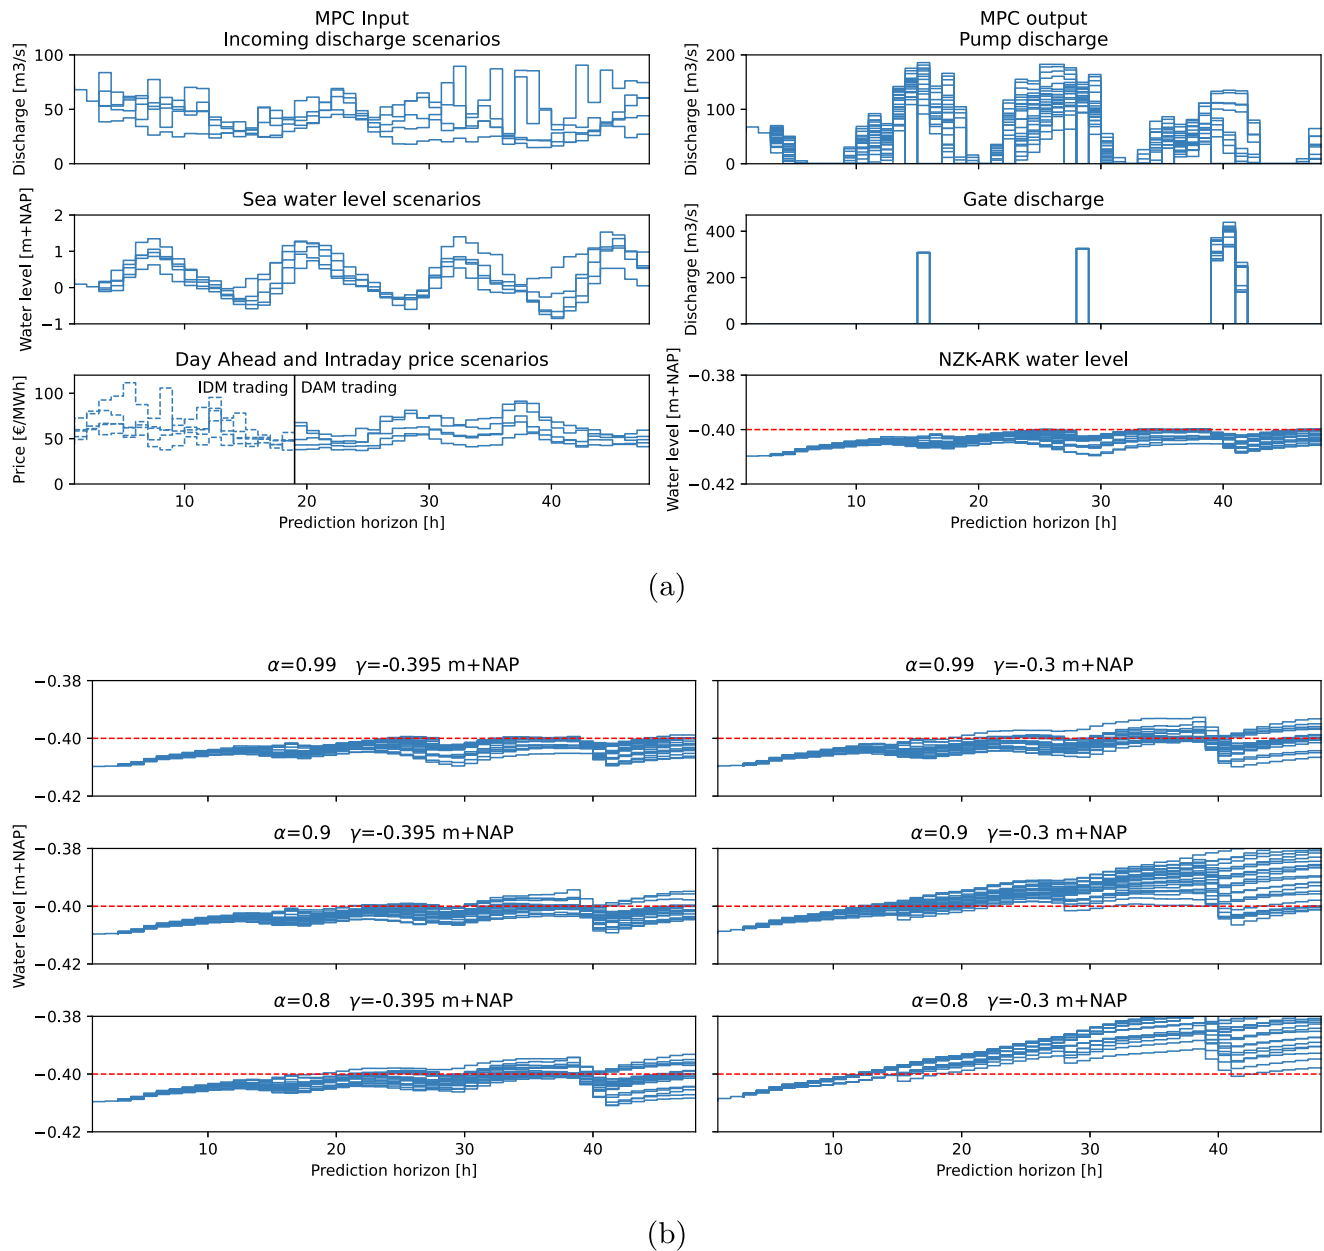


Figure 8. (a) The MPC model input and output from the internal model with constraints that are robust toward the worst case in the reduced set, and (b) the output water level from MPCs with different ER α and γ constraint settings and the same input. The water levels are shown in m + NAP, which is the standard reference for water levels in the Netherlands, equivalent to Amsterdam Ordnance Datum.

Figure 8b shows the water level for a single control timestep of the simulation with receding horizon implementation. We can see how, as the ER α and γ change, the output water level changes, effectively translating ER parameters into operational consequences.

5.3. Closed-Loop Simulation Testing

We perform closed-loop simulations of energy market participation for January, April, July, and October for 2019 and 2020, and January for the year 2021 (due to data availability) at hourly simulation and control time steps, with a control horizon of 2 days. The months are simulated separately due to the framework's runtime resulting from the computational demands of the optimization problem. The simulations were run on Snellius, the Dutch National Supercomputer (Surf Cooperation, 2023), and a single month of simulation still took 22.5 hr to run on




Figure 9. Relative yearly cost of the simulated months compared to the simulations with a perfect forecast. The left figure shows the relative cost for the simulations with ER-constrained water levels, while the right figure shows the relative cost of the simulations with scenario-robust control. Scenario-robust with one scenario is the same as a deterministic problem and is depicted as “Det.”. Red indicates that the simulated costs were higher than the cost for the simulation with perfect forecasts, while blue indicates the cost were lower.

average. The MPC problem was solved using Gurobi (Gurobi Optimization, LLC, 2023) with the relative MIPGap set to 0.01, the absolute MIPGap to €1, the NonConvex parameter set to 2, and a time limit of 15 min if a solution was found, else the optimization would continue until the first feasible solution was found. An overview of the control and simulation settings can be found in Appendix C2.

