UNIVERSITY OF OKLAHOMA
GRADUATE COLLEGE

PHYSICS-AWARE AND MARKET-RESPONSIVE OPTIMIZATION FOR HYDROPOWER STORAGE: A CONVERGENCE-BASED DECISION SUPPORT FRAMEWORK

A THESIS
SUBMITTED TO THE GRADUATE FACULTY
in partial fulfillment of the requirements for the
Degree of
MASTER OF SCIENCE

By DAVID SOLANO
Norman, Oklahoma
2026

PHYSICS-AWARE AND MARKET-RESPONSIVE OPTIMIZATION FOR HYDROPOWER STORAGE: A CONVERGENCE-BASED DECISION SUPPORT FRAMEWORK

A THESIS APPROVED FOR THE
GALLOGLY COLLEGE OF ENGINEERING
BY THE COMMITTEE CONSISTING OF

Dr. Kash Barker, Chair

Dr. Sina Khanmohammadi

Dr. Talayeh Razzaghi

Contents

Abstract

The increasing penetration of variable renewable energy into modern power grids has fundamentally altered the operational context of hydropower, necessitating a shift from conventional load-peaking strategies toward adaptive, market-responsive dispatch. This thesis develops and evaluates the foundational components of a physics-aware, market-responsive optimization framework for hydropower storage, demonstrated through application to DeGray Dam a 74 MW United States Corps of Engineers (USACE)-operated facility within the Midcontinent Independent System Operator (MISO) South region and marketed by Southwestern Power Administration (SWPA).

The framework integrates two independent analytical components coupled through a composite multi-objective function balancing economic reward maximization against physical risk penalization. The market component employs a walk-forward validated statistical model forecasting daily maximum Locational Marginal Prices (LMP) using publicly available MISO Market Operating Margin (MOM) reports and Energy Information Administration (EIA) commodity price data across a 978-day dataset spanning January 2023 through September 2025. Ridge Regression was selected as the production model after nine-month holdout validation, achieving an aggregate RMSE of 0.284 on the log scale and demonstrating superior robustness in high-volatility regimes. The physical component employs an Hydrologic Engineer Center-Hydrologic Modeling Software (HEC-HMS) simulation calibrated to the Caddo River watershed, translating dispatch schedules into daily reservoir storage trajectories and quantifying risk penalties for energy-in-storage violations.

Two operational scenarios were evaluated: a conservation baseline and a market-dispatch strategy triggering four-hour full-turbine releases when the predicted LMP exceeded the 75th percentile threshold ($65.45/MWh). The market-dispatch scenario generated $1,415,531 in additional revenue on 62 dispatch days at a capture rate of 54.8%, while incurring a $455,506 risk penalty for energy storage of 144 days below the storage threshold of 50%. Under a base-case weighting of $\lambda_1 = 0.60$, market dispatch dominated the conservation baseline by a composite margin of $667,116. Sensitivity analysis identified a crossover at $\lambda_1^* \approx 0.25$, indicating the dispatch strategy is preferred across a wide range of operator risk tolerances.

The results establish a quantitative foundation for convergence-based decision support in federal hydropower operations, with direct applicability to other grid operator facilities and broader relevance to hydropower systems operating in organized energy markets.

1 Introduction

The increasing integration of Variable Renewable Energy (VRE) sources presents significant challenges for traditional power system components, necessitating a re-evaluation of long-standing operational practices, particularly for hydropower. Conventional hydropower operations, often designed for historical conditions to minimize load demand peaking, are no longer optimally aligned with modern grids that feature high VRE penetration [9]. The shift towards a VRE-dominant system requires a new focus on how flexible resources like hydropower can provide balancing services that reduce overall system-wide costs—a concept referred to as adaptive hydro, distinct from the historical fixed hydro focus on peak shaving [9]. This critical transition raises the fundamental question: are the current operational models and objective functions capturing the right economic and system-wide priorities?

In response to the increasing complexity and non-linear relationships within modern energy systems, the use of advanced techniques like Machine Learning (ML) and Deep Learning (DL) has become a major research focus [2, 3]. Studies have explored the application of ML for predictive tasks, such as enhancing stream flow forecasts via Bayesian Model Averaging [15]. Other research has demonstrated the potential of ML over traditional statistical regressions, especially in capturing detailed patterns at the hourly level for maximum energy production, though this work was not explicitly tied to market conditions [6]. Additionally, the implementation of sophisticated ML/DL models faces significant hurdles, including the critical need for large amounts of quality data, a challenge highlighted as a persistent issue in the energy sector [3, 2]. Further concerns include computational requirements, forecast uncertainty, and a lack of interoperability for real-time operational decisions [2, 10].

Beyond pure data-driven forecasting, research has also investigated the optimization of hydropower systems. Some approaches combine physical knowledge with machine learning, such as hybrid models where a Neural Network is used to predict physical property coefficients for a hydrological model [7]. In the realm of pure optimization, studies have applied evolutionary algorithms, which are often adaptations of heuristic approaches, with a primary objective function being to maximize electricity generation [14]. Another key area of development is the direct application of Deep Reinforced Learning (DRL) to find optimal individual strategies for storage units within electricity market models, moving beyond simple generation maximization [4]. Furthermore, understanding the true value of forecast skill is crucial, as the performance increase from forecast-informed decisions is highly dependent on both a dam's design specifications (e.g., storage capacity) and the regional forecast skill [8]. These factors determine the ultimate cost-benefit of incorporating forecasting into operation.

This literature review highlights a clear disconnect: while ML is successfully being applied to both hydrological forecasting [15] and generation maximization [6, 14], and DRL shows promise in market modeling [4], a comprehensive framework is needed to integrate accurate operational details—such as the start-stop costs associated with wearing of hydropower units [1] or fuel storage—into a holistic, market-driven, system-wide optimization objective function [9]. Addressing the persistent challenges of data quality, forecast uncertainty, and the need for a truly "adaptive" dispatch remains paramount for integrating hydropower effectively into the VRE-dominated energy market.

This thesis addresses that gap by developing and evaluating the foundational components of a physics-aware, market-responsive optimization framework, demonstrated through application to the DeGray hydroelectric facility within MISO South.

2 Project Location and Description

The United States Corp of Engineers (USACE) owns and operates the Degray dam and powerhouse which consists of a multipurpose reservoir with project purposes of flood control, power, water supply, recreation, and navigation. A downstream re-regulation dam impounds tailrace flows for continuous water supply and pump-back storage for power. The initial construction of the project began on 29 June 1962. The river diversion was made in May 1966. [17]. The Caddo River has its source in the mountains of Montgomery County, Arkansas, and flows in a southeasterly direction a distance of 78 miles to its junction with Ouachita River at mile 426.0 above the mouth of the Black River and approximately 5.4 miles above Arkadelphia, Arkansas. The Degray Damsite is located about 7.9 miles above the mouth of the Caddo River. The general location of Caddo River and Degray Damsite can be seen in Figure 2.1.

2.1 Operating Conditions

Degray's power is marketed by Southwestern Power Administration (SWPA) as a Power Marketing Agency (PMA) under the Department of Energy (DOE). It is within the MISO footprint and its energy is offered into their market. Degray consists of two generating units with a nameplate capacity of roughly 38 MW each and average annual generation of 47,000 MWh [5].

2.2 Reservoir Storage Allocations

The reservoir's storage is divided into specific pools to manage operational requirements. The specific elevation limits and capacities are detailed in Table 2.1.

2.2.1 Storage Descriptions

- **Flood Control Pool:** Containing 227,200 acre-feet of storage, this pool is designed to control the flood of record (April 1927) and reduce stages on the Ouachita River.

- **Joint Use Pool:** This 393,200 acre-foot zone provides the active storage for hydro-electric power generation and assures a firm water supply yield of 250 million gallons per day (average 387 c.f.s.).

Figure 2.1: Overlay of DeGray Dam indicated by the green triangle and the hashed area indicating the watershed

- **Minimum Pool:** The 261,500 acre-feet below elevation 367.0 serves to maintain the hydraulic head required for power generation and provides dead storage for sediment accumulation.

Table 2.1: Reservoir Pool Allocations

Pool Designation	**Elevation Range** (ft m.s.l.)	**Storage** (acre-feet)	**Primary Function**
Surcharge Pool	423.0 – 447.5	495,100	Dam Safety (Extreme Floods)
Flood Control	408.0 – 423.0	227,200	Storm Runoff Control
Joint Use	367.0 – 408.0	393,200	Power & Water Supply
Minimum Pool	< 367.0	261,500	Power Head & Sediment

Note: The total storage capacity below the spillway crest (El. 423.0) is 881,900 acre-feet. Storage estimates from original Water Control Manual

3 Multi-Objective Function

The goal of the optimization model is to find the operational strategy that yields the highest total weighted score, balancing economic reward incentive with the imperative to manage system risk.

3.1 Overall Objective Function: Maximizing the Weighted Goal

The core of the model is the weighted sum approach, which converts the two competing objectives (reward maximization and risk minimization) into a single objective function **F** to be maximized.

$$\max \mathbf{F} = \lambda_1 F_{\text{Reward}} - \lambda_2 F_{\text{Risk}} \tag{3.1}$$

Key Components and Weights

- **Reward Term (F_{Reward}):** Maximized, as it represents the desired financial outcome.
- **Risk Term (F_{Risk}):** Subtracted, as it represents a penalty (cost) that should be minimized.
- **Weighting Factors (λ_1, λ_2):** These factors determine the trade-off. They must satisfy $\lambda_1 + \lambda_2 = 1$.
  - A higher λ_1 prioritizes reward.
  - A higher λ_2 prioritizes risk avoidance (i.e., respecting constraints).

3.1.1 Economic Efficiency Objective: Maximizing Reward (F_{Reward})

This objective calculates the total net revenue from electricity generation over the dispatch period T.

$$\mathbf{F}_{\text{Reward}} = \sum_{t=1}^{T} (C_{\text{sale},t} P_{\text{gen},t} - C_{\text{o\&m}}) \tag{3.2}$$

Explanation

- **Revenue ($C_{\mathbf{sale},t} P_{\mathbf{gen},t}$):** This encourages generation ($P_{\text{gen},t}$) when the electricity price ($C_{\text{sale},t}$) is highest, maximizing revenue.
- **Operational Cost ($C_{\mathbf{o\&m}}$):** A deduction for the ongoing operational and maintenance cost at each time step.
- **Summation ($\sum_{t=1}^{T}$):** Totals the reward across all time steps, allowing the model to optimize water usage strategically over the entire time horizon T.

3.2 Risk Minimization Objective: Minimizing Penalty ($F_{\mathbf{Risk}}$)

This objective uses a penalty function (or soft constraint) to minimize the cost incurred from violating non-power operational limits or strategic storage levels for capacity considerations.

$$\mathbf{F}_{\text{Risk}} = F_{\text{Flood}} + F_{\text{Quality}} + F_{\text{Storage}} \tag{3.3}$$

Penalty Function Rationale

- Violating a constraint (e.g., exceeding $H_{\max}$) results in a penalty cost, which the optimizer seeks to minimize.
- The use of large penalty coefficients (α, β) ensures that violations, while technically allowed, are extremely expensive, forcing the solution to remain within the safe, feasible operational range.

3.2.1 Flood Control Risk Penalty ($F_{\mathbf{Flood}}$)

This term penalizes the system when the reservoir level exceeds the maximum safe elevation ($H_{\max}$), which is critical for mitigating flood risk. This equation is designated as a theoretical to account for competing uses of the reservoir, in most dams flood control will supersede all other operations so normally not violated for competing uses.

$$\mathbf{F}_{\text{Flood}} = \sum_{t=1}^{T} \alpha_{\text{H}} \left[\max\left(0, H_t - H_{\max}\right)\right]^2 \tag{3.4}$$

Explanation

- **Violation Term ($\max(0, \dots)$):** Only a positive violation (when $H_t > H_{\max}$) is counted. If the reservoir is below the limit, the term is zero.
- **Quadratic Penalty ($\dots^2$):** The squared term applies a **non-linear, steeply increasing penalty**. This means that small violations are moderately penalized, but large violations become prohibitively expensive, ensuring strong adherence to the flood limit.

3.2.2 Water Quality/Environmental Flow Risk Penalty ($F_{\textbf{Quality}}$)

This term ensures that minimum operational and environmental needs are met by penalizing low reservoir levels and insufficient downstream flows.

$$\mathbf{F}_{\text{Quality}} = \sum_{t=1}^{T} \left(\alpha_{\text{L}} \left[\max\left(0, H_{\min} - H_t\right) \right]^2 + \beta_{\text{Q}} \left[\max\left(0, Q_{\min} - Q_t\right) \right]^2 \right) \tag{3.5}$$

Explanation

- **Low-Level Penalty (First Term):** Penalizes $H_t < H_{\min}$. The minimum level is necessary to ensure reliable water intake, preserve water quality, and maintain a minimal pool for environmental stability.
- **Environmental Flow Penalty (Second Term):** Penalizes $Q_t < Q_{\min}$. This enforces the requirement for a minimum release to sustain the downstream aquatic ecosystem.
- **Penalty Coefficients ($\alpha_{\mathbf{L}}, \beta_{\mathbf{Q}}$):** These are tuned to reflect the severity of violating low-level and minimum-flow requirements.

3.2.3 Storage Risk Penalty ($F_{\textbf{Storage}}$)

This term is corresponding to the relative energy in storage percentages compared to the entire power pool. Often this threshold is based on longer term capacity strategies and confidence in hydrologic cycles for storage replenishment. It is subjective to the risk tolerance of the operator and operational criteria of the organization, but for this study penalties were considered below 50 percent.

$$\mathbf{F}_{\text{Storage}} = \sum_{t=1}^{T} \alpha_{\text{L}} \left[\max\left(0, EIS_{50\%} - EIS_{\text{current}}\right) \right]^2 \tag{3.6}$$

Explanation

- **Violation Term ($\max(0, \dots)$):** Only a positive violation (when $EIS_{\text{current}} < EIS_{50\%}$) is counted. If the reservoir is above the limit, the term is zero.
- **Quadratic Penalty ($\dots^2$):** The squared term applies a **non-linear, steeply increasing penalty**. This means that small violations are moderately penalized, but large violations become prohibitively expensive, ensuring strong adherence to the storage limit.

3.3 Integration with the Predictive and Physical Models

The complete dispatch framework relies on two distinct, non-interacting models a statistical/machine learning (ML) model and a physical/rule-based model which concurrently provide the two primary components of the multi-objective function. The final optimal dispatch is determined by a central optimization solver that seeks to maximize the unified objective **F**.

3.3.1 The Statistical/Machine Learning Model and $F_{\textbf{Reward}}$

The statistical or machine learning model is purely concerned with maximizing economic efficiency based on market dynamics. Its analysis provides the necessary data and an initial optimal dispatch trajectory for the reward component of the objective function.

$$\mathbf{F}_{\text{Reward}} \xleftarrow{\text{Informed by}} \text{Statistical/ML Predictive Model} \tag{3.7}$$

Mechanism

- **Data Independence:** This model operates solely on market data, including forecasts for electricity prices ($C_{\text{sale},t}$) and load, independent of any real-time hydrological or physical system state.
- **Objective:** The model's analysis essentially defines the highest possible reward achievable by optimizing the power generation schedule ($P_{\text{gen},t}$) against the forecast prices.

3.3.2 The Physical Model and $F_{\textbf{Risk}}$

The physical model is a rule-based or simulation model focused on hydrological and regulatory compliance. It provides the crucial penalty structure for constraint violations that define the risk component of the objective function.

$$\mathbf{F}_{\text{Risk}} \xleftarrow{\text{Informed by}} \text{Physical/Rule-Based Model} \tag{3.8}$$

Mechanism

- **Data Independence:** This model operates solely on physical and regulatory inputs, such as forecast inflows, maximum flood control levels ($H_{\max}$), and minimum environmental flows ($Q_{\min}$), independent of market price forecasts.
- **Constraint Structure:** The physical model establishes the functional form and the high penalty coefficients (α, β) that quantify the cost of violating the key physical variables (H_t and Q_t).

Symbol	Definition
$\mathbf{F}$	Overall weighted objective function to be maximized.
F_{Reward}	Objective to Maximize Economic Efficiency (Reward/Revenue).
F_{Risk}	Objective to Minimize Risk (Penalty of Constraint Violations).
λ_1, λ_2	Weighting factors (where $\lambda_1 + \lambda_2 = 1$ and $\lambda_1, \lambda_2 \geq 0$).
T	Total number of time steps (e.g., 24 hours).
t	Index for the time step, $t = 1, 2, \ldots, T$.
$C_{\text{sale},t}$	Price of electricity at time t (\$/MWh).
$P_{\text{gen},t}$	Power generated by the plant at time t (MW).
$C_{\text{o\&m}}$	Operational and maintenance cost for the time step (\$).
H_t	Reservoir water level at time t (m).
$H_{\max}$	Maximum allowable reservoir level (m) for flood control.
$H_{\min}$	Minimum allowable reservoir level (m) for water quality/intake.
Q_t	Total water released/discharged at time t (m^3/s).
$Q_{\min}$	Minimum required environmental flow/release (m^3/s).
$\alpha_{\text{H}}, \alpha_{\text{L}}$	Penalty coefficients for violating reservoir level limits ($\$/\text{m}^2$).
β_{Q}	Penalty coefficient for violating minimum release limit ($\$/(\text{m}^3/\text{s})^2$).

Table 3.1: Definition of Variables in the Multi-Objective Function

4 Data Sources

Data was collected across many publicly available sources. Collection needs were different for the physical model and the market model.

