NYISO

Zonal Merchant Forecast

December 2025

noreva

Noreva Merchant Forecasting

Noreva produces long-horizon merchant power price forecasts for hourly deliveries across NYISO Zones A-K through 2050, grounded in the firm's comprehensive price-formation modeling framework. These curves are published on the Noreva Hub and undergo full methodological updates twice per year, with interim revisions released when significant market developments warrant refreshed forward views. This document outlines the forecasting methodology applied to the NYISO merchant curve as of December 2025, detailing the analytical structure, data inputs, and modeling components that underpin Noreva's long-term outlook.

Executive Summary

NYISO is in transition. Renewable buildout, rising electrification, new transmission investments, and increasingly volatile fuel markets are changing the dynamics of price formation. Capacity markets, environmental attributes (RECs), and energy prices - once modeled independently - are now linked through a grid that is becoming more weather-sensitive and capacity-constrained.

With trading insight and market experience, our model captures this nuance. Traditional statistical models struggle as historical patterns break down, while pure fundamental models often miss short-term price behavior that still drives near-term economics. As renewables flatten midday prices but deepen ramps, load growth introduces sharper peaks, and gas–power coupling remains a dominant driver of extreme events, volatility will persist.

Noreva's hybrid model combines machine learning and economic dispatch fundamentals to construct hourly LBMP forecasts across all eleven NYISO zones through 2050. We evaluate three futures by varying the underlying assumptions that drive NYISO price formation. Each case applies a different combination of load growth, gas price trajectories, renewable buildout pace, and transmission or import availability - shaping the balance between thermal generation, renewable penetration, and system stress. Constrained transmission or tighter fuel markets create more scarcity conditions and higher peak pricing, while stronger efficiency gains, lower fuel costs, or accelerated renewable additions lead to flatter dispatch curves and fewer high-priced hours.

Figure 1: Zone J Monthly Averages

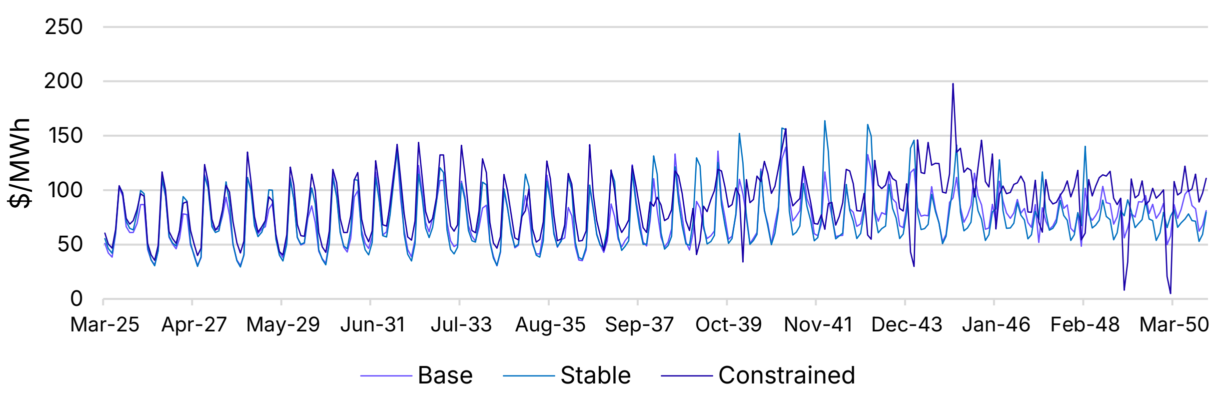

Our model is designed to support investment decisions, hedging, asset valuation, and policy analysis in a grid where fundamentals and price patterns develop at different speeds. Our process is customizable -s we adapt to client needs, test scenarios, and iteratively improve our methodology as markets evolve.

Methodology

The forecasting methodology combines historical data analysis, machine-learning models, and a fundamentals-based economic-dispatch into long-horizon, zone-specific power price scenarios. The process begins by generating consistent hourly forecasts for load, temperature, gas prices, and renewable output. These component forecasts feed into two complementary price-formation models: a short-term ML model that captures recent behavioral dynamics and a long-term stack model that reflects structural system conditions, generator costs, and transmission constraints. The final price trajectory is produced by blending these two approaches with time-varying weights, ensuring near-term responsiveness while maintaining long-term coherence with fundamentals.

Load Component Forecast

Our load-forecasting module produces hourly demand profiles that match both long-term annual energy expectations and seasonal peak targets while preserving the characteristic shapes observed in recent years. The framework begins by analyzing historical load data to extract stable zonal monthly energy shares and hour-of-day patterns for weekdays and weekends. These patterns serve as the structural backbone of each future year's forecast. Annual energy, summer peaks, and winter peaks sourced from the NYISO Gold Book are then allocated across months, days, and hours using these historical distributions. The model iteratively rescales the resulting hourly profiles so that each forecasted year hits its specified annual energy target and achieves the desired summer and winter peak magnitudes.

Beyond the baseline approach, the framework can also incorporate long-term seasonal evolution driven by electrification or other structural changes. When enabled, monthly energy shares shift gradually from a summer-peaking profile to a winter-dominant pattern, guided by smooth transitions centered around a configurable inflection year. This allows the model to reflect emerging trends - such as heat-pump adoption or electrified heating - without discarding the realistic intraday and intra-seasonal structure derived from history. The final output is a zone-specific, hourly load forecast spanning multiple decades, aligned with the broader scenario assumptions.

Figure 2: Seasonal Load Patterns

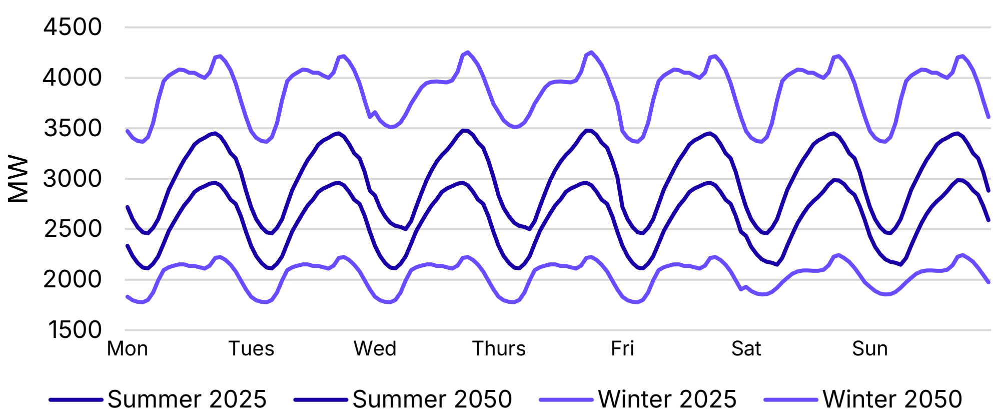

Temperature Component Forecast

We model synthetic future temperature trajectories that preserve the statistical behavior of historical data while allowing for climate-aligned adjustments and controlled stress events. The process begins by decomposing historical NOAA temperature records into a climatological baseline - average conditions by month and hour - and the residual variability around that baseline. By resampling these residuals within each month and recombining them with the climatology, the model produces future temperature profiles that maintain realistic intraday and seasonal patterns. Optional scenario-level adjustments, such as warming offsets, can be applied to reflect long-term climate trends or user-defined sensitivities.

Gas Component Forecast

Our gas-price scenarios provide both synthetic and market-aligned pathways for future fuel cost assumptions. The synthetic approach blends long-term annual price targets with historically grounded day-to-day volatility patterns