We perform a sensitivity analysis on the amount and shape of the scenarios, where we consider one (e.g., a deterministic forecast with deterministic constraints), 3, and 5 scenarios per source of uncertainty in either an ensemble with fan or tree shape, and where all scenarios are selected based on the energy distance subset selection method described in Algorithm 1. We then compare a scenario-robust approach with several ER settings ($\alpha \in \{0.99, 0.9, 0.8\}$, $\gamma \in \{-0.395, -0.3\}$), and we compare all simulated scenarios and energy costs with a simulation using a deterministic perfect forecast. Figure 9 shows the relative yearly cost of the simulations of market participation compared to the perfect forecast.

Figure 9 shows a clear value in a probabilistic approach compared to a deterministic (i.e., scenario-robust with 1 scenario) approach, especially in 2019. All considered probabilistic approaches result in lower energy costs than the deterministic approach. However, large variations in relative performance can be seen over the years. The lowest performance increase was seen in 2020 when the COVID pandemic hit. One possible explanation is that the difficulty of forecasting electricity prices increased in 2020 and 2021, leading to suboptimal operational uncertainty estimation.

To better understand the observed variability in performance over the years, we compare the forecasting accuracy of key variables using the CRPS and the PICP of the 80% prediction interval. Figure 6c presents the CRPS and PICP for combined discharge, North Sea water levels, and DAM prices for the years 2019–2021. The results indicate that forecasting performance for DAM prices deteriorated significantly in 2020, with a CRPS of 5.91 and a PICP of 0.60, compared to 2019 (CRPS of 5.47, PICP of 0.74). This decline likely impacted the relative performance of the stochastic methods, as accurate price forecasts are crucial for optimal DAM bidding strategies. Interestingly, these results indicate that operational performance is more dependent on the correctness of the outer quantiles (PICP of the 80% prediction interval) than the accuracy of all combined quantiles (CRPS).