4.1 Market Model Inputs

MISO market data is available within the market reports portal. There are many market reports available, so the focus was on reports that are available with predictions and information that is multi-day that would influence day-ahead decisions. Therefore, the chosen reports were Market Operating Margin Report [11] and historical LMPs [12]. In addition to the MISO data fuel spot prices were collected from Energy Information Administration (EIA) for natural gas and petroleum [20],[20]. For extra context, Degray historical energy was collected, however it was only publicly available at a monthly aggregate. [19].

4.2 Physical Model Data Inputs

In order to have a hydrologic model that would be able to forecast inflow, a watershed model was created in (HEC-HMS) [16]. This model was created off terrain data from United States Geological Survey USGS [21]. To recreate the scenarios that would be experienced during the time of actual decision making, it was necessary to compile data that would allow the anticipation of inflows into the reservoir. Therefore, historical gridded precipitation data was necessary to vary run-off conditions. This data came from climate prediction center [13]. For the Hydrologic model to be better calibrated historically observed stream flows was taken from Caddo River Near Caddo Gap, AR [22]. This data was collected with a focus on 2023-2025 to be the same time period as the available market data.

5 Exploratory Data Analysis and Feature Engineering

Prior to model development, a comprehensive exploratory analysis was conducted on the 978-day dataset (January 2023–September 2025) to assess data quality, distributional properties, and temporal dynamics. This phase determined the necessary transformations and feature engineering strategies required to model the MISO South market effectively.

5.1 Data Quality and Missingness

The dataset exhibited high overall completeness (99.8%). A missingness analysis by column type revealed distinct patterns:

- **Target Variables:** The `max_daily_LMP` vector contained 29 missing values (2.97%), primarily concentrated in a specific operational window in April 2024 and the final days of the dataset in September 2025.
- **Multi-Day Operating Margin (MOM) Predictors:** 88 of the daily predictor variables had less than 2% missingness.
- **Imputation Strategy:** Given the low missingness rate, a median imputation strategy was applied to predictor variables to preserve the distribution without introducing synthetic variance. Rows with missing target variables were excluded from the training set.

5.2 Univariate Analysis: Distributional Properties

5.2.1 Target Variable Skewness

A Shapiro-Wilk assessment and skewness calculation confirmed that the target variable, `max_daily_LMP`, followed a highly right-skewed distribution (skewness = 5.11), characteristic of electricity price spikes.

To address this, a logarithmic transformation was tested. As shown in Table C.2, the log-transformation successfully reduced skewness to manageable levels for linear modeling.

Consequently, all subsequent modeling utilized `max_daily_LMP_log` as the dependent variable.

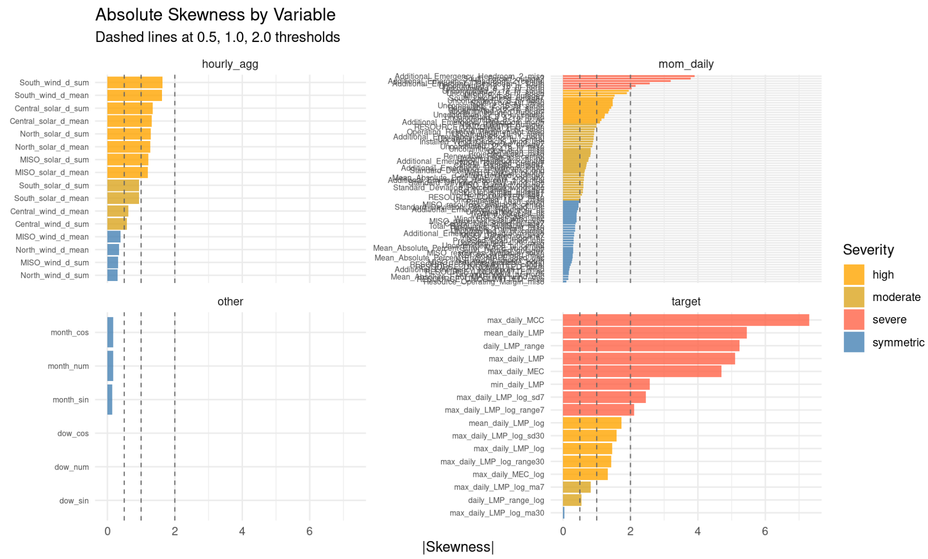

Figure 5.1: Visual representation of the various types of predictors and target variables skewness

5.3 Multivariate Analysis: Collinearity and Redundancy

The dataset contains significant multicollinearity, this is for a multitude of reasons, this dataset was taking into account predictors within the entire MISO, so regional predictors could be very similar to the entire system, also some renewable resources are dominant in one load zone, so their influence on the regional level is still the near the same on the entire system.

5.3.1 Correlation Structure

A pairwise correlation analysis identified 70 pairs of predictors with a Pearson correlation coefficient $|r| > 0.90$.

- **Wind Forecasts:** Extremely high correlation was observed between MISO-wide wind forecasts and regional (North/Central) wind forecasts ($r \approx 1.0$).
- **Load vs. Obligation:** The `Projected_Load` and `Obligation` variables exhibited near-perfect collinearity.

5.3.2 Dimensionality Assessment (PCA)

Principal Component Analysis (PCA) was performed to quantify the redundancy. The analysis revealed that:

- The first 15 components capture 80% of the dataset's variance.

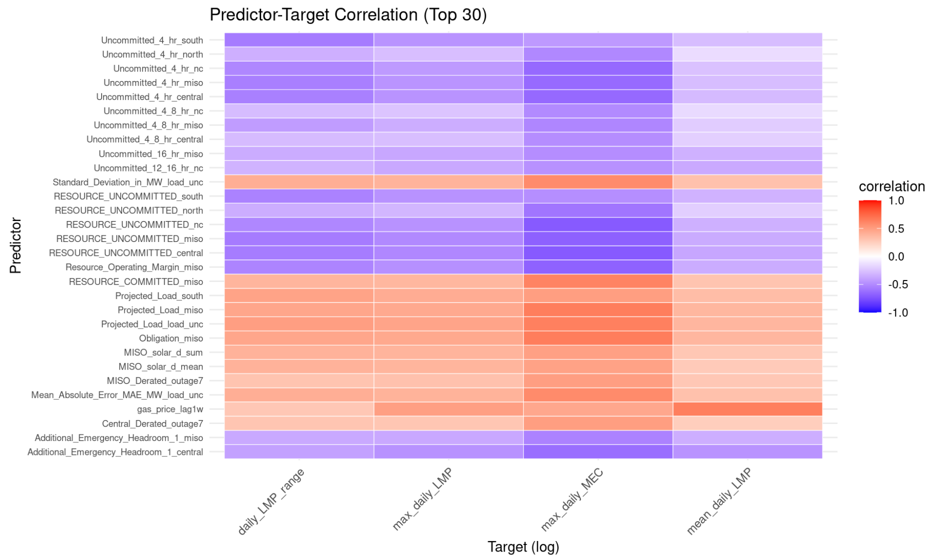

Figure 5.2: Heat correlation map of the top 30 predictors compared to the target variables

- The first 26 components capture 90% of the variance.

Despite this redundancy, a "Strategic Retention" approach was selected, retaining 112 predictors rather than dropping correlated features. This decision relied on the ability of the selected models (Ridge Regression and XGBoost) to handle multicollinearity via regularization and tree-splitting mechanisms, respectively as shown in 5.3 and 5.4.

5.4 Temporal Dynamics and Stationary

5.4.1 Stationarity

An Augmented Dickey-Fuller (ADF) test was performed on the log-transformed target series.

- **ADF Statistic:** -7.14
- **p-value:** 0.01

The null hypothesis of a unit root was rejected, indicating the time series is **stationary**. Although the target variable was determined to be stationary, there were most likely non-stationary predictors that comprised the total training set and validation set as the market tries to minimize volatility the trigger of each volatile event may differ.

5.4.2 Autocorrelation and Seasonality

Autocorrelation Function (ACF) analysis revealed significant lag correlations up to 10 days, confirming the persistence of price regimes see 5.5,5.6 and 5.7.

Seasonal decomposition highlighted:

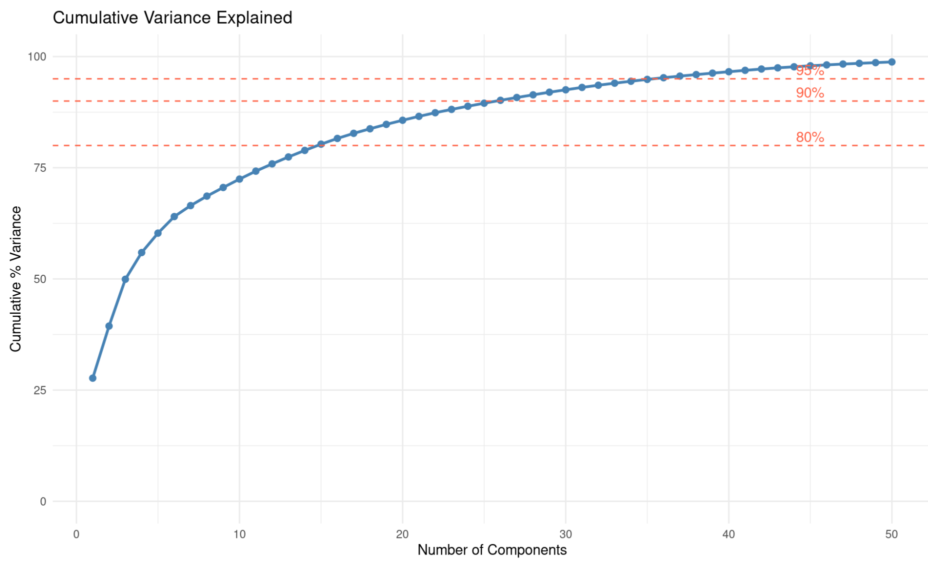

Figure 5.3: Cumulative variance of principal component analysis

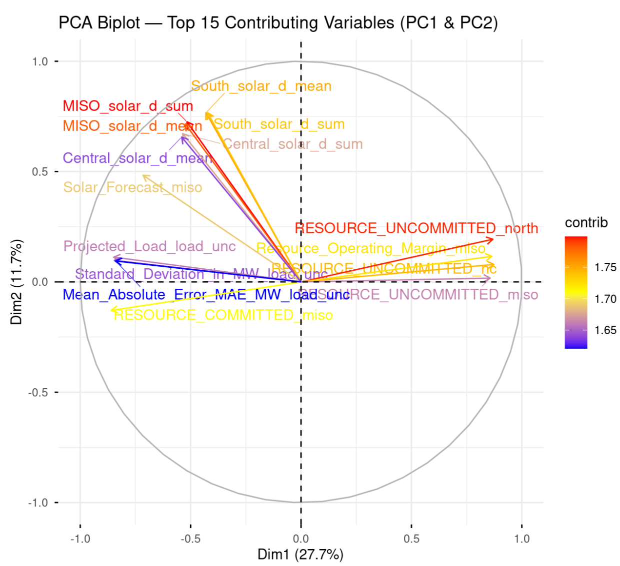

Figure 5.4: Biplot of principal component 1 and 2 showing the contributions of the predictors

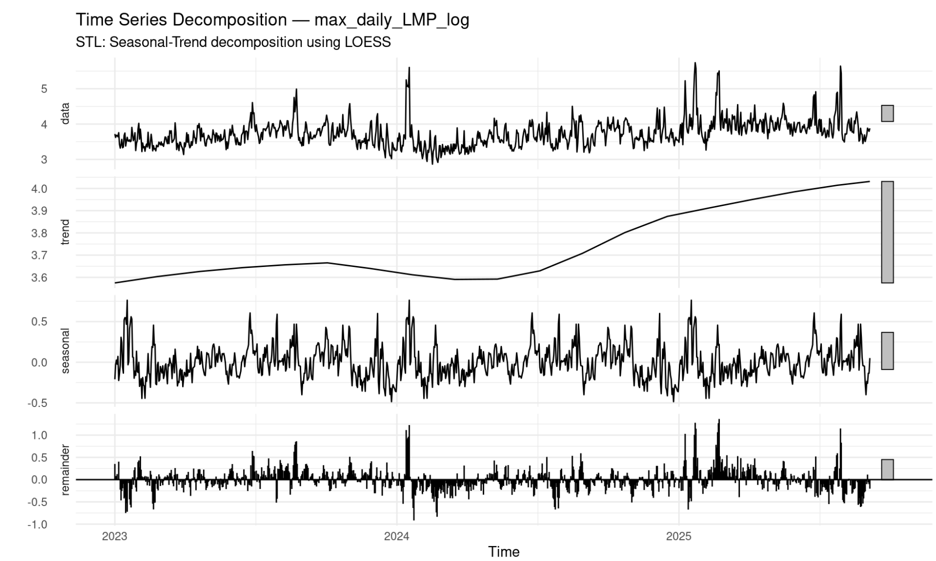

Figure 5.5: Time-series decomposition showing the trend and seasonal component extracted from the observed and then the residuals have decomposition

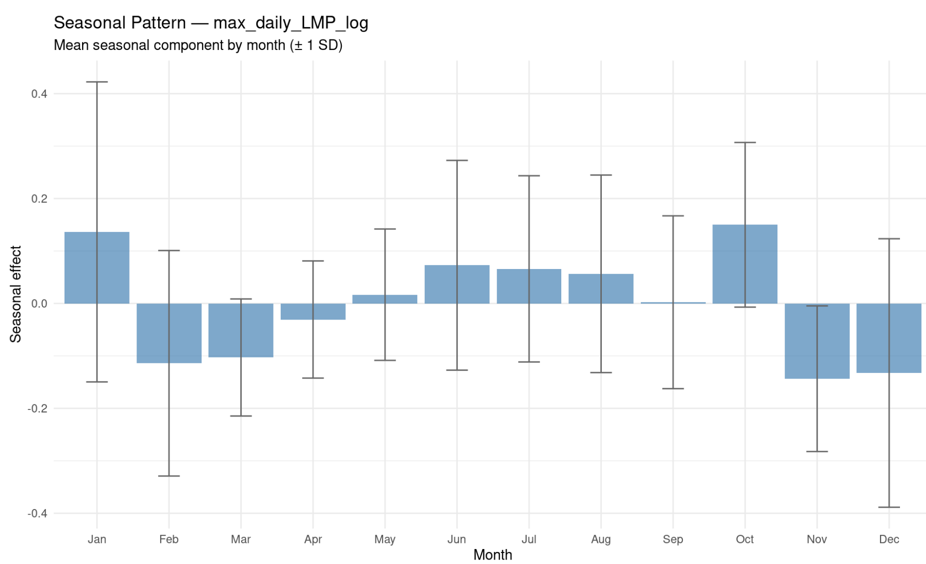

Figure 5.6: Box plot separated by month to the data is centered so 0 indicates the mean of the entire population and the box-plots show how months respective means deviate from that

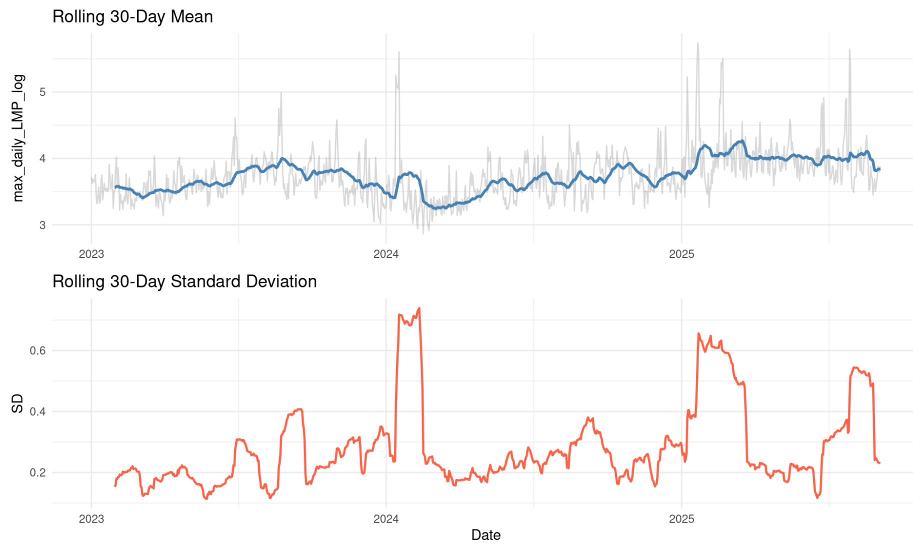

Figure 5.7: Rolling 30-day averages for both mean and standard deviation to better appreciate overall behavior away from daily spikes

- **Weekly Seasonality:** Distinct variations by Day of Week (Mean log-LMP is lowest on Saturdays/Sundays).
- **Monthly Seasonality:** Peaks observed in October and January, this is not-intuitive to the expected summer peak, this can indicate that the net load volatility may be driving prices more than just load obligations considered in isolation.

5.5 Data Preprocessing and Feature Engineering

The data processing pipeline for the MISO South market analysis was designed to ensure temporal integrity and prevent data leakage, particularly for forecasting applications. The raw data, consisting of Multi-day Operating Margin (MOM) sheets, ex-post Locational Marginal Pricing (LMP), and commodity data, was transformed into three purpose-built datasets: `master_hourly`, `master_daily`, and `master_full`.

5.5.1 Data Acquisition and Temporal Integrity

A critical component of the preprocessing pipeline was the enforcement of "Day-Ahead" (D-1) constraints. To accurately model the information available to market participants prior to the Day-Ahead market close (09:30 CT), a strict filter was applied to all ex-ante predictors:

$$\text{Publication Date} = \text{Target Date} - 1 \tag{5.1}$$

This ensures that all MOM predictors (Solar, Wind, Load Forecasts) represent the state of knowledge exactly one day prior to the target operating hour.