In contrast, water level forecasts show relatively stable performance across all years, with PICP values consistently close to the target. These results highlight the importance of forecasting accuracy for DAM prices, as their variability is directly linked to the profitability of multi-market trading. The difficulty of forecasting DAM prices in 2020 and 2021 may explain the lower cost savings observed that year.

When comparing ER approaches with scenario-robust approaches, the figure shows that, on average, ER approaches result in lower energy cost than scenario-robust approaches. Less risk-tolerant ER approaches seem to perform worse than the scenario-robust approaches, which a lack of on-time intraday trading could explain. As the future probability of a high water level becomes more tolerable, the scenario will become more likely at some point, activating the upper-bound constraint. Suppose this happens too close to the execution time. In that case, more energy has to be traded on the IDM since the DAM is already closed, limiting the economic potential of multi-market trading.




Figure 10. Histogram and KDEs of the simulated and observed water levels in the NZK-ARK. From left to right, the figure shows the distribution of the simulated water level with an optimization using deterministic forecasts, scenario-robust optimization with 5 scenario forecasts, three optimization problems with ER-constrained water level, and the observed water level in the NZK-ARK. The water levels are shown in m + NAP, which is the standard reference for water levels in the Netherlands, equivalent to Amsterdam Ordnance Datum.

Increasing the size of the scenario set does not always lead to better results. When the MPC problem is hard to solve, the optimal solution would not be reached within the cut-off time. A more scarce scenario-tree representation, therefore, generally performs better with 5 scenarios. In 2021, which only consists of January, the difference between a fan- or tree uncertainty representation is the largest. January 2021 was a wet month, with an average simulation runtime for the 5-scenario fan approach of 58hr, while the average simulation runtime for the 5-scenario tree approach was 27hr. This is consistent with our hypothesis that time constraints can affect the optimal scenario set size for uncertainty representation.

In one simulation with actual forecasts, we observe a lower cost than in the perfect forecast scenario. This can be explained by a forecast error at advantageous time, leading to more intraday trading when IDM prices were advantageous. Even with the perfect forecast, IDM prices are only known after DAM closure, leading to the absence of speculation on IDM prices. This could lead to more profitable energy on the IDM, even though the initial DAM bid was suboptimal. The portfolio could be further optimized if IDM and DAM prices were available for the optimization problem at the same time. However, that would increase the complexity of the optimization problem and might not be practically feasible using the same problem formulation.

In general, it can be seen that the benefits of the proposed stochastic method vary over the years. The year 2019, when electricity prices were easiest to forecast, seemed to have the most optimal results with up to 44 p.p. cost savings compared to the deterministic forecast, and even contained a simulation with lower cost than the perfect forecast scenario. In 2020, when the COVID pandemic just started, the profitability of the proposed strategy decreased, but the stochastic method still allows 10 p.p. cost savings compared to a deterministic approach. We also see that the value of taking risks decreases in 2020, while that benefit is present in 2019 and 2021 when the electricity markets were relatively calm compared to 2020.

Figure 10 shows the probability distributions of the observed water levels of the NZK-ARK and a set of simulated water levels. The figure shows that a deterministic forecast can lead to a sharp distribution around the target water level. This is possibly caused by forecast inaccuracies, causing the control system to correct when water levels rise. This is consistent with the higher energy cost since flexibility can't be exploited well. The difference between scenario-robust and ER constraints can be seen in the figure. While both strategies have a similar distribution shape, ER constraints allow for a shift of the distribution and a slight flattening of the shape. When the complexity of the problem is lower, it seems like there is slightly more constraint violation occurring, as depicted by the distribution of the simulated water level with a tree-shaped uncertainty representation. None of the strategies will lead to a significant added risk compared with the observed water levels. Arguably, there can be a large difference between modeled and actual water levels, but we believe the main conclusions stand. When comparing the distributions with the regime depicted in Figure 1b, we conclude that a lot of flexibility is left unexploited and that having less conservative control could well benefit both the energy transition as Rijkswaterstaat, the operator of the IJmuiden pumping station.

6. Summary and Conclusions

This study has presented a framework for incorporating risk into the operational management of water resources, focusing on the Netherlands' unique water management challenges. Our approach centered around the NZK-ARK system and effectively integrates probabilistic forecasting and optimal control strategies in a real-world setting, integrating uncertainty from various sources.