5.5.2 MISO Market Data Processing

The raw MISO data was classified into two distinct granularities:

- **Hourly Sheets:** Genuinely hourly data (24 observations per day), specifically for Solar and Wind forecasts. Completeness checks verified approximately 22,000 rows per sheet, covering the period from January 2023 to August 2025.

- **Daily Peak Sheets:** Data reported as a single daily peak value (e.g., Load Uncertainty, Outage Outlooks, Regional forecasts for North/Central/South). These contained approximately 959 rows each.

5.5.3 Ex-Post and Commodity Data

Ex-post target variables, including LMP and its components (Marginal Congestion Component, Marginal Loss Component, Marginal Energy Component), were loaded and cleaned. Commodity prices (Oil and Gas) were ingested and lagged to align with the D-1 decision horizon (1-day lag for oil, 1-week lag for gas).

5.5.4 Feature Engineering

Lag Generation

To capture temporal dependencies in the pricing signals, autoregressive features were generated for the target variables (LMP, MCC, MLC, MEC). The following lags were computed:

- $t-1$ (Previous hour)

- $t-24$ (Same hour, previous day)

- $t-168$ (Same hour, previous week)

Temporal Features

Time-based features were derived to capture seasonality and market cycles, including: Hour, Day of Week, Quarter, Season, Day of Year, and Boolean flags for weekends, business hours, and peak periods.

Aggregation

For the daily datasets, hourly predictors (Solar and Wind) were aggregated to create daily summary statistics. For every target date, the pipeline calculated the Minimum, Maximum, Mean, and Sum of the hourly ex-ante forecasts.

5.5.5 Final Dataset Construction

The pipeline culminated in the export of three synchronized datasets, as detailed in Table 5.1.

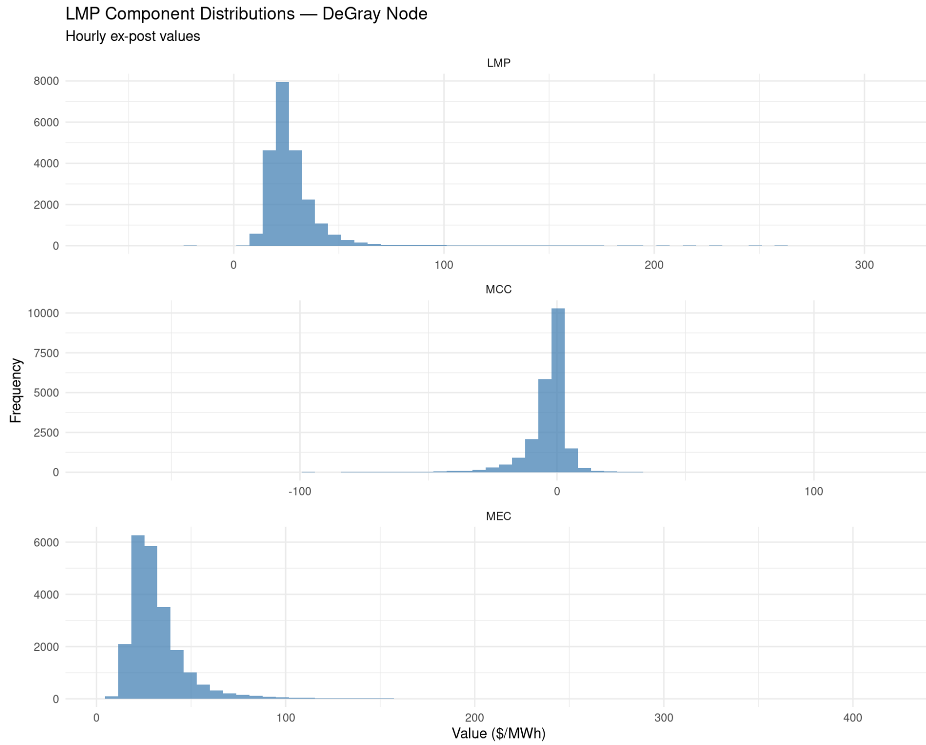

Figure 5.8: Histogram of Degray Node LMP broken down by the components MEC, MCC

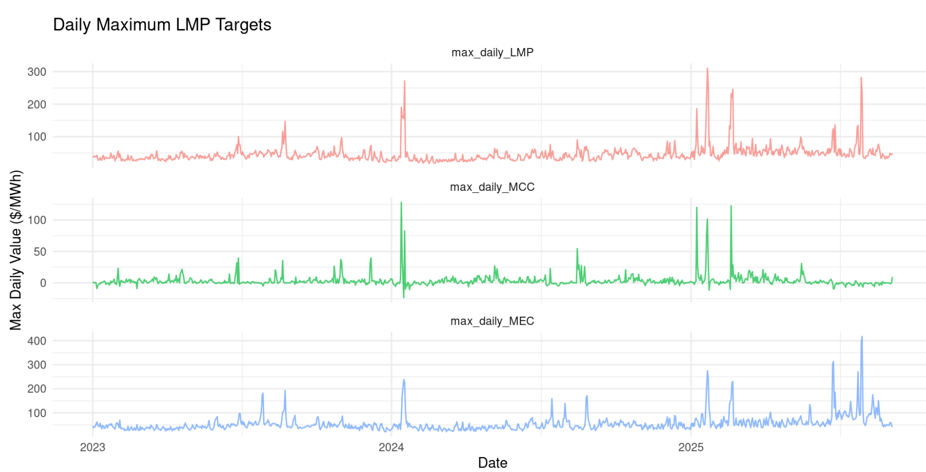

Figure 5.9: Degray Node LMP and its components MEC and MCC aggregated to a daily max over the period of collection

Table 5.1: Description of Final Processed Datasets

Dataset	**Dimensions**	**Description**
`master_hourly`	23,471 rows 36 columns	Granular target data. Contains one row per target hour. Includes ex-ante hourly sheets (Solar/Wind), ex-post LMP targets, lagged features, and temporal markers.
`master_daily`	978 rows 163 columns	Macro-level view. Contains one row per target date. Includes all daily-peak MOM sheets, aggregated hourly forecasts, and daily LMP summary targets (Min/Max/Mean).
`master_full`	23,471 rows 190 columns	The comprehensive modeling set. Constructed by joining `master_hourly` with `master_daily`. Daily predictors were broadcast to each constituent hour of the target date.

6 Market Model Development and Validation

Given the MISO market's operational focus on peak load planning, this study shifts the prediction target from granular hourly LMPs to the `max_daily_LMP`. This metric represents the single most critical price point for daily resource commitment and reward incentive analysis.

6.1 Methodology and Experimental Design

6.1.1 Target Variable Transformation

The target variable, `max_daily_LMP`, exhibits a log-normal distribution with significant right-skewness due to price spikes. To stabilize variance and ensure valid error metrics, a logarithmic transformation was applied:

$$y_t = \ln(\mathrm{MaxLMP}_t) \tag{6.1}$$

All performance metrics reported in this chapter are calculated on the log scale unless otherwise noted as "Original Scale ($/MWh)."

6.1.2 Validation Strategy: The Walk-Forward Approach

Time-series data violates the independence assumption of standard K-fold cross-validation. To rigorously test model performance, this study employed a two-phase validation strategy:

1. **Development Phase (2023–2024):** Used for hyperparameter tuning (e.g., Ridge λ, XGBoost depth). XGBoost was originally used with 10-fold and performed poorly, switching to 3-fold. Consider making it explicit that the 10-fold CV finding applies to Ridge tuning, and that XGBoost required 3-fold due to training set size sensitivity.

2. **Walk-Forward Validation (2025):** Used for final performance assessment. The model was trained on all data prior to January 1, 2025, and predicted the first month of 2025. The training window was then expanded to include January, and the model was retrained to predict February. This simulates the actual operational workflow of a Data Scientist updating the model monthly.

6.1.3 Rationale for Regression-Based Approach over Autoregressive Models

A natural alternative to the regression-based forecasting approach developed here would be a purely autoregressive (AR) or integrated moving-average (ARIMA/SARIMA) model, which generates predictions from the historical trajectory of the price series itself. Such models were considered and deliberately excluded for two reasons grounded in the operational context of this framework.

The primary purpose of the market model is not prediction accuracy alone but decision support for hydropower dispatch in a federal power operations setting. In this context, the interpretability of forecast drivers is operationally essential: an operator or regulator must be able to understand *why* the model recommends dispatch on a given day — that natural gas prices are elevated, that wind forecast uncertainty is high, or that load obligations are projected to strain reserve margins. Autoregressive models produce forecasts as weighted functions of lagged price observations, offering no direct mapping to observable market conditions that operators can independently verify or challenge. The Ridge Regression formulation, by contrast, produces coefficients with physical interpretations tied to publicly available MISO market reports, making the dispatch signal transparent and auditable.

The temporal dependencies present in the series are instead captured through explicit autoregressive *features* (lagged values at $t-1$, $t-7$, and $t-30$) incorporated directly as predictors within the Ridge Regression framework. This approach retains the persistence signal that AR models exploit while embedding it within a broader feature set that includes economic and market structure variables, yielding a model that is simultaneously temporally aware and operationally interpretable.

6.2 Baseline Performance (Development Phase)

Three baseline models were established to define the "floor" of acceptable performance using the 2023–2024 development set.

- **Persistence:** Assumes $y_t = y_{t-1}$.
- **Seasonal Naive:** Assumes $y_t = y_{t-7}$ (same day last week).
- **Ridge Regression (Baseline):** A linear model with L_2 regularization using all available MOM predictors.

As shown in Table 6.1, the Ridge Regression baseline significantly outperformed naive approaches, indicating strong signal availability in the MOM forecast sheets.

6.3 Advanced Modeling and Comparison

Two primary candidate architectures were evaluated: XGBoost (Gradient Boosted Decision Trees) to capture non-linearities and regime-switching behavior, and Regularized Linear

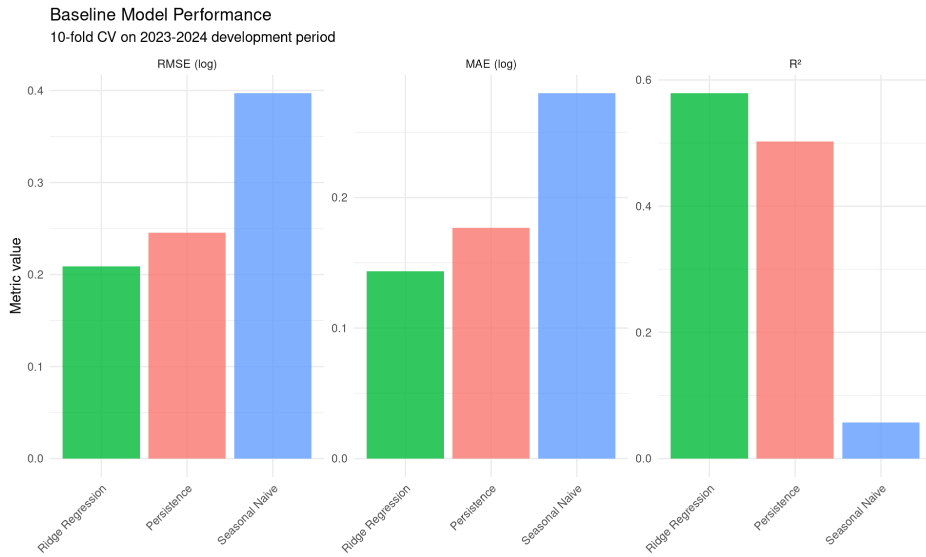

Figure 6.1: Baseline Model Performance Metrics

Table 6.1: Baseline Model Performance (2023–2024 Development Set)

Model	RMSE (log)	MAE (log)	MAPE (Original Scale)
Seasonal Naive	0.397	0.280	31.1%
Persistence	0.245	0.177	18.6%
Ridge Regression	**0.209**	**0.144**	**14.2%**

Models (Ridge/Elastic Net) to handle the high multicollinearity of the 112 predictor variables. From these and the models for benchmarking, it was intended to be used as inference if more sophisticated machine learning techniques would be appropriate.

6.3.1 XGBoost Implementation

An XGBoost regressor was trained with a manually validated 3-fold cross-validation scheme. XGBoost was initially tuned with 10-fold cross-validation; however, this produced training folds too small for the model to learn effectively, yielding degraded performance. Switching to 3-fold CV resolved this issue. Hyperparameters were tuned to a maximum depth of 6 and a learning rate (η) of 0.3.

6.3.2 Linear Model Tuning

Extensive diagnostics were performed on the linear approach:

- **PCR vs. Ridge:** Principal Component Regression (PCR) achieved an RMSE of 0.211 using 43 components. Ridge Regression achieved a lower RMSE of 0.207. This suggests that "soft shrinkage" of variables is superior to the "hard truncation" of PCR for this dataset; low-variance components still contain valuable price signals.

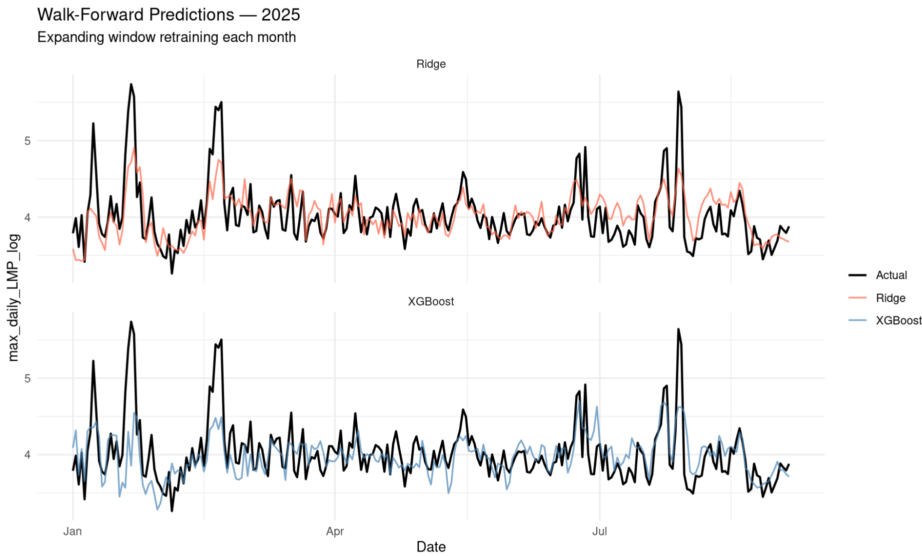

Figure 6.2: Model predictions plotted with actual observations

- **Elastic Net:** A grid search for the mixing parameter α found that $\alpha = 0$ (pure Ridge) minimized error. Sparsity (Lasso) was not beneficial, implying that retaining all MOM predictors—even highly correlated ones—improves forecast accuracy.
- **Interaction Terms:** Physical interaction terms were engineered to model supply curve dynamics (e.g., Load $\times$ Gas Price). While theoretically sound, these improved performance only marginally, suggesting the linear model was already capturing these proxies.

6.4 Walk-Forward Results (2025 Holdout)

The definitive test of model capability was the 2025 walk-forward validation. Surprisingly, the simpler Ridge Regression model outperformed the XGBoost ensemble across almost all metrics.

Table 6.2: Walk-Forward Validation Results (Jan–Sept 2025)

Model	RMSE (log)	MAE (log)	R^2
XGBoost	0.326	0.225	0.348
Ridge Regression	**0.284**	**0.210**	**0.497**

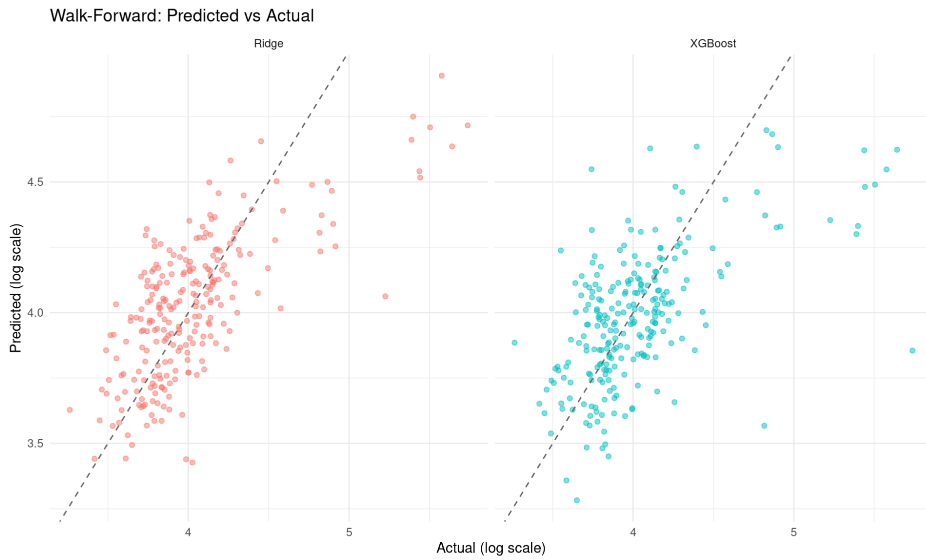

Figure 6.3: Actual versus Predicted variables

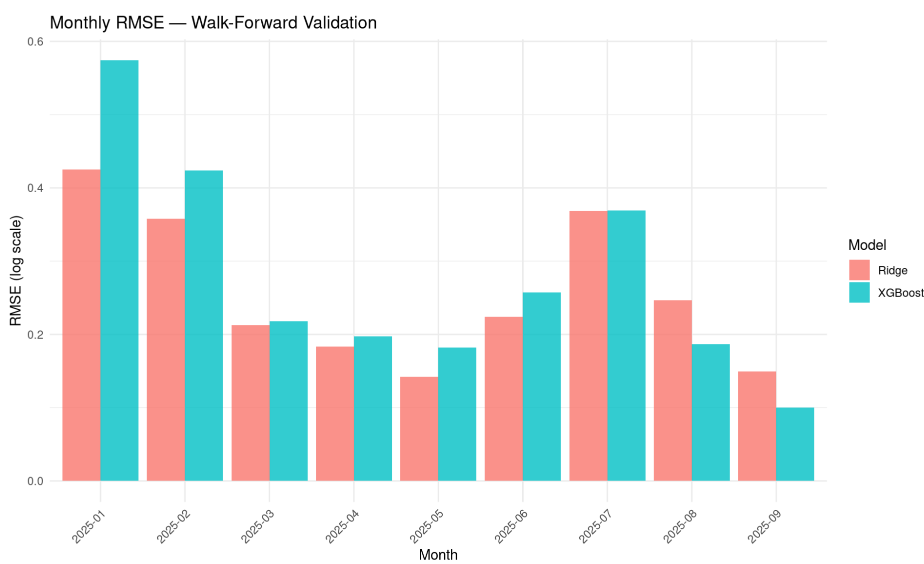

Figure 6.4: Residual Means Squared Error RMSE as it predicted each month for both models

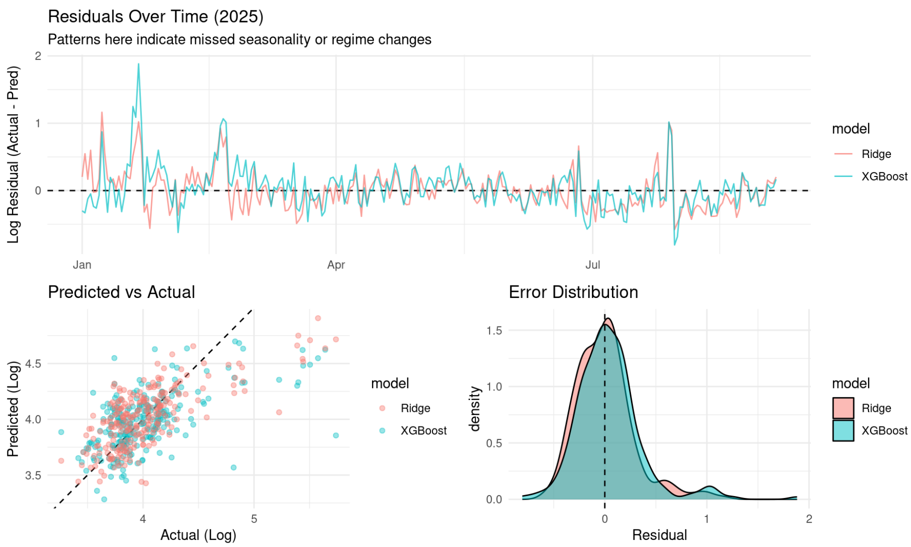

Figure 6.5: Residuals of the model over time, showing large residuals during volatile events

6.4.1 Regime Analysis

To understand the divergence in performance, errors were analyzed based on market volatility. A "High Volatility" regime was defined based on the magnitude of price spikes.

- In Normal Regimes, Ridge and XGBoost performed comparably (RMSE 0.25 vs 0.27).
- In High Volatility Regimes, Ridge significantly outperformed XGBoost (RMSE 0.349 vs 0.438).

This counter-intuitive result suggests that during extreme events, the tree-based model tended to overfit to noise or fail to extrapolate beyond the training bounds, whereas the linear constraints of Ridge Regression provided a safer, more robust extrapolation.

6.5 Final Model Selection and Interpretation

Based on the superior performance in the walk-forward validation and robust handling of high-volatility regimes, the Ridge Regression model was selected for production.

6.5.1 Feature Importance

The coefficients of the final model ($\lambda = 0.0201$) reveal the primary drivers of Daily Max LMP:

1. **Gas Price (Lagged):** The strongest positive driver ($\beta = 0.137$).

   This confirms that marginal fuel costs remain the dominant setter of peak prices.

2. **Seasonality:** Cosine and Sine terms for Month were highly significant, capturing the macro-seasonal drifts in demand.

3. **Wind Uncertainty:** The `MAPE_wind_unc` feature had a positive coefficient, indicating that higher uncertainty in wind forecasts correlates with higher risk premiums in the Day-Ahead market.

4. **Load Uncertainty:** Interestingly, `Std_Dev_load_unc` showed a negative coefficient. One plausible interpretation is that in the context of MISO, this may imply that periods of high predicted load variance trigger conservative operator behavior (over-commitment of reserves), which inadvertently dampens the final LMP.

7 Physical Environmental Constraint Model

7.1 HEC-HMS Model Configuration

To evaluate the physical consequences of each release strategy, the two operational scenarios were simulated using the HEC-HMS. HEC-HMS was selected for its capacity to model hydrologic runoff then route prescribed outflow schedules through the reservoir water balance and return daily end-of-day storage volumes, which serve as the primary physical state variable in the risk component of the objective function. The model was configured for Degray Lake using the USACE Vicksburg District regulation parameters, with the conservation pool bounded between elevation 367 ft MSL (bottom of power pool, $V_{\text{bottom}} \approx 261{,}500$ ac-ft) and elevation 408 ft MSL (top of power pool, $V_{\text{top}} = 654{,}700$ ac-ft), yielding an active power pool storage of 393,200 ac-ft.

The simulation period spans January 1 through December 30, 2025, corresponding to the operational year covered by the Ridge regression LMP prediction model. Daily inflow was computed from historical gridded precipitation within the basin model, and losses were calibrated from observed streamflow records. The water balance at each daily time step follows:

$$S_{t+1} = S_t + (I_t - R_t) \cdot \Delta t \tag{7.1}$$

where S_t is reservoir storage (ac-ft) at the beginning of time step t, I_t is total inflow (CFS), R_t is prescribed outflow (CFS), and Δt is the time step duration converted to acre-feet (1 CFS $\cdot$ day = 1.9835 ac-ft).

7.2 Scenario Outflows

Two distinct outflow time series were evaluated in the HMS model, corresponding to the two operational scenarios evaluated in this study.

Scenario A (Baseline — Flood and Minimum Release Only) prescribes a constant minimum environmental release of 157 CFS on all non-flood days, consistent with USACE minimum flow requirements for DeGray. During hydrologically significant inflow events — identified as days where the required outflow to maintain pool elevation below the top of flood pool (423 ft MSL) exceeded 6,000 CFS — flood control releases were prescribed according to the regulation manual routing schedule. This scenario represents the conservative operational

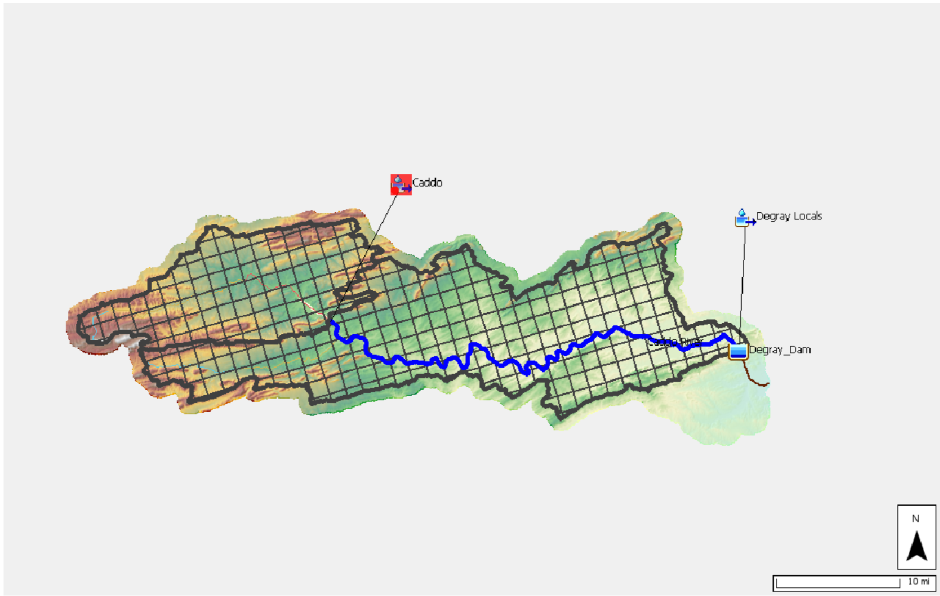

Figure 7.1: Basin Model within HEC-HMS

baseline: no market participation, full conservation of energy storage, and strict adherence to physical constraints. Scenario A thus provides the upper bound on energy-in-storage throughout the year.

Scenario B (Market-Dispatch) retains all flood control and minimum release operations from Scenario A as a base layer, and superimposes additional turbine dispatch on days where the Ridge regression model predicted LMP exceeding the 75th percentile threshold (\$65.45/MWh). On qualifying dispatch days excluding any day already subject to flood control release — both generating units were assumed to operate at max turbine release (combined 6,000 CFS) for a four-hour peak window. For simplicity the model operates at a daily time step, this four-hour dispatch was converted to a daily average outflow equivalent:

$$\bar{R}_{\text{dispatch}} = \frac{Q_{\text{full}} \cdot t_{\text{dispatch}} + Q_{\text{min}} \cdot (24 - t_{\text{dispatch}})}{24} = \frac{6{,}000 \cdot 4 + 157 \cdot 20}{24} \approx 1{,}131 \text{ CFS} \quad (7.2)$$

This daily average release preserves the volumetric equivalence of a four-hour max-turbine release dispatch event while remaining compatible with the HMS daily time step. A total of 62 dispatch days were initially flagged by the Ridge model's Q4 signal; 2 were suppressed due to concurrent flood control operations, yielding a net active dispatch count of 60 days out of the 365 days of analysis period.

7.3 Energy-in-Storage Computation

The HMS-simulated daily storage volumes were used to compute an energy-in-storage percentage for each scenario at each time step. Rather than converting directly to megawatt-hours which would require a precise head-discharge-efficiency relationship across varying pool elevations a normalized storage percentage was adopted as the primary physical state metric:

$$E_{\%,t} = \frac{S_t - V_{\text{bottom}}}{V_{\text{top}} - V_{\text{bottom}}} \times 100 \tag{7.3}$$

where $E_{\%,t}$ represents the fraction of usable power pool storage remaining at time t, expressed as a percentage of the active conservation pool volume. A value of 100% corresponds to the reservoir at full pool (408 ft MSL, 654,700 ac-ft), while 0% corresponds to the bottom of the power pool (367 ft MSL, 261,500 ac-ft). This metric is directly analogous to the state of charge concept used in battery storage optimization and provides an intuitive basis for comparing the two scenarios across seasonal inflow regimes.

The elevation-area-storage relationship used to contextualize storage volumes relative to pool elevations was digitized from the USACE Vicksburg District Degray Lake regulation manual and is provided in Appendix A.

7.4 Storage Trajectory Comparison

The HMS simulation results reveal a systematic divergence in reservoir storage between the two scenarios, accumulating progressively through the low-inflow summer months before partially recovering during fall recharge. Figure 7.2 illustrates the daily storage trajectories for both scenarios alongside the pool elevation reference bounds.

On individual dispatch days, Scenario B released an incremental $(\bar{R}_{\text{dispatch}} - Q_{\text{min}}) \cdot \Delta t \approx (1{,}131 - 157) \cdot 1.9835 \approx 1{,}932$ ac-ft above the Scenario A baseline. Summed across all 60 active dispatch days, the total incremental volumetric release attributable to market dispatch is approximately **115,000** ac-ft, representing **35%** of the usable conservation pool.

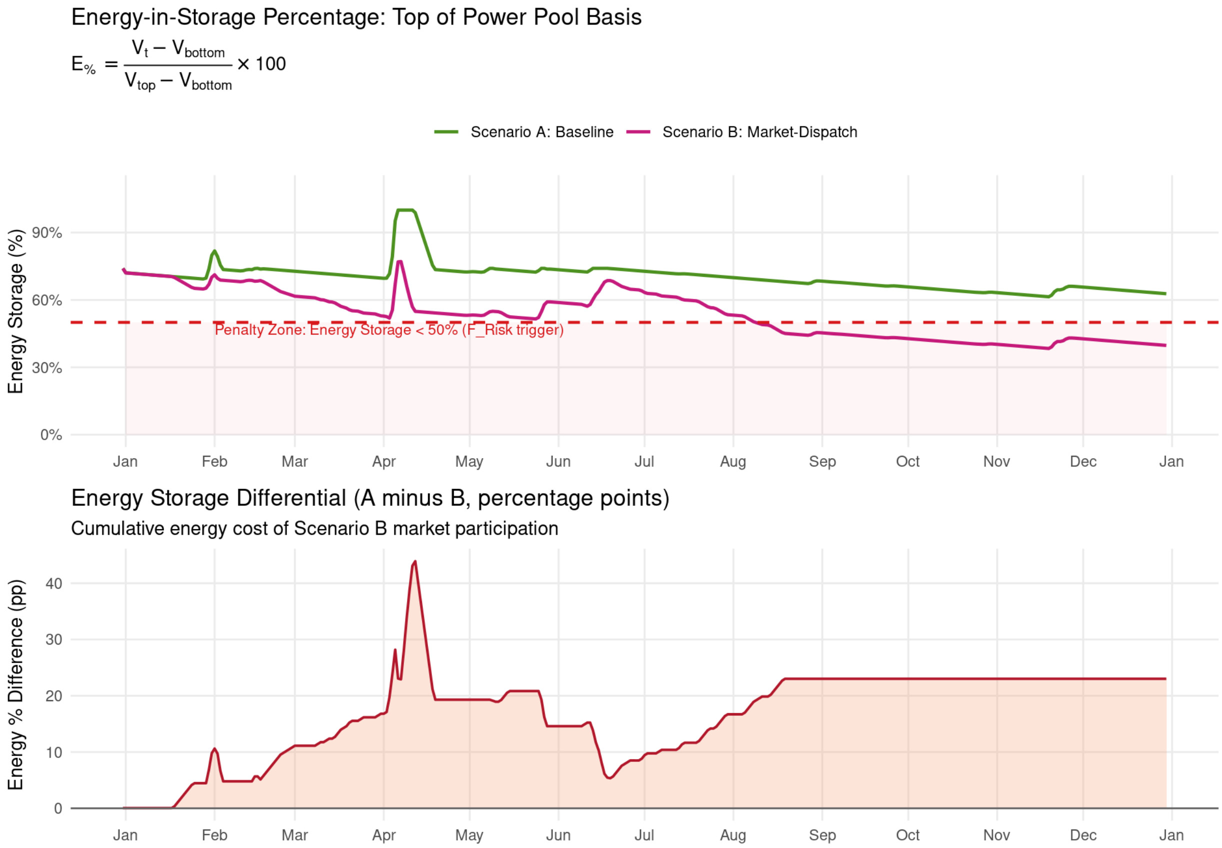

Figure 7.2: Daily end-of-day reservoir storage (acre-feet) for Scenario A and Scenario B simulated in HEC-HMS, January–December 2025 (upper panel). The lower panel shows the cumulative storage differential (Scenario A minus Scenario B), representing the volumetric cost of market-dispatch operations relative to the conservation baseline.

8 Operational Scenario Analysis: Objective Function Evaluation and Market Performance

8.1 Overview

This chapter presents the quantitative evaluation of the multi-objective framework across the two HMS-simulated operational scenarios. The analysis proceeds in three stages: (1) computation of the risk penalty term F_{Risk} from the physical storage trajectories returned by HEC-HMS, (2) estimation of the Reward term F_{Reward} from market revenue realized on dispatch days, and (3) assembly of the composite objective function F and sensitivity analysis across the Reward-risk weighting space. A monthly dispatch performance assessment then characterizes where and when the Ridge model's Q4 signal captured or missed actual market value.

8.2 Risk Penalty Evaluation

8.2.1 Penalty Structure

The risk component of the objective function is defined as:

$$F_{\text{Risk}} = F_{\text{Flood}} + F_{\text{Quality}} + F_{\text{Storage}} \tag{8.1}$$

Because both scenarios were constructed to satisfy flood pool and minimum environmental flow constraints by design, $F_{\text{Flood}} = 0$ and the environmental flow penalty within F_{Quality} equals zero across the entire simulation period for both scenarios. The operative penalty term is therefore the energy storage component, which activates when the normalized energy-in-storage percentage $E_{\%}$ falls below the 50% soft constraint threshold:

$$F_{\text{Storage}} = \sum_{t=1}^{T} \alpha_L \left[\max\left(0,\ 50 - E_{\%,t}\right)\right]^2 \tag{8.2}$$

8.2.2 Penalty Results

Scenario A remained above the 50% energy storage threshold for all 364 days of the analysis period, incurring zero penalty ($F_{\text{Risk},A} = \$0$). Scenario B violated the threshold on 144 days, concentrated primarily in the low-inflow summer and fall period following the accumulation of dispatch-driven drawdown through the spring. The maximum observed violation reached 11.6 percentage points below the threshold, occurring in September–October as the reservoir continued drawing down with no new dispatch signal and insufficient inflow to recover. The total energy storage penalty assessed against Scenario B was:

$$F_{\text{Risk},B} = \sum_{t=1}^{364} \alpha_L \left[\max\left(0,\ 50 - E_{\%,t}\right)\right]^2 = \$455{,}506 \tag{8.3}$$

This penalty is not a cash payment but a representation of operational risk cost embedded in the objective function — its magnitude relative to the reward term determines whether the market-dispatch strategy is net-beneficial under the chosen weighting.