In our proposed framework, we apply a Combined Quantile Regression Deep Neural Network (CQRDNN) to estimate operational uncertainty, while a Non-Parametric Bayesian Network (BN) was applied to model auto-temporal relationships of forecast variables. Our methodology for scenario reduction, employing the Energy distance metric, has shown that it is feasible to distil a comprehensive set of scenarios into an optimal subset, while closed-loop performance does seem to be partially dependent on the number of scenarios in the subset.

By formulating Stochastic Model Predictive Control (MPC) with Exceedance Risk (ER) constraints, we provide an advancement over traditional deterministic MPC approaches. ER constraints allow for nuanced risk management by modulating the control response based on the severity and likelihood of upper-bound violations of water levels. This novel approach not only manages risks more effectively but also provides a mechanism for adjusting risk preferences dynamically, illustrated by our ability to warp and shift the distribution of the simulated water levels as shown in Figure 10. Besides leading to a nuanced description of risk, the constraints are formulated linearly, giving considerable computational advantages compared to integer programming.

We use 'historical simulations' with real water system and electricity market data to demonstrate the advantages of multi-market participation under uncertainty. The optimal control of the NZK-ARK system, simulated over several years, revealed that stochastic MPC could yield substantial energy cost savings compared to deterministic strategies. These savings highlighted the effectiveness of our approach in adapting to variabilities in climate and market conditions, thereby providing a more resilient and economically efficient operational framework that can potentially be applied in many fields. Furthermore, our results show the utility of employing larger scenario sets for enhancing decision-making optimality, although we noted the practical constraints of optimization under tight time frames. The use of scenario trees are shown to be an effective strategy under these circumstances, facilitating more manageable computations while minimizing compromising the quality of outcomes.

An important aspect of this study is its scalability and adaptability to other water management systems. The computational complexity of the current formulation arises primarily from the complementarity constraints on pump and gate discharge, modeled using a big-M approach. These constraints are specific to the NZK-ARK system but may not be necessary for other canal systems, particularly in local or regional contexts within the Netherlands, where a positive head can typically be assured. Without these constraints, the computational burden is significantly reduced, making it feasible to scale the framework to larger networks of canals or systems with delayed flows. The modular and flexible design of the framework also makes it suitable for application in other domains, such as hydropower and reservoir management. These systems present different types of uncertainty over longer relevant horizons, such as inflow variability or market dynamics, but future work could be directed at tailoring the framework's underlying structure to address these challenges effectively. Additionally, we have successfully applied a similar approach to battery energy storage systems (van der Heijden et al., 2023), further demonstrating the framework's versatility and its potential for broader applications across diverse domains.

In examining water levels, our simulated results stayed within acceptable ranges, indicating the operational viability of our approach. While there are observable differences between simulated and actual water levels, our findings also show that adjusting the ER parameters effectively modifies risk profiles, offering avenues for further refinement in future applications.

While this study did not explicitly benchmark the ML-based probabilistic forecasts or the MPC approach against alternative methods, this was a deliberate decision to maintain a focused scope. The goal was to validate the framework as an integrated, modular approach for risk-aware decision-making. Benchmarking individual components, such as the forecasting methods or control strategies, could provide additional insights but would have significantly broadened the study's scope. Future research could address these comparisons, evaluating the forecasts against ensemble weather models or assessing the MPC approach against other state-of-the-art control methods. The modular nature of the framework ensures that specific components can be benchmarked or replaced without altering the overarching structure, allowing for continuous improvement and validation across different contexts.

To conclude, this research advances the field of operational water resource management by delivering a robust, adaptable, and economically viable framework capable of managing the uncertainty and variability inherent in weather and electricity price conditions. This framework not only supports Rijkswaterstaat in optimizing energy costs but also enhances the safety and reliability of water and energy systems. While our study focused on the Dutch context, the methodologies and findings are applicable globally, offering valuable insights for regions

facing similar challenges. The proposed methods also extend beyond water resources, possibly benefiting any application facing similar uncertainty and risk.

Appendix A: Scenario Subset Selection Example: The Bernoulli Walk

A1. Scenario Reduction and Weight Redistribution

Figure A1




Figure A1. Scenario reduction and weight redistribution of the Bernoulli walk using the Forward Selection algorithm and the Energy distance metric.

A2. Scenario Tree Reduction

Figure A2




Figure A2. Tree construction for the Bernoulli walk using the Energy distance.