8.3 Market Revenue Assessment

8.3.1 Revenue Methodology

On days where Scenario B dispatched, revenue was estimated as the product of energy generated during the four-hour dispatch window and a weighted peak LMP. Because the dispatch was assumed to target the peak hours of the day, the actual daily average LMP was adjusted by a conservative multiplier of 0.8 to approximate the on-peak price realized during the dispatch window:

$$\text{Revenue}_t = P_{\text{installed}} \cdot t_{\text{dispatch}} \cdot (0.8 \times \text{LMP}_{\text{actual},t}) \tag{8.4}$$

where $P_{\text{installed}} = 74$ MW (combined nameplate capacity), $t_{\text{dispatch}} = 4$ hours, and $\text{LMP}_{\text{actual},t}$ is the observed day-ahead LMP for Degray Node on dispatch day t.

8.3.2 Prediction Signal Performance

Over the 62 dispatch days triggered by the Ridge model's Q4 signal, the model correctly identified days where the actual LMP exceeded the \$65.45/MWh threshold on 34 occasions (54.8%), while dispatching on 28 days where actual prices ultimately settled below the threshold. This 54.8% capture rate reflects a meaningful but imperfect signal the Ridge model's mean prediction error on dispatch days was \$16.86/MWh, indicating a systematic tendency to under-predict prices on the days it chose to dispatch, which makes sense when considering the highest prices come from system anomalies not in common trends within the dataset, with an RMSE of \$7.22/MWh.

The 28 missed days where dispatch occurred but actual LMP fell below the threshold represent the most consequential failure mode of the single-model strategy. These events constitute a dual loss: incremental storage was released without commensurate revenue, simultaneously increasing F_{Risk} (through additional drawdown) while contributing below-threshold returns to F_{Reward}. This dynamic is precisely the operational condition that an ensemble framework is designed to address. Independent signal from another predictive component would provide a check against high-confidence but ultimately incorrect Q4 predictions.

8.3.3 Total Revenue

Total revenue attributable to Scenario B dispatch operations over the 2025 analysis period was \$1,415,531, decomposed as \$1,048,398 generated on the 34 correctly-timed dispatch days and \$367,134 generated on the 28 days where dispatch occurred but actual LMP fell below the Q4 threshold.

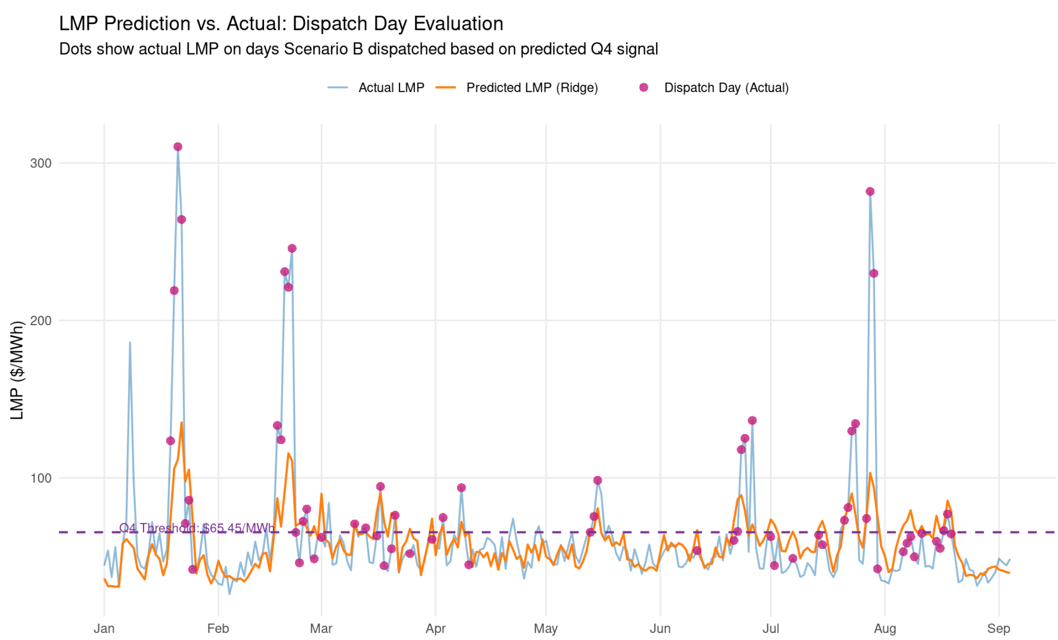

Figure 8.1: Predicted versus actual LMP on Scenario B dispatch days. Points above the horizontal dashed line (\$65.45/MWh) indicate days where actual prices justified the dispatch signal (captured, $n = 34$); points below represent dispatch events where actual prices failed to meet the threshold (missed, $n = 28$)

8.4 Monthly Dispatch Performance

Table 8.1 presents the month-by-month decomposition of dispatch activity, prediction accuracy, energy storage cost, and revenue generation. Several operational regimes are evident in the data.

January and February produced the highest per-dispatch revenues (\$379,803 and \$431,561 respectively) and the strongest prediction accuracy (86% and 70% capture rates),

Table 8.1: Monthly Dispatch Performance: Prediction Signal vs. Actual Market Outcomes

Month	Dispatch Days	Avg Pred LMP ($/MWh)	Avg Actual LMP ($/MWh)	% Cap.	Avg ΔPP	Monthly Revenue
Jan 2025	7	100.44	159.39	86%	1.79	$379,803
Feb 2025	10	82.47	126.78	70%	6.82	$431,561
Mar 2025	10	76.29	64.70	40%	13.60	$220,235
Apr 2025	3	69.03	71.08	67%	25.16	$72,590
May 2025	3	71.84	79.77	67%	19.17	$81,457
Jun 2025	6	77.40	93.25	67%	10.76	$190,450
Jul 2025	13	78.44	101.82	54%	12.14	$450,569
Aug 2025	10	73.49	61.15	20%	20.76	$208,161
Sep 2025	0	—	—	—	23.02	$0

Note: "Avg Pred LMP" and "Avg Actual LMP" are in $/MWh. "% Cap." = percentage of dispatch days where actual LMP exceeded the $65.45/MWh threshold. "Avg ΔPP" = average percentage-point difference in energy-in-storage (Scenario A minus Scenario B). Revenue figures represent incremental amounts above the conservation baseline, computed using a 0.8× peak-hour LMP multiplier over the 4-hour dispatch window at 74 MW installed capacity. September records no dispatch activity due to the Ridge model prediction dataset boundary (September 4, 2025).

reflecting the winter price volatility driven by heating load and the cold-snap events characteristic of MISO South in early 2025. The average actual LMP on dispatch days during these months ($159.39/MWh and $126.78/MWh) substantially exceeded the Ridge model's predictions ($100.44/MWh and $82.47/MWh), indicating that the model conservatively estimated winter price spikes and that the realized revenue exceeded what the dispatch signal anticipated. The energy storage cost during these months was comparatively modest the high-inflow winter hydrology partially offset dispatch drawdown, with average energy differentials of 1.79 and 6.82 percentage points respectively.

March and August illustrate the opposite regime. In March, the model dispatched on 10 days at an average predicted LMP of $76.29/MWh, but actual prices averaged only $64.70/MWh below the dispatch threshold yielding a 40% capture rate and contributing to a 13.60 percentage-point average energy differential. August repeated this pattern more severely: 10 dispatch events at a 20% capture rate, with actual prices averaging $61.15/MWh against predictions of $73.49/MWh. By August, cumulative dispatch drawdown had already reduced Scenario B's energy storage to below the 50% threshold, meaning each additional missed dispatch event compounded both the revenue shortfall and the penalty accrual simultaneously.

July was the highest-volume dispatch month (13 days) and produced the largest monthly revenue ($450,569) at a 54% capture rate, consistent with summer air conditioning load driving sustained elevated prices in MISO South. However, the 13 dispatch days also contributed the largest monthly volumetric drawdown, with the average energy differential reaching 12.14 percentage points.

September through December recorded no dispatch activity, as the Ridge model produced no Q4 predictions after September 4, 2025 the boundary of the prediction dataset. During this period, the 70,450 ac-ft storage gap between scenarios remained essentially fixed, as both scenarios released only the minimum environmental flow with no inflow sufficient to close the deficit before year end.

8.5 Composite Objective Function

8.5.1 Base Case Results

Under the base case Reward-risk weighting of $\lambda_1 = 0.60$ and $\lambda_2 = 0.40$, the composite objective function evaluates as:

$$F_A = \lambda_1 \cdot F_{\text{Reward},A} - \lambda_2 \cdot F_{\text{Risk},A} = 0.60 \cdot \$0 - 0.40 \cdot \$0 = \$0 \tag{8.5}$$

$$F_B = \lambda_1 \cdot F_{\text{Reward},B} - \lambda_2 \cdot F_{\text{Risk},B} = 0.60 \cdot \$1{,}415{,}531 - 0.40 \cdot \$455{,}506 = \$667{,}116 \tag{8.6}$$

Scenario B dominates Scenario A under the base weighting by a margin of \$667,116, indicating that at equal-to-moderate Reward preference, the market-dispatch strategy generates sufficient revenue to justify its energy storage cost. It is important to note, however, that this result is sensitive to the penalty coefficient α_L and the reward weight λ_1, both of which require calibration against operator risk tolerance.

8.5.2 Lambda Sensitivity

As λ_1 decreases (increasing risk aversion), the penalty assessed against Scenario B's 144 violation-days increasingly offsets its revenue advantage. The crossover point the weighting at which $F_A = F_B$ can be derived analytically:

$$\lambda_1^* = \frac{F_{\text{Risk},B}}{F_{\text{Reward},B} + F_{\text{Risk},B}} = \frac{\$455{,}506}{\$1{,}415{,}531 + \$455{,}506} \approx 0.25 \tag{8.7}$$

At $\lambda_1^* \approx 0.25$, Scenario A and Scenario B produce equal composite scores. For any $\lambda_1 > 0.25$, market dispatch (Scenario B) is preferred; below this threshold, the conservation baseline (Scenario A) dominates. The relatively low crossover value suggests that the market-dispatch strategy is robust across a wide range of Reward-risk preferences — only a strongly risk-averse operator weighting physical storage at more than 75% of the total objective would prefer the baseline. This result reflects the favorable January–February revenue performance outweighing the summer penalty accumulation under most reasonable weighting schemes.

8.5.3 Pareto Trade-off Analysis

The composite weighting λ_1 alone does not fully characterize the operational trade-off. To examine the underlying structure of the reward-risk relationship, a Pareto frontier was constructed by sweeping the LMP dispatch threshold across the 65th through 95th percentile

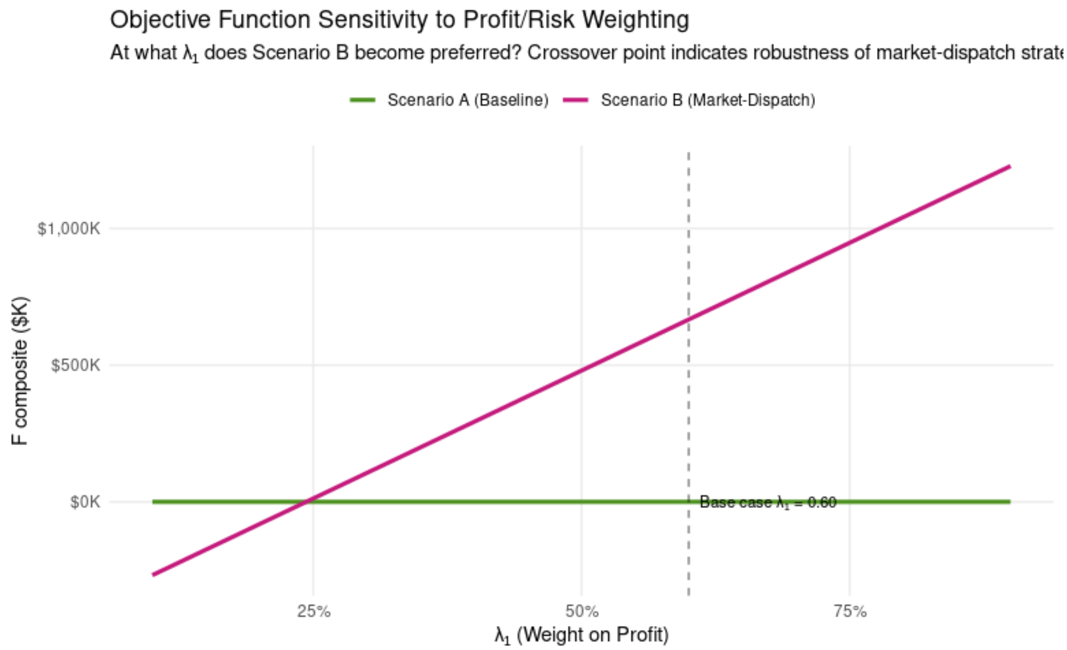

Figure 8.2: Composite objective function F as a function of the Reward weight λ_1 for Scenario A (baseline, green) and Scenario B (market-dispatch, pink). The crossover at $\lambda_1^* \approx 0.25$ identifies the weighting below which risk-averse preferences favor the conservation baseline. The vertical dashed line marks the base case $\lambda_1 = 0.60$.

of the predicted price distribution ($61.50–$89.60/MWh). At each candidate threshold, the full annual dispatch sequence was re-evaluated: revenue was computed from actual LMP outcomes on qualifying dispatch days, and the storage risk penalty was re-derived by applying marginal storage adjustments to the HMS-simulated Scenario B trajectory, anchored to the verified $F_{\text{Risk},B} = \$455{,}506$ result at the 75th percentile.

The frontier reveals a pronounced elbow near the 75th–80th percentile threshold range, below which each additional dollar of revenue is accompanied by a disproportionate increase in cumulative storage penalty. Above this elbow – that is, at higher dispatch thresholds – the frontier is nearly flat in the risk dimension, indicating that reducing the number of dispatch days from the 75th percentile level yields minimal additional storage benefit while meaningfully sacrificing revenue. This asymmetry identifies the 75th percentile as a well-calibrated operating point: it captures the majority of available market value while remaining just above the threshold where storage risk begins to accelerate sharply.

Notably, the 75th percentile threshold was selected on purely distributional grounds prior to this analysis, as the upper quartile of the predicted LMP distribution. The Pareto analysis confirms retroactively that this selection aligns with the frontier elbow, providing independent validation of the threshold choice through a multi-objective lens. The crossover point derived analytically in the lambda sensitivity analysis — $\lambda_1^* \approx 0.25$ — corresponds to the weighting at which the Scenario B operating point (Figure 8.3) becomes indifferent to Scenario A, and is consistent with the frontier geometry: Scenario B's storage penalty is modest relative to its revenue advantage under all but the most risk-averse operator

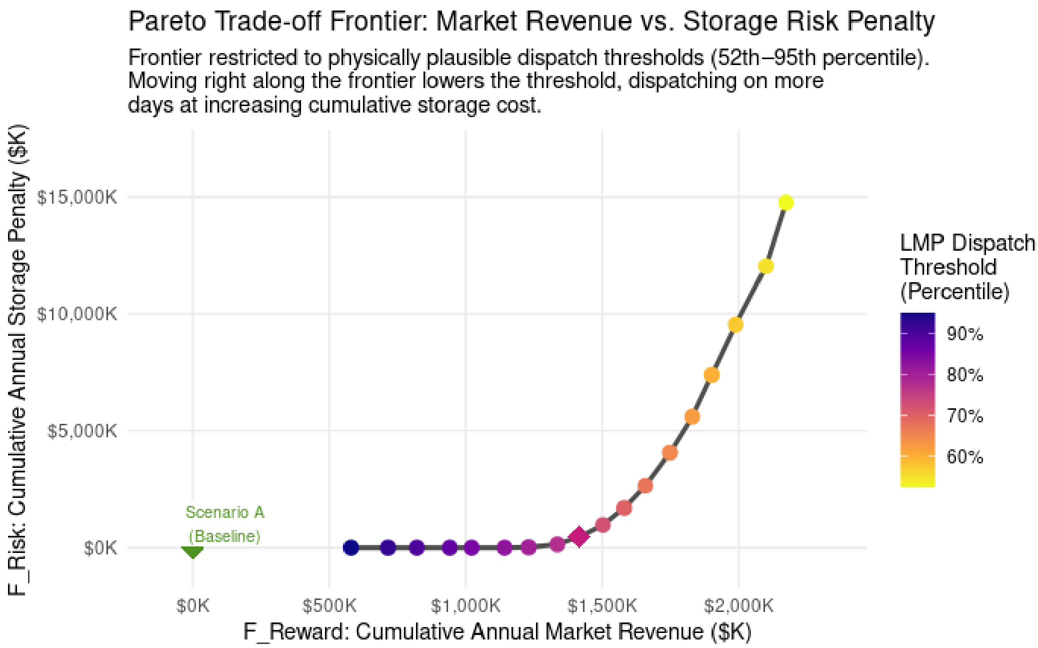

Figure 8.3: Pareto trade-off frontier between cumulative annual market revenue (F_{Reward}) and cumulative annual storage risk penalty (F_{Risk}) across candidate LMP dispatch thresholds (65th–95th percentile). Each point represents a distinct threshold operating policy evaluated over the 2025 analysis period. Scenario A (Baseline) named and Scenario B (75th percentile, $65.45/MWh) diamond symbol. The preferred direction (higher revenue, lower penalty). The frontier elbow near the 75th–80th percentile identifies the threshold region below which marginal storage risk accelerates disproportionately relative to incremental revenue gains.

preferences.

8.6 Framework Implications

The results of this operational analysis surface three findings directly relevant to the convergence-based decision support framework.

First, the 54.8% dispatch capture rate establishes a clear performance baseline for the single-model Q4 threshold strategy. A convergence framework that requires agreement between the Ridge regression signal and a second independent market model before committing to dispatch would be expected to improve this capture rate by filtering out the 28 missed events, at the cost of also potentially suppressing some correctly-timed dispatches. The net effect on F_{Reward} depends on the specificity of the second signal, but the directional improvement in F_{Risk} (fewer unwarranted drawdown events) is near-certain.

Second, the 144-day penalty accumulation in Scenario B is almost entirely a consequence of summer and fall operations, where the Ridge model ceased producing Q4 predictions but the storage deficit from prior dispatch events persisted without hydrologic recovery. This temporal asymmetry dispatch decisions made in winter and spring bearing storage consequences through the following fall underscores the importance of multi-period lookahead in the optimization horizon. A seven-day forecast window, as proposed in the Forecasted Informed Reservoir Operations FIRO integration component of the framework, would not have resolved this seasonal carry-over effect alone; a storage trajectory constraint on the optimization horizon is needed.

Third, the $\lambda_1^* \approx 0.25$ crossover indicates that the operational tradeoff is less balanced than the symmetric 0.60/0.40 base case weighting suggests. In practice, calibrating λ_1 and λ_2 to reflect Degray's actual regulatory obligations including USACE conditions, downstream water supply commitments, and recreation pool targets would replace the illustrative penalty coefficient α_L with operationally grounded values, providing a defensible basis for the crossover threshold in the decision support context.

Finally, The Pareto frontier more fully characterizes this asymmetry, showing that storage risk accumulates disproportionately at lower dispatch thresholds while revenue gains below the 75th percentile elbow increase without much risk.

9 Discussion

9.1 Synthesis of Framework Components

The preceding chapters developed three distinct but tightly coupled analytical layers.: an exploratory data analysis and feature engineering pipeline grounded in MISO South market structure (Section 5), a walk-forward validated statistical market model whose Ridge Regression formulation outperformed gradient-boosted alternatives on out-of-sample data (Section 6), and a physics-aware hydrologic simulation via HEC-HMS that translated market dispatch decisions into reservoir state trajectories with measurable risk penalties (Section 7–8). This chapter synthesizes those findings, examines the implications for the convergence-based decision support framework motivating this research, and reflects on the limitations that bound the current results.

9.1.1 Model Interoperability

A central design principle of the proposed framework is that the statistical/ML component and the physical/rule-based component remain independent neither model ingests outputs from the other during the prediction stage; instead, both contribute independently to the composite objective function $F = \lambda_1 F_{\text{Reward}} - \lambda_2 F_{\text{Risk}}$. The operational scenario analysis in Chapter 8 demonstrates that this architecture is practically feasible using existing, publicly available data and open-source software.

The independence of the two models also presents a practical advantage: each can be updated and retested without invalidating the other. This modularity is essential for long-term maintainability of any decision support tool in a federal power operations context.

9.1.2 Temporal Scale Mismatches and Multi-Period Planning

The seasonal carry-over effect identified in Section 8.6 is perhaps the most operationally consequential finding of this study. Dispatch decisions concentrated in January through July produced a storage deficit that persisted through the low-inflow summer and fall months without sufficient hydrologic recovery, ultimately triggering 144 penalty-days despite the absence of any market dispatch after early September. This result exposes a structural limitation of threshold-based dispatch rules that treat each day as conditionally independent of past decisions: in a system with limited active storage (approximately 393,200 acre-feet of usable conservation pool at Degray), each high-value dispatch event consumes a meaningful fraction of a shared, seasonally-recharged resource.

This finding has direct implications for the Forecast Informed Reservoir Operations (FIRO) integration component proposed in the broader research programs. A seven-day precipitation forecast horizon, as traditionally associated with FIRO implementations [8], would improve short-term inflow anticipation but would not resolve a carry-over effect that operates at seasonal scales. What is required is a multi-period storage trajectory constraint embedded directly in the optimization horizon—effectively, a minimum end-of-season storage requirement that limits cumulative dispatch during periods when hydrologic recovery is projected to be insufficient. Implementing such a constraint requires coupling the market model's dispatch sequence to a seasonal storage forecast, which is precisely the integration challenge that motivates this research.

9.1.3 Market Model Robustness and Non-Stationarity

The counter-intuitive finding that Ridge Regression outperformed XGBoost in the 2025 walk-forward holdout—particularly in high-volatility regimes (RMSE 0.349 vs. 0.438)—warrants careful interpretation. The result is consistent with the general machine learning literature on the bias-variance tradeoff under distribution shift: when out-of-sample conditions differ from the training distribution, the additional variance introduced by a flexible model can outweigh its in-sample bias advantage [2]. In the context of MISO South price dynamics, where price spikes are driven by supply-demand imbalances that may not recur with precisely the same feature signatures as in the training period, the linear constraints of Ridge Regression appear to provide a more stable extrapolation surface.

This observation supports the research framework's emphasis on model robustness under non-stationarity as a primary concern. However, it should not be interpreted as a general argument against tree-based methods for electricity price forecasting. The comparison was conducted on a 2025 holdout period that included unusual winter events and a transition in the solar generation mix that may have systematically disadvantaged the XGBoost model. A more definitive assessment would require evaluating both models across multiple calendar years with different hydrologic and market regimes. The walk-forward protocol implemented here—monthly retraining with an expanding training window—partially mitigates this concern by ensuring that the model's parameter estimates are continuously updated as new data accumulate.

9.1.4 Feature Importance and Market Dynamics

The Ridge model's coefficient structure confirms several theoretically motivated expectations about MISO South price drivers. The dominant role of lagged gas prices ($\beta = 0.137$) is consistent with the well-established finding that natural gas-fired peaking units set marginal prices during high-demand periods in most US market regions. The significance of wind uncertainty (positive coefficient) and load uncertainty (negative coefficient) reflects the increasingly important role of variable renewable energy in shaping day-ahead price risk premiums a dynamic that is expected to intensify as the MISO South footprint continues its renewable integration trajectory [9].

9.2 Implications for Decision Support Practice

9.2.1 Operator Decision Protocols

The lambda sensitivity analysis (Section 8.5.2) provides a practical tool for translating operator risk preferences into dispatch guidance. The crossover point at $\lambda_1^* \approx 0.25$ implies that only an operator who weights physical storage preservation at nearly more than three times the importance of market revenue would prefer the conservative baseline over the market-dispatch strategy. In practice, the appropriate weighting for the Degray project would be informed by USACE's water control plan, downstream water supply commitments, recreation pool management targets, and the regulatory constraints that govern Southwestern Power Administration's marketing obligations. Replacing the illustrative penalty coefficient α_L with operationally grounded values—derived from the economic cost of water supply curtailments or the regulatory penalties for pool elevation violations would transform the sensitivity curve from an analytical illustration into a defensible operational input. The Pareto frontier in Figure 8.3 extends this analysis by identifying not just whether to dispatch, but at what price threshold — information that is more directly actionable for operators calibrating the market trigger.

9.2.2 Implementation Pathway

The modular architecture of the proposed framework suggests a phased implementation pathway. In an initial phase, the Ridge Regression market model and HEC-HMS simulation can be deployed as standalone tools, providing daily dispatch recommendations and weekly storage trajectory forecasts to operators who retain final decision authority. This phase would also generate the operational performance data needed to calibrate the penalty coefficients with real-world consequences. In a subsequent phase, the convergence component—requiring a second market model would be added, with the agreement between models serving as an explicit confidence indicator that operators can use to adjust their discretionary dispatch decisions. The final phase would integrate precipitation forecast data from National Oceanic Atmosphere Association NOAA's Climate Prediction Center to extend the inflow anticipation horizon and reduce the seasonal carry-over risk identified in this study.

9.2.3 Broader Applicability

While this study is scoped to Degray Lake within the MISO South, the framework components are generalizable across several dimensions. The statistical market model relies exclusively on publicly available MISO market reports and commodity price data; the same pipeline could be applied to any hydropower asset within the MISO footprint with node-level LMP access. The HEC-HMS water balance formulation is standard and transferable to any single-reservoir system for which an existing regulation manual provides the necessary elevation-area-storage relationships. The multi-objective function structure separating reward maximization from risk penalization and coupling them through a weighted composite reflects a design pattern that is applicable to pumped storage, run-of-river, and multi-reservoir cascades, though the specific penalty terms would require adaptation to each system's regulatory context.

9.3 Limitations

Data Horizon. The market model was trained on 978 days of MISO market data spanning January 2023 through September 2025. This period is characterized by a specific pattern of natural gas price volatility, solar and wind penetration growth, and extreme weather events that may not be representative of future conditions. The walk-forward validation demonstrates reasonable out-of-sample stability over nine months of 2025, but longer evaluation periods across multiple weather and market cycles would strengthen the confidence in the model's operational readiness.

Daily Time Step. Both the market model and the HEC-HMS simulation operate at a daily temporal resolution. This is consistent with the day-ahead market's commitment horizon but precludes analysis of intra-day dispatch optimization, ramp rate constraints, and hour-to-hour price arbitrage opportunities that could be material for a 38 MW unit with limited storage. A sub-daily formulation would require hourly reservoir routing and a more detailed turbine-generator model than the simplified daily average outflow conversion used here.

Simplified Revenue Calculation. The revenue estimation in Section 8.3 applies a fixed 0.80 multiplier to the daily average LMP to approximate the on-peak price during the dispatch window. This multiplier is a conservative engineering assumption rather than a calibrated estimate. In practice, the realized dispatch price depends on the specific hours selected, the unit's offer curve, and MISO's settlement procedures for dispatchable units with multiple generating configurations.

Penalty Coefficient Calibration. The energy storage penalty coefficient α_L is specified as an illustrative value. The sensitivity analysis demonstrates that the qualitative conclusion—Scenario B dominates Scenario A for $\lambda_1 > 0.25$—is robust to moderate changes in the penalty structure, but the quantitative crossover threshold is sensitive to α_L. Operational deployment would require explicit calibration against the economic consequences of pool elevation violations as defined in the USACE water control plan.

10 Conclusion

10.1 Research Summary

This thesis developed and evaluated the foundational components of a physics-aware, market-responsive optimization framework for hydropower storage at Degray Lake, a 74 MW USACE-operated facility within the MISO South balancing authority. The research addressed a well-documented gap in the hydropower optimization literature: the absence of an integrated framework that simultaneously accounts for the economic objectives of electricity market participation and the physical constraints imposed by reservoir regulation, while providing operators with an interpretable confidence signal for day-ahead dispatch decisions.

The analytical work proceeded through four interconnected contributions. First, a comprehensive data assembly and feature engineering pipeline was constructed from publicly available MISO market reports, EIA commodity price series, and USACE operational records, yielding a 978-day dataset spanning January 2023 through September 2025 with 99.8% overall completeness. Second, a walk-forward validated market model was developed and benchmarked, with Ridge Regression selected over XGBoost as the production forecasting component on the basis of superior out-of-sample performance and robustness in high-volatility regimes. Third, a HEC-HMS hydrologic simulation was configured for the Degray watershed, enabling translation of market dispatch schedules into reservoir storage trajectories under two operational scenarios. Fourth, the resulting storage trajectories and market revenues were assembled into a composite multi-objective function and evaluated across a Reward-risk weighting space, establishing a quantitative foundation for future convergence-based decision support.

10.2 Recommendations

10.2.1 For Future Research

The most immediate research priority is implementation of the second optimization component required to operationalize the convergence mechanism.

Longer-term research directions include the extension of the validation period across multiple years with distinct hydrologic and market regimes; the development of a sub-daily dispatch formulation that captures intra-day price arbitrage; the calibration of penalty coefficients against documented economic consequences of USACE regulatory violations; and the integration of NOAA ensemble precipitation forecasts into a seasonal storage trajectory constraint.

10.3 Closing Remarks

The transition to a variable renewable energy-dominated grid creates both challenges and opportunities for flexible resources like hydropower. Conventional operational frameworks that treat hydropower as a dispatchable but largely passive load-following resource leave economic value unrealized and fail to exploit the market intelligence embedded in publicly available day-ahead forecast data. This thesis demonstrates, with a specific facility and a nearly three-year empirical dataset, that a data-driven market model coupled with a physics-grounded risk framework can identify actionable dispatch opportunities that are robust to operator risk preferences across a wide weighting range.

Bibliography

[1] Sibel Akkaya Oy, Serkan Inal, and Ali Ekber Özdemir. Operation of hydroelectric power plants, dam reservoirs, and energy trade using artificial neural networks. *Applied Sciences*, 15(1):183, 2025.

[2] Saad Aslam, Pyi Phyo Aung, Ahmad Sahban Rafsanjani, and Anwar P. P. Abdul Majeed. Machine learning applications in energy systems: current trends, challenges, and research directions. *Energy Informatics*, 8(1):62, 2025.

[3] Natei Ermias Benti, Mesfin Diro Chaka, and Addisu Gezahegn Semie. Forecasting renewable energy generation with machine learning and deep learning: Current advances and future prospects. *Sustainability*, 15(9):7087, 2023.

[4] Nick Harder, Anke Weidlich, and Philipp Staudt. Finding individual strategies for storage units in electricity market models using deep reinforcement learning. In *The 12th DACH+ Conference on Energy Informatics 2023*, Vienna, Austria, 2023.

[5] Hydropower Reform Coalition. Degray hydroelectric project, n.d.

[6] Jun Jia, Guangming Zhang, Xiaoxiong Zhou, Zhihan Shi, Mingxiang Zhu, and Xiaodong Lv. Research on joint dispatch of wind, solar, hydro, and thermal power based on pumped storage power stations. *Frontiers in Energy Research*, 2024.

[7] Basil Kraft, Martin Jung, Marco Körner, Sujan Koirala, and Markus Reichstein. Towards hybrid modeling of the global hydrological cycle. *Hydrology and Earth System Sciences*, 26:1579–1614, 2022.

[8] Donghoon Lee, Jia Yi Ng, Stefano Galelli, and Paul Block. Unfolding the relationship between seasonal forecast skill and value in hydropower production: a global analysis. *Hydrology and Earth System Sciences*, 26:2431–2448, 2022.

[9] Zhanwei Liu and Xiaogang He. Balancing-oriented hydropower operation makes the clean energy transition more affordable and simultaneously boosts water security. *Nature Water*, 2023.

[10] Álvaro Lorca, Marcel Favereau, and Daniel Olivares. Challenges in the management of hydroelectric generation in power system operations. *Current Sustainable/Renewable Energy Reports*, 7:94–99, 2020.

[11] Midcontinent Independent System Operator. MISO Market Reports Archive. https://docs.misoenergy.org/marketreports/, 2023–2025. Daily Market Operating Margin (MOM) and Locational Marginal Price (LMP) reports. Accessed: 2023-01-01 through 2025-12-02.

[12] Midcontinent Independent System Operator. Locational Marginal Pricing Methodology. Market operations manual, MISO, 2024. Details calculation methodology for Energy Component (MEC), Congestion Component (MCC), and Loss Component (MLC).

[13] NOAA Physical Sciences Laboratory. CPC Global Unified Precipitation Data. https://psl.noaa.gov/data/gridded/data.cpc.globalprecip.html, 2026.

[14] Mohammad Reza Sharifi, Saeid Akbarifard, Mohamad Reza Madadi, Kourosh Qaderi, and Hossein Akbarifard. Optimization of hydropower energy generation by 14 robust evolutionary algorithms. *Scientific Reports*, 12(1):7739, 2022.

[15] Francisca Lanai Ribeiro Torres, Luana Medeiros Marangon Lima, Michelle Simoes Reboita, Anderson Rodrigo de Queiroz, and José Wanderley Marangon Lima. Integrating hydrological and machine learning models for enhanced streamflow forecasting via bayesian model averaging in a hydro-dominant power system. *Water*, 16(4):586, 2024.

[16] U.S. Army Corps of Engineers, Hydrologic Engineering Center (HEC). HEC-HMS: Hydrologic Modeling System. https://www.hec.usace.army.mil/software/hec-hms/, 2026.

[17] U.S. Army Corps of Engineers (USACE). DeGray Dam and Reservoir Water Control Manual. https://water.usace.army.mil/cda/documents/wc/2284/Degray_ResRegManual_1969.pdf, 1969.

[18] U.S. Army Corps of Engineers, Vicksburg District. *Reservoir Regulation Manual: DeGray Dam and Reservoir, Caddo River, Arkansas*. Department of the Army, Vicksburg, Mississippi, September 1969.

[19] U.S. Energy Information Administration. Electricity API: Facility Level Generation. https://api.eia.gov/v2/electricity/facility-fuel/data/, 2025. Monthly net generation by facility and prime mover. Plant Code 187 (DeGray), Prime Mover HY (Hydroelectric). Coverage: 2001-01 through 2025-06. Frequency: monthly. Query parameters: facets[plantCode][]=187, facets[primeMover][]=HY.

[20] U.S. Energy Information Administration. Petroleum & Other Liquids API: WTI Crude Oil Spot Price. https://api.eia.gov/v2/petroleum/pri/spt/data/, 2025. Daily West Texas Intermediate (WTI) crude oil spot prices (dollars per barrel). Series: RWTC. Frequency: daily. Length: 5000 records per query.

[21] U.S. Geological Survey (USGS). The National Map Downloader. https://apps.nationalmap.gov/downloader/, 2026.