Appendix B: Optimization Problem Details

B1. Overview of Control Variables

Table B1

Table B1

Table With Control Variables, Units, Mathematical Notation, Variable Type, and Their Domains

Definition	Unit of measurement	Symbol	Type	Domain
Pump discharge	$\text{m}^3 \text{ s}^{-1}$	q_p	Decision variable	$[0, 260]$
Gate discharge	$\text{m}^3 \text{ s}^{-1}$	q_g	Decision variable	$[0, 500]$
NZK-ARK water level	$\text{m} + \text{NAP}$	h_{nzk}	State variable	$[-0.6, -0.4]$
Pump discharge indicator	-	z_p	Auxiliary variable	$\{0, 1\}$
Gate discharge indicator	-	z_g	Auxiliary variable	$\{0, 1\}$
Upper water level bound violation	m	z_{wl}	Auxiliary variable	$[0, 1]$
Bilinear reformulation variable	m^2	H	Auxiliary variable	$[0, 20]$
Timestep size	s	Δt	Constant	3,600

B2. Full MPC Control Problem

The full MPC control problem is defined below. With.

$$J_{IDM} := \min_{q_g, q_p} \mathbb{E} |\Delta E_{t, w_1} c^{ID} [t, w_2]|, \quad \forall t \in T_{IDM}, w_1 \in \mathcal{W}_{q \times h}, w_2 \in \mathcal{W}_{idm|dam}, \quad (\text{B1})$$

$$J_{DAM} := \min_{q_g, q_p} \mathbb{E} |E_{t, w_1} c^{DA} [t, w_3]|, \quad \forall t \in T_{DAM}, w_1 \in \mathcal{W}_{q \times h}, w_3 \in \mathcal{W}_{dam}, \quad (\text{B2})$$

subject to.

$$h_{nzk}[t, w] = h_{nzk}[t-1, w] + (q_{in}[t-1, w] - q_g[t-1, w] - q_p[t-1, w]) \cdot \frac{\Delta t}{A_{nzk}} \quad \forall t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (\text{B3a})$$

$$dh[t, w] := h_{nzk}[t, w] - h_{NS}[t, w], \quad (\text{B3b})$$

$$dh[t, w] - dh_g^- + (1 - z_g[t, w]) \cdot M_g \geq 0 \quad \forall t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (\text{B3c})$$

$$dh[t, w] - dh_g^- - z_g[t, w] \cdot M_g \leq 0 \quad \forall t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (\text{B3d})$$

$$q_g[t, w] \leq (a_g \cdot dh[t, w] + b_g) \cdot z_g[t, w] \quad \forall t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (\text{B3e})$$

$$-dh[t, w] - dh_p^- + (1 - z_p[t, w]) \cdot M_p \geq 0 \quad \forall t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (\text{B3f})$$

$$-dh[t, w] - dh_p^- - z_p[t, w] \cdot M_p \leq 0 \quad \forall t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (\text{B3g})$$

$$q_p[t, w] \leq z_p[t, w] \cdot \sum_{i=1}^6 (a_p[i] \cdot -dh[t, w] + b_p[i]) \quad \forall t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (\text{B3h})$$

$$h_{nzk}[t, w] \geq h^- \quad \forall t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (\text{B3i})$$

$$H[t, w] = (h_{NS}[t-1, w] - h_{nzk}[t-1, w])^2 \quad \forall t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (\text{B3j})$$

$$h^+ + (1 - \alpha)^{-1} \cdot \sum_{w \in \mathcal{W}^*} z_{wl}[t, w] \cdot p[w] \leq \gamma \quad \forall t \in T, \quad (\text{B3k})$$

$$z_{wl}[t, w] \geq h_{nzk}[t, w] - h^+ \quad \forall t \in T \text{ and } w \in \mathcal{W}_{q \times h}, \quad (\text{B3l})$$

where.

$$E_{t, w} := a_p q_p[t, w]^2 + b_p H[t, w] q_p[t, w] + c_p q_p[t, w] - dh[t, w] \Delta t, \quad (\text{B3m})$$

$$\Delta E_{t,w} := E_{t,w} - E_t^{\text{bid}}, \quad (\text{B3n})$$

$$E_t^{\text{bid}} := \mathbb{E}_w |E_{t,w}|, \quad (\text{B3o})$$

where decision variables, q_g (gate discharge), q_p (pump discharge), system state h_{nzk} (water level of the NZK-ARK), and auxiliary variables z_g (binary indicator for gate discharge opportunities), z_p (binary indicator for pump discharge opportunities), H (quadratic constraint for to make dh^2 -term bi-linear in the objective), and z_{wl} (the slack variable used for the ER constraint) are optimized to minimize the expected energy cost of pumping under uncertainty.