[22] U.S. Geological Survey (USGS). USGS 07359610: Streamflow in Caddo River near Caddo Gap, AR (Daily Data). https://waterdata.usgs.gov/monitoring-location/USGS-07359610/#dataTypeId=daily-00060-0&period=periodOfRecord&showFieldMeasurements=true, 2026.

A Operational Reference Plates

The following operational plates and tables are reproduced from the *DeGray Reservoir Regulation Manual* (1969) to support the analysis of reservoir allocations and discharge characteristics.

Table A.1: Stage reductions at Arkadelphia and Camden, Arkansas, attributed to the operation of DeGray, Blakely Mountain, and Narrows Dams during historical floods. [18]

Flood Year	**Location**	**Uncontrolled** (ft)	**Regulated** (ft)	**Reduction** (ft)
1927	Arkadelphia	29.2	24.2	5.0
	Camden	41.9	38.5	3.4
1945	Arkadelphia	30.3	24.3	6.0
	Camden	44.8	43.0	1.8
1958	Arkadelphia	30.2	26.0	4.2
	Camden	44.7	43.0	1.7

Source: Adapted from Table 3, DeGray Reservoir Regulation Manual [18, p. 15].

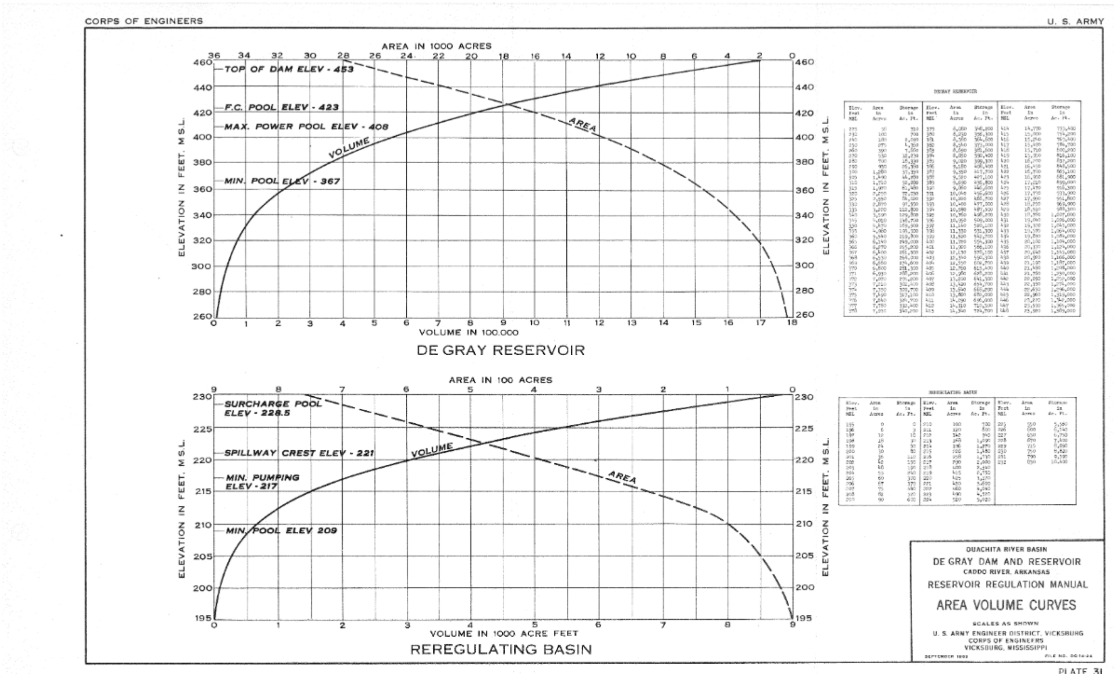

Figure A.1: Area and Capacity Curves for DeGray Reservoir and the Reregulating Basin. This plate defines the storage-elevation relationship for the Minimum, Power, and Flood Control pools. [18]

Source: Plate 31, DeGray Reservoir Regulation Manual [18].

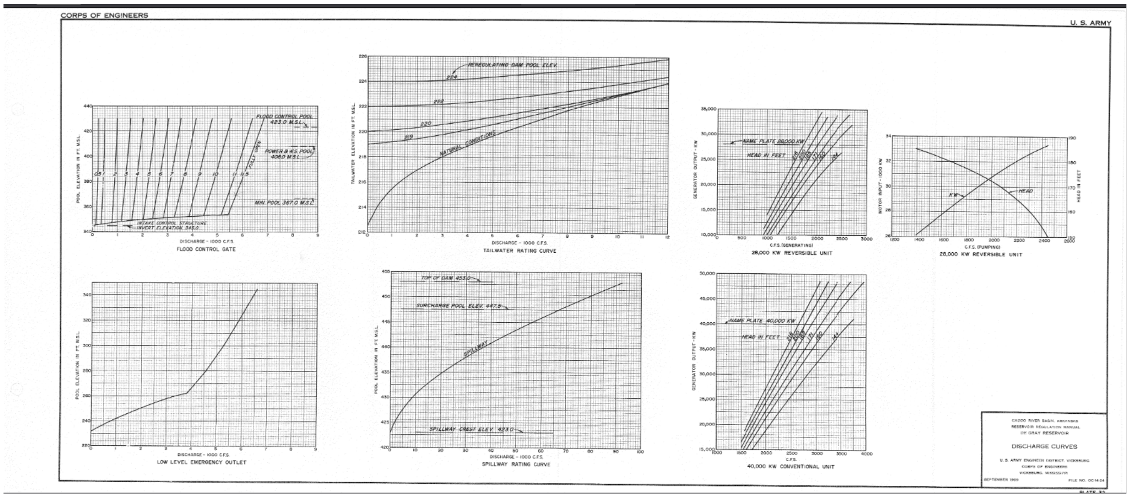

Figure A.2: Discharge capabilities for the Main Dam, including the Intake Control Structure, Flood Control Gate (6' x 11.5'), and the Uncontrolled Spillway. [18]

Source: Plate 32, DeGray Reservoir Regulation Manual [18].

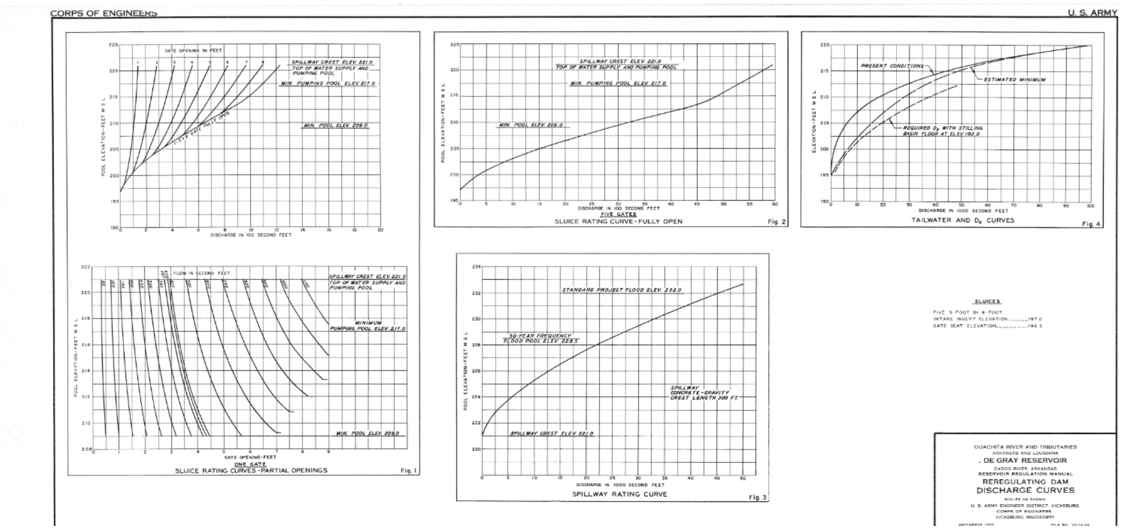

Figure A.3: Rating curves for the Reregulating Dam sluice gates (5' x 9') and spillway, utilized for maintaining continuous water supply releases. [18]

Source: Plate 33, DeGray Reservoir Regulation Manual [18].

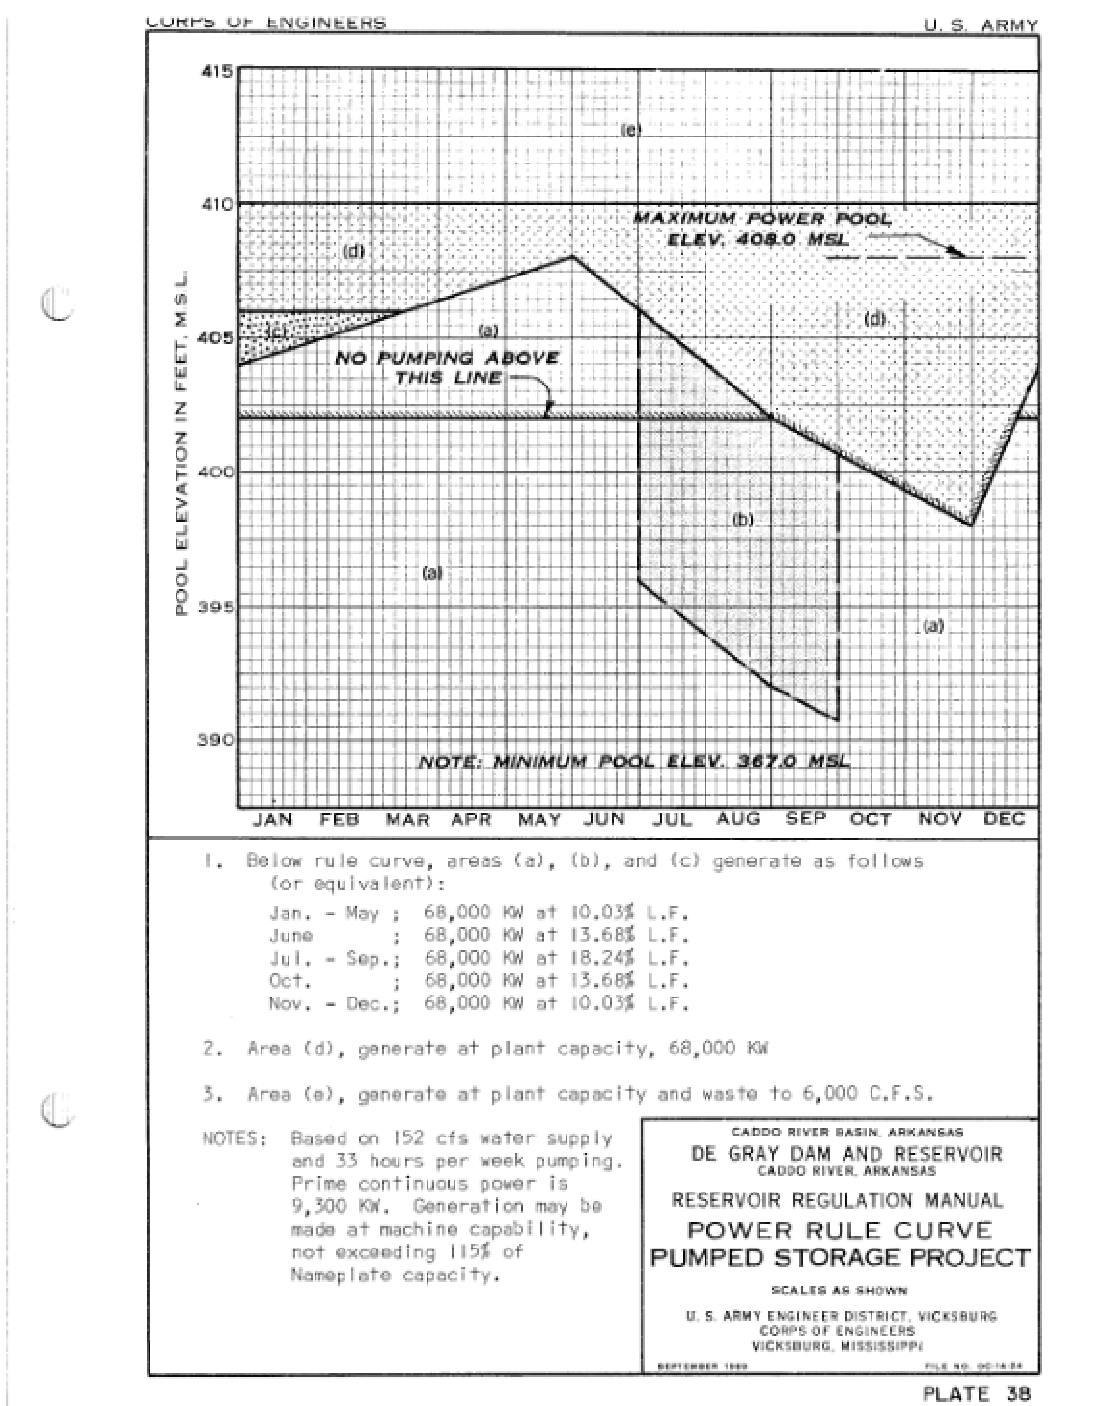

Figure A.4: Power Rule Curve guiding seasonal pool elevations and generation schedules. Areas (a) through (e) dictate generation intensity based on current pool levels. [18]

Source: Plate 38, DeGray Reservoir Regulation Manual [18].

B Data Sources and Processing Documentation

This appendix provides complete documentation of all data sources, retrieval parameters, temporal coverage, and preprocessing decisions used to construct the three analytical datasets (`master_hourly`, `master_daily`, and `master_full`) described in Chapter 5. The goal is to provide sufficient detail for replication using only publicly available resources.

B.1 MISO Market Data

B.1.1 Market Operating Margin (MOM) Reports

Multi-day Operating Margin (MOM) sheets were the primary source of day-ahead forecast predictors. All MOM data were retrieved from the MISO Market Reports Archive portal.
Access: `https://docs.misoenergy.org/marketreports/`
Date range accessed: 2023-01-01 through 2025-12-02
Completeness: 88 of the daily predictor variables had less than 2% missingness; overall dataset completeness 99.8%

B.1.2 Temporal Integrity Enforcement (D-1 Filter)

To prevent data leakage and accurately reflect the information available to market participants prior to the day-ahead submission deadline (09:30 CT), a strict publication date filter was applied to all ex-ante MOM predictors:

$$\text{Publication Date} = \text{Target Date} - 1 \tag{B.1}$$

This ensures that Solar, Wind, and Load forecast values for a given target operating date reflect the MOM sheet published the preceding day, consistent with the 08:00 CT data pull assumed in the operational decision framework.

B.1.3 Locational Marginal Price (LMP) Data

Ex-post LMP data for the DeGray node were retrieved from the MISO historical LMP archive. The LMP is decomposed into three components:

- **MEC** – Marginal Energy Component ($/MWh)

Table B.1: MISO Market Operating Margin (MOM) Report Inventory

Report Name	**Granularity**	**Coverage**	**Key Variables**
Solar Forecast	Hourly	Jan 2023–Sep 2025	Solar generation (MW) by region: North, Central, South, MISO-wide
Wind Forecast	Hourly	Jan 2023–Sep 2025	Wind generation (MW) by region: North, Central, South, MISO-wide
Load Uncertainty	Daily peak	Jan 2023–Sep 2025	Projected load, load uncertainty, obligation (MW)
Outage Outlook (7-day)	Daily peak	Jan 2023–Sep 2025	Committed, uncommitted capacity by region (MW)
Outage Outlook (30-day)	Daily peak	Jan 2023–Sep 2025	*Empty; excluded from analysis*
Resource Operating Margin	Daily peak	Jan 2023–Sep 2025	Resource uncommitted MW, operating margin, emergency headroom
Additional Emergency Headroom	Daily peak	Jan 2023–Sep 2025	Emergency headroom 1–4 hour, 4–8 hour, 8–16 hour windows

- **MCC** – Marginal Congestion Component ($/MWh)
- **MLC** – Marginal Loss Component ($/MWh)

Node identifier: DeGray (MISO node name: EAI.DEGRAY1$_{G}ennode$)
Balancing area:$LRZ8(LoadResourceZone8), MISOSouth$
Granularity:$Hourlyex-postsettlementvalues$
Access:
Date range: 2023-01-01 through 2025-09-04
Records: 23,471 hourly observations (`master_hourly` dataset)

B.2 Commodity Price Data

B.2.1 Natural Gas Spot Price

Source: U.S. Energy Information Administration (EIA)
Series: RNGWHHD "Henry Hub Natural Gas Spot Price"
API endpoint: `https://api.eia.gov/v2/`
Frequency: Daily
Units: $/MMBtu
Lag applied: 1-week lag (aligned to D-1 decision horizon)

Date range: 2023-01-01 through 2025-09-04

B.2.2 Crude Oil Spot Price (WTI)

Source: U.S. Energy Information Administration (EIA)
Series: RWTC (West Texas Intermediate crude oil spot price)
API endpoint: `https://api.eia.gov/v2/petroleum/pri/spt/data/`
Query parameters: Series RWTC; frequency daily; 5,000 records per query
Frequency: Daily
Units: $/barrel
Lag applied: 1-day lag (aligned to D-1 decision horizon)
Date range: 2023-01-01 through 2025-09-04

B.3 DeGray Historical Generation Data

Source: U.S. Energy Information Administration (EIA) Electricity API
API endpoint: `https://api.eia.gov/v2/electricity/facility-fuel/data/`
Plant Code: 187 (DeGray)
Prime Mover: HY (Hydroelectric)
Frequency: Monthly aggregate net generation (MWh)
Coverage: 2001-01 through 2025-06
Usage note: Monthly aggregate only; used for contextual validation and annual generation benchmarking (reported average annual generation ≈ 47,000 MWh). Not used as a model predictor due to aggregation level.