Appendix C: Experiment Settings

C1. Feature and Hyperparameter Search Space

Table C1

Table C1

Feature and Hyperparameter Search Space Dimensions for the Optimization of the CQRDNN in Several Forecasting Tasks

Hyperparameters	Waterboard discharge	North sea water level	Day ahead market price
N layers	Discharge lag range	Lagged water level range	Lagged price range
Neurons per layer	Lagged wind data range	Lagged wind data range	Include prices of d-7
Dropout rate	Rolling window size	Rolling water level window size	Load data
Regularization	Include precipitation window	Include hourly wind data	Lagged load range
Batch normalization	Include evaporation window	Include 10 min wind data	Include load of d-7
Random seed	Include temperature window	Include wind direction data	Include load forecast for targets
	Include discharge window	Include day of the year	Include onshore wind generation data
	Include wind forecast	Include hour of the day	Include offshore wind generation data
	Include precipitation forecast		Include solar generation data
	Include sea water level ^a		Market integration

^aSea water level is only included in the feature search space for waterboards that have the option to discharge into the North Sea directly.

C2. Simulation and Control Settings

Table C2

Table C2

Table With Simulation and Control Settings

Simulation settings		Solver settings	
Simulation length	9 × 1 month	Solver	Gurobi v10.0.0.2
Simulation timestep	1 hr	Relative MIPGap	0.01
Control horizon	48 hr	Absolute MIPGap	€1
Control timestep	1 hr	Feasibility tolerance	1e-6
		NonConvex parameter	2
		Time limit	15 min
		Solution limit ^a	1

^aSolution limit is only applied if a time out was reached while no feasible solution was found yet, forcing the optimization to continue until a feasible solution was found.

Data Availability Statement

The data that was used for this work is available on van der Heijden (2024).

Acknowledgments

This work was supported by TKI Water technologie and Optimal control of open canal system for demand response: Enabling and preparing for the energy transition.