B.4 Physical Model Data

B.4.1 Gridded Precipitation Data

Source: NOAA Physical Sciences Laboratory, Climate Prediction Center (CPC) Global Unified Precipitation
URL: `https://psl.noaa.gov/data/gridded/data.cpc.globalprecip.html`
Product: Daily Total Precip
Spatial resolution: .5 degree latitude x .5 degree longitude global grid
Temporal resolution: Daily
Domain extracted: Contiguous US
Date range used: 2023-01-01 through 2025-09-30
Accessed: 2026-02-16

B.4.2 Streamflow Calibration Data

Source: U.S. Geological Survey (USGS) National Water Information System
Station: Caddo River near Caddo Gap, AR
Station ID: 07356000
URL: `https://waterdata.usgs.gov/`
Parameter: Streamflow, mean daily (00060), ft^3/s
Date range used: 2023-01-01 through 2025-09-30
Usage: HEC-HMS model calibration and inflow prescription
Note: The bibliography reference [22] cites USGS station 07359610.

B.4.3 Terrain Data

Source: U.S. Geological Survey (USGS), The National Map Downloader
URL: `https://apps.nationalmap.gov/downloader/`
Product: 1/3 arc-second DEM Spatial resolution: resampled to 2000 x 2000 M
Usage: HEC-HMS basin delineation and watershed boundary definition
Accessed: 2025-11-13

B.5 Missingness Summary

Table B.2 summarises the missingness profile of the final `master_daily` dataset (978 rows × 163 columns) prior to imputation.

Table B.2: Dataset Missingness Summary by Variable Group

Variable Group	Count	Missing (%)	Imputation
Max Daily LMP (target)	978	2.97% (29 obs.)	Rows excluded from training
MOM daily predictors	88 vars	<2% each	Median imputation
Outage 30-day sheet	–	100%	Excluded entirely

April 2024 gap: A concentrated block of missing values was identified in the target LMP series and several MOM predictor sheets during April 2024. These rows were excluded from model training, but were retained in the time index for continuity assessment.

B.6 Final Dataset Dimensions

File format: rds
Software environment: R 4.5.2; Python 3.11.5, HEC-DSSVue

Table B.3: Final Processed Dataset Specifications

Dataset	**Rows**	**Cols**	**Description**
`master_hourly`	23,471	36	One row per target hour; hourly Solar/Wind forecasts, ex-post LMP, lags, temporal markers
`master_daily`	978	163	One row per target date; all daily MOM sheets, aggregated hourly forecasts, daily LMP targets
`master_full`	23,471	190	Join of `master_hourly` and `master_daily`; daily predictors broadcast to each hour

C Model Specifications and Hyperparameters

This appendix documents the complete specification of both the market prediction models and the HEC-HMS physical simulation, including all hyperparameters, retained predictors, validation schedules, and software dependencies. The intent is to provide sufficient detail for independent replication.

C.1 Software Environment

Table C.1: Software Dependencies

Package / Tool	Version	Language	Usage
R	4.5.2	R	Primary analysis environment
Python	3.11.5	Python	Data preprocessing pipelines
HEC-HMS	4.13	–	Hydrologic simulation

C.2 Target Variable Transformation

The target variable for all market models is `max_daily_LMP_log`, defined as:

$$y_t = \ln\big(\mathrm{MaxLMP}_t\big) \tag{C.1}$$

Table C.2 reproduces the skewness reduction achieved by the log transformation across all candidate target variables.

Table C.2: Effect of Logarithmic Transformation on Target Skewness

Variable	Original Skewness	Log Skewness	Reduction
Max Daily LMP	5.11	1.46	71.5%
Max Daily MEC	4.69	1.33	71.5%
Mean Daily LMP	5.45	1.73	68.3%

All model performance metrics reported in the thesis body are calculated on the log scale unless explicitly noted as "Original Scale ($/MWh)."

C.3 Feature Inventory

C.3.1 Feature Groups and Construction Logic

A total of 112 predictors were retained under the "Strategic Retention" approach described in Section 5.3.2. Table C.3 summarises the feature groups; the full predictor list is given in Table C.4.

Table C.3: Feature Group Summary

Feature Group	Description
MOM Solar forecasts	Daily min/max/mean/sum of hourly ex-ante solar forecasts (MW) for North, Central, South, MISO-wide regions
MOM Wind forecasts	Daily min/max/mean/sum of hourly ex-ante wind forecasts (MW) for North, Central, South, MISO-wide regions
MOM Load / Obligation	Projected load, obligation, load uncertainty, standard deviation of load uncertainty
Uncommitted capacity	Resource uncommitted MW at 4 hr, 8 hr, 12 hr, 16 hr horizons by region
Operating margin	Resource operating margin, additional emergency headroom (1-hr, 4-hr, 8-hr windows)
Outage indicators	Central derates, committed capacity flags
Autoregressive lags (LMP)	$t-1$ (prior day), $t-7$ (prior week), $t-30$ (prior month) of `max_daily_LMP_log`
Autoregressive lags (MCC)	$t-1$, $t-7$, $t-30$ of `max_daily_MCC_log`
Autoregressive lags (MEC)	$t-1$, $t-7$, $t-30$ of `max_daily_MEC_log`
Commodity prices	Natural gas spot price (1-week lag), WTI crude oil (1-day lag)
Cyclical temporal	Sine/cosine encoding of Day-of-Week; Sine/cosine encoding of Month
Volatility regime flag	Binary: 1 if prior-day rolling SD of LMP exceeds 90th percentile
Interaction terms	Physical interaction terms (e.g., Load $\times$ Gas Price, Wind $\times$ Load Uncertainty)
Total retained	**112**

C.3.2 Volatility Regime Flag Details

The "High Volatility" binary flag was constructed as follows:

1. Compute the rolling 30-day standard deviation of `max_daily_LMP_log`.
2. Define the threshold as the 90th percentile of that rolling SD series across the full dataset.
3. Assign flag = 1 on days where the prior day's rolling SD exceeds the threshold.

This flag identified 95 high-volatility days across the full dataset, primarily concentrated between January 2024 and August 2025.

C.3.3 Full Predictor List

Table C.4: Complete List of 112 Retained Predictors

Feature Name	**Type**	**Construction**
Autoregressive Lags		
`max_daily_LMP_log_lag1`	Continuous	y_{t-1}
`max_daily_LMP_log_lag7`	Continuous	y_{t-7}
`max_daily_LMP_log_lag30`	Continuous	y_{t-30}
`max_daily_MCC_log_lag1`	Continuous	MCC_{t-1}
`max_daily_MCC_log_lag7`	Continuous	MCC_{t-7}
`max_daily_MCC_log_lag30`	Continuous	MCC_{t-30}
`max_daily_MEC_log_lag1`	Continuous	MEC_{t-1}
`max_daily_MEC_log_lag7`	Continuous	MEC_{t-7}
`max_daily_MEC_log_lag30`	Continuous	MEC_{t-30}
Cyclical Temporal Encodings		
`dow_sin`	Continuous	$\sin(2\pi \cdot \mathrm{DayOfWeek}/7)$
`dow_cos`	Continuous	$\cos(2\pi \cdot \mathrm{DayOfWeek}/7)$
`month_sin`	Continuous	$\sin(2\pi \cdot \mathrm{Month}/12)$
`month_cos`	Continuous	$\cos(2\pi \cdot \mathrm{Month}/12)$
Volatility / Regime Flag		
`high_volatility_flag`	Binary	1 if rolling-30d SD > 90th pctile
Commodity Prices		
`gas_price_log`	Continuous	ln(Henry Hub spot), lagged 7 days
`oil_price_log`	Continuous	ln(WTI spot), lagged 1 day

Continued on next page...

Table C.4 – continued from previous page

Feature Name	Type	Construction
MOM Solar Forecasts (aggregated to daily)		
South_solar_d_min	Continuous	Daily min of hourly ex-ante South solar (MW)
South_solar_d_max	Continuous	Daily max of hourly ex-ante South solar (MW)
South_solar_d_mean	Continuous	Daily mean of hourly ex-ante South solar (MW)
South_solar_d_sum	Continuous	Daily sum of hourly ex-ante South solar (MW)
Central_solar_d_min	Continuous	same pattern, Central region
Central_solar_d_max	Continuous	same pattern, Central region
Central_solar_d_mean	Continuous	same pattern, Central region
Central_solar_d_sum	Continuous	same pattern, Central region
North_solar_d_min	Continuous	same pattern, North region
North_solar_d_max	Continuous	same pattern, North region
North_solar_d_mean	Continuous	same pattern, North region
North_solar_d_sum	Continuous	same pattern, North region
MISO_solar_d_min	Continuous	same pattern, MISO-wide
MISO_solar_d_max	Continuous	same pattern, MISO-wide
MISO_solar_d_mean	Continuous	same pattern, MISO-wide
MISO_solar_d_sum	Continuous	same pattern, MISO-wide
MOM Wind Forecasts (aggregated to daily)		
South_wind_d_min	Continuous	Daily min of hourly ex-ante South wind (MW)
South_wind_d_max	Continuous	Daily max
South_wind_d_mean	Continuous	Daily mean
South_wind_d_sum	Continuous	Daily sum
	[Repeat for Central, North, MISO-wide wind]	
MOM Load / Capacity Predictors		
Projected_Load_south	Continuous	Projected load (MW), South region
Projected_Load_miso	Continuous	Projected load (MW), MISO-wide
Projected_Load_load_unc	Continuous	Load uncertainty (MW)
Obligation_miso	Continuous	Obligation (MW), MISO-wide

Continued on next page...

Table C.4 – continued from previous page

Feature Name	Type	Construction
Std_Dev_MW_load_unc	Continuous	SD of load uncertainty (MW)
MAE_MW_load_unc	Continuous	MAE of load forecast (MW)
RESOURCE_UNCOMMITTED_south	Continuous	Uncommitted resources at 4-hr horizon (MW)
RESOURCE_UNCOMMITTED_north	Continuous	Same, North region
RESOURCE_UNCOMMITTED_miso	Continuous	Same, MISO-wide
Uncommitted_4_hr_south	Continuous	As observed in MOM
Uncommitted_4_hr_north	Continuous	As observed in MOM
Uncommitted_4_hr_miso	Continuous	As observed in MOM
Uncommitted_8_hr_south	Continuous	As observed in MOM
Uncommitted_8_hr_north	Continuous	As observed in MOM
Uncommitted_8_hr_miso	Continuous	As observed in MOM
Uncommitted_12_hr_south	Continuous	As observed in MOM
Uncommitted_12_hr_miso	Continuous	As observed in MOM
Uncommitted_16_hr_south	Continuous	As observed in MOM
Uncommitted_16_hr_miso	Continuous	As observed in MOM
Resource_Op_Margin_miso	Continuous	Resource operating margin (MW)
RESOURCE_COMMITTED_miso	Continuous	Committed resources (MW)
Central_Derates_outage	Continuous	As observed in MOM
MISO_Derates_outage	Continuous	As observed in MOM
Additional MOM predictors to total 112		

C.4 Ridge Regression Model

C.4.1 Model Specification

The Ridge Regression model minimises the L2-penalised objective:

$$\hat{\boldsymbol{\beta}} = \arg\min_{\boldsymbol{\beta}} \left\{ \sum_{t=1}^{T} (y_t - \mathbf{x}_t^\top \boldsymbol{\beta})^2 + \lambda \|\boldsymbol{\beta}\|_2^2 \right\} \tag{C.2}$$

where λ is the regularisation strength selected via cross-validation.

Table C.5: Ridge Regression Final Hyperparameters

Parameter	Value	Selection Method
Regularisation strength (λ)	0.0201	10-fold CV on development set
Elastic Net mixing (α)	0.0 (pure Ridge)	Grid search over $\alpha \in [0, 1]$; Lasso not beneficial
Predictors retained	112	Strategic retention; all MOM predictors kept
Target variable	`max_daily_LMP_log`	Log-transformed
PCR alternative tested	43 components, RMSE 0.211	Ridge (RMSE 0.207) preferred

C.4.2 Hyperparameters

C.4.3 Key Coefficient Values (Final Production Model)

Table C.6: Top Ridge Regression Coefficients by Absolute Magnitude

Predictor	Coefficient ($\hat{\beta}$)	Interpretation
`gas_price_log`	0.137	Dominant positive driver; marginal fuel cost
`month_cos`	0.130849	Macro-seasonal drift
`MAPE_wind_unc`	0.046341	Higher wind forecast uncertainty $\rightarrow$ higher risk premium
`Standard_Deviation_in_MW_load_unc`	-0.1346140	Negative: high load variance $\rightarrow$ conservative operator commitment, dampened LMP

C.5 XGBoost Model

C.5.1 Hyperparameters

Table C.7: XGBoost Hyperparameters

Parameter	Value	Notes
Maximum tree depth	6	Manually validated
Learning rate (η)	0.3	Manually validated
Cross-validation scheme	3-fold	Development phase 2023–2024 only

C.6 Walk-Forward Validation Schedule

The walk-forward validation simulates operational deployment by retraining both models monthly on all available data prior to each test month. No future data are accessible to the model during any given prediction window.

Table C.8: Walk-Forward Validation Monthly Schedule (2025 Holdout)

Test Month	Training Window Start	Training Window End	Ridge RMSE	XGB RMSE
Jan 2025	2023-01-01	2024-12-31	$\approx$0.40	$\approx$0.58
Feb 2025	2023-01-01	2025-01-31	$\approx$0.33	$\approx$0.33
Mar 2025	2023-01-01	2025-02-28	$\approx$0.22	$\approx$0.22
Apr 2025	2023-01-01	2025-03-31	$\approx$0.20	$\approx$0.21
May 2025	2023-01-01	2025-04-30	$\approx$0.14	$\approx$0.15
Jun 2025	2023-01-01	2025-05-31	$\approx$0.20	$\approx$0.20
Jul 2025	2023-01-01	2025-06-30	$\approx$0.37	$\approx$0.37
Aug 2025	2023-01-01	2025-07-31	$\approx$0.30	$\approx$0.38
Sep 2025	2023-01-01	2025-08-31	$\approx$0.15	$\approx$0.09
Full holdout aggregate (Jan–Sep 2025)			**0.284**	**0.326**

C.7 Baseline Model Specifications

Table C.9: Baseline Model Specifications and Development-Phase Performance

Model	Rule	RMSE (log)	MAE (log)	MAPE (orig.)
Seasonal Naive	$\hat{y}_t = y_{t-7}$	0.397	0.280	31.1%
Persistence	$\hat{y}_t = y_{t-1}$	0.245	0.177	18.6%
Ridge Regression	L2-regularised linear	0.209	0.144	14.2%
Evaluation period: 2023–2024 development set, 10-fold cross-validation				

C.8 HEC-HMS Model Configuration

C.8.1 Model Parameters

Table C.10: HEC-HMS DeGray Lake Model Configuration

Parameter	Value	Source
Simulation period	2025-01-01 to 2025-12-30	Matches Ridge model forecast year
Time step	Daily	USACE regulation manual frequency
Conservation pool bottom	El. 367 ft MSL; $V_{\text{bot}} \approx$ 261,500 ac-ft	USACE Vicksburg District
Conservation pool top	El. 408 ft MSL; $V_{\text{top}} =$ 654,700 ac-ft	USACE regulation manual
Active power pool storage	$\approx$ 393,200 ac-ft	$V_{\text{top}} - V_{\text{bot}}$
Energy storage penalty threshold	50% of active pool	Defined in objective function
Inflow source	Observed streamflow, Caddo River headwater gage	USGS
Evaporation	Monthly pan evaporation coefficients	USACE regulation manual
Min. env. release (Scenario A)	157 CFS	USACE minimum flow requirement
Flood control trigger	Required outflow >6,000 CFS; maintain pool <El. 423 ft MSL	USACE regulation manual
Scenario B dispatch trigger	Predicted LMP >$65.45/MWh (75th percentile)	Ridge model Q4 threshold
Scenario B dispatch duration	4 hours at full gate	Operational assumption
Scenario B max turbine release	6,000 CFS (combined)	USACE nameplate
Scenario B daily avg. outflow	$\approx$1,131 CFS	Eq. 7.2

C.9 Objective Function Parameter Values

Table C.11: Objective Function Parameters: Specified and Illustrative Values

Parameter	Value Used	Units	Notes
λ_1 (Reward weight)	0.60	–	Base case; sensitivity range 0–1
λ_2 (risk weight)	0.40	–	$\lambda_1 + \lambda_2 = 1$
λ_1^* (crossover point)	≈0.25	–	Analytical result; Eq. 8.7
Energy storage threshold	50%	% of active pool	Soft constraint trigger level
α_L (low-storage penalty)	50	$\$/(\%\,\mathrm{pp})^2$	Illustrative; requires operational calibration
α_H (flood penalty)	N/A (never triggered)	$\$/\mathrm{m}^2$	$F_{\mathrm{Flood}} = 0$ for both scenarios
β_Q (min-flow penalty)	N/A (never triggered)	$\$/(\mathrm{m}^3/\mathrm{s})^2$	Environmental flow met by design
Installed capacity ($P_{\mathrm{installed}}$)	74 MW (combined)	MW	Nameplate; Max approx. 38 MW single-unit
Dispatch window (t_{dispatch})	4 hours	hours	Operational assumption
LMP peak-hour multiplier	0.80	–	Conservative approximation; Eq. 8.4
Q4 dispatch threshold	$65.45/MWh	$/MWh	75th percentile of 2023–2024 LMP distribution