References

- Bergstra, J., Yamins, D., & Cox, D. D. (2013). Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *International conference on machine learning* (Vol. 28, pp. 115–123). PMLR.
- Bessa, R. J., Möhrlein, C., Fundel, V., Siefert, M., Browell, J., Haglund El Gaidi, S., et al. (2017). Towards improved understanding of the applicability of uncertainty forecasts in the electric power industry. *Energies*, 10(9), 1402. <https://doi.org/10.3390/en10091402>
- Castelletti, A., Ficchi, A., Cominola, A., Segovia, P., Giuliani, M., Wu, W., et al. (2023). Model predictive control of water resources systems: A review and research agenda. *Annual Reviews in Control*, 55, 442–465. <https://doi.org/10.1016/j.arcontrol.2023.03.013>
- Castelletti, A., Pianosi, F., & Soncini-Sessa, R. (2008). Integration, participation and optimal control in water resources planning and management. *Applied Mathematics and Computation*, 206(1), 21–33. <https://doi.org/10.1016/j.amc.2007.09.069>
- Conejo, A. J., Carrión, M., & Morales, J. M. (2010). *Decision making under uncertainty in electricity markets*. Springer. <https://doi.org/10.1007/978-1-4419-7421-1>
- Dang Doan, M., Giselsson, P., Keviczky, T., De Schutter, B., & Rantzer, A. (2013). A distributed accelerated gradient algorithm for distributed model predictive control of a hydro power valley. *Control Engineering Practice*, 21(11), 1594–1605. <https://doi.org/10.1016/j.conengprac.2013.06.012>
- De Jong, J., Hassel, A., Jansen, J., Egenhofer, C., & Xu, Z. (2017). Improving the market for flexibility in the electricity sector. Retrieved from <https://www.ceps.eu/ceps-publications/improving-market-flexibility-electricity-sector/>
- ENTSO-E. (2018). Entso-e transparency platform. Retrieved from <https://transparency.entsoe.eu/>
- Giuliani, M., Lamontagne, J. R., Reed, P. M., & Castelletti, A. (2021). A state-of-the-art review of optimal reservoir control for managing conflicting demands in a changing world. *Water Resources Research*, 57(12), e2021WR029927. <https://doi.org/10.1029/2021WR029927>
- Goedbloed, A. (2006). Kwaliteitsanalyse beslissingen ondersteunend systeem noordzeekanaal amsterdam-rijnkanaal.
- Growe-Kuska, N., Heitsch, H., & Roemisch, W. (2003). Scenario reduction and scenario tree construction for power management problem. In *2003 IEEE Bologna Power Tech Conference Proceedings* (Vol. 3, p. 7). IEEE. <https://doi.org/10.1109/PTC.2003.1304379>
- Gurobi Optimization, LLC. (2023). Gurobi optimizer reference manual. Retrieved from <https://www.gurobi.com>
- Hanea, A., Morales Napoles, O., & Ababei, D. (2015). Non-parametric Bayesian networks: Improving theory and reviewing applications. *Reliability Engineering & System Safety*, 144, 265–284. <https://doi.org/10.1016/j.res.2015.07.027>
- Horváth, K., van Esch, B., Vreeken, T., Piovesan, T., Talsma, J., & Pothof, I. (2022). Potential of model predictive control of a polder water system including pumps, weirs and gates. *Journal of Process Control*, 119, 128–140. <https://doi.org/10.1016/j.jprocont.2022.10.003>
- Janssen, H. (2017). *Effect selectieve onttrekking IJmuiden op waterbeheer* (Tech. Rep.). Rijkswaterstaat. Retrieved from <https://www.platformparticipatie.nl/selectieve-onttrekking/beroep+1/relevante+documenten+104/HandlerDownloadFiles.aspx?id=1258968>
- Joe, H. (1997). *Multivariate models and multivariate dependence concepts*. CRC Press.
- Jordehi, A. R. (2019). Optimisation of demand response in electric power systems, a review. *Renewable and Sustainable Energy Reviews*, 103(July 2018), 308–319. <https://doi.org/10.1016/j.rser.2018.12.054>
- Koenker, R., & Bassett, G. (1978). Regression quantiles. *Econometrica*, 46(1), 33. <https://doi.org/10.2307/1913643>
- Krokhmal, P., Palmquist, J., & Uryasev, S. (2003). Portfolio optimization with conditional value-at-risk objective and constraints. *Journal of Risk*, 4(2), 43–68. <https://doi.org/10.21314/JOR.2002.057>
- Lago, J., De Ridder, F., & De Schutter, B. (2018). Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms. *Applied Energy*, 221(January), 386–405. <https://doi.org/10.1016/j.apenergy.2018.02.069>
- Maestre, J. M., Raso, L., Van Overloop, P., & De Schutter, B. (2013). Distributed tree-based model predictive control on a drainage water system. *Journal of Hydroinformatics*, 15(2), 335–347. <https://doi.org/10.2166/hydro.2012.125>
- Makridakis, S., Spiliotis, E., Assimakopoulos, V., Chen, Z., Gaba, A., Tsetlin, I., & Winkler, R. L. (2022). The m5 uncertainty competition: Results, findings and conclusions. *International Journal of Forecasting*, 38(4), 1365–1385. <https://doi.org/10.1016/j.ijforecast.2021.10.009>
- Mendoza-Lugo, M. A., & Morales-Nápoles, O. (2024). Mapping hazardous locations on a road network due to extreme gross vehicle weights. *Reliability Engineering & System Safety*, 242, 109698. <https://doi.org/10.1016/j.res.2023.109698>
- Mesbah, A. (2016). Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Systems Magazine*, 36(6), 30–44. <https://doi.org/10.1109/MCS.2016.2602087>
- Pape, C. (2018). The impact of intraday markets on the market value of flexibility — Decomposing effects on profile and the imbalance costs. *Energy Economics*, 76, 186–201. <https://doi.org/10.1016/j.eneco.2018.10.004>
- Pianosi, F., & Soncini-Sessa, R. (2009). Real-time management of a multipurpose water reservoir with a heteroscedastic inflow model. *Water Resources Research*, 45(10), W10430. <https://doi.org/10.1029/2008WR007335>
- Pour, F. K., Segovia, P., Duviella, E., & Puig, V. (2022). A two-layer control architecture for operational management and hydroelectricity production maximization in inland waterways using model predictive control. *Control Engineering Practice*, 124, 105172. <https://doi.org/10.1016/j.conengprac.2022.105172>
- Raso, L., Schwanenberg, D., van de Giesen, N. C., & van Overloop, P. J. (2014). Short-term optimal operation of water systems using ensemble forecasts. *Advances in Water Resources*, 71, 200–208. <https://doi.org/10.1016/j.advwatres.2014.06.009>
- Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2(3), 21–41. <https://doi.org/10.21314/JOR.2000.038>
- Sahu, R. K., & McLaughlin, D. B. (2018). An ensemble optimization framework for coupled design of hydropower contracts and real-time reservoir operating rules. *Water Resources Research*, 54(10), 8401–8419. <https://doi.org/10.1029/2018WR022753>
- Schwabeneder, D., Corinaldesi, C., Lettner, G., & Auer, H. (2021). Business cases of aggregated flexibilities in multiple electricity markets in a european market design. *Energy Conversion and Management*, 230, 113783. <https://doi.org/10.1016/j.enconman.2020.113783>
- Surf Cooperation. (2023). Snellius: The national supercomputer. Retrieved from <https://www.surf.nl/en/services/snellius-the-national-supercomputer>
- Sweeney, C., Bessa, R. J., Browell, J., & Pinson, P. (2020). The future of forecasting for renewable energy. *WIREs Energy and Environment*, 9(2), e365. <https://doi.org/10.1002/wene.365>
- Székely, G. J., & Rizzo, M. L. (2013). Energy statistics: A class of statistics based on distances. *Journal of Statistical Planning and Inference*, 143(8), 1249–1272. <https://doi.org/10.1016/J.JSPI.2013.03.018>

- TenneT, & Gasunie (2022). *Infrastructure outlook 2050 (Tech. Rep.)*. Author.
- van der Heijden, T. (2024). Dataset underlying the PHD thesis “unlocking flexibility: Risk-aware operational water and energy management”. <https://doi.org/10.4121/e747fa10-af31-41b8-b984-79132a1efb0.v1>
- van der Heijden, T., Lago, J., Palensky, P., & Abraham, E. (2021a). Electricity price forecasting in European day ahead markets: A greedy consideration of market integration. *IEEE Access*, 9, 119954–119966. <https://doi.org/10.1109/ACCESS.2021.3108629>
- van der Heijden, T., Lugt, D., van Nooijen, R., Palensky, P., & Abraham, E. (2022a). Multi-market demand response from pump-controlled open canal systems: An economic MPC approach to pump-scheduling. *Journal of Hydroinformatics*, 24(4), 838–855. <https://doi.org/10.2166/hydro.2022.018>
- van der Heijden, T., Palensky, P., & Abraham, E. (2021b). Probabilistic dam price forecasting using a combined quantile regression deep neural network with less-crossing quantiles. In *Iecon 2021 – 47th annual conference of the IEEE industrial electronics society* (pp. 1–6). <https://doi.org/10.1109/IECON48115.2021.9589097>
- van der Heijden, T., Palensky, P., van de Giesen, N., & Abraham, E. (2022b). Day ahead market price scenario generation using a combined quantile regression deep neural network and a non-parametric Bayesian network. In *2022 IEEE international conference on power systems technology (powercon)* (pp. 1–5). <https://doi.org/10.1109/POWERCON53406.2022.9929940>
- van der Heijden, T., Palensky, P., van de Giesen, N., & Abraham, E. (2023). Closed-loop simulation testing of a probabilistic DR framework for day ahead market participation applied to battery energy storage systems. In *2023 IEEE 32nd international symposium on industrial electronics (isie)* (pp. 1–6). <https://doi.org/10.1109/ISIE51358.2023.10227921>
- Venkatasubramanian, J., Rostampour, V., & Keviczky, T. (2020). Stochastic MPC for energy management in smart grids with conditional value at risk as penalty function. In *2020 IEEE PES innovative smart grid technologies Europe (ISGT-Europe)*, (pp. 309–313). IEEE Computer Society. <https://doi.org/10.1109/ISGT-Europe47291.2020.9248769>
- Verhagen, J., van der Zaag, P., & Abraham, E. (2021). Operational planning of WEF infrastructure: Quantifying the value of information sharing and cooperation in the eastern Nile basin. *Environmental Research Letters*, 16(8), 085006. <https://doi.org/10.1088/1748-9326/ac1194>
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., et al. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Zhao, P., Wang, Q. J., Wu, W., & Yang, Q. (2021). Which precipitation forecasts to use? Deterministic versus coarser-resolution ensemble NWP models. *Quarterly Journal of the Royal Meteorological Society*, 147(735), 900–913. <https://doi.org/10.1002/qj.3952>
- Ziel, F. (2020). The energy distance for ensemble and scenario reduction. *Philosophical Transactions of the Royal Society A: Mathematical, Physical & Engineering Sciences*, 379(2202), 20190431. <https://doi.org/10.1098/rsta.2019.0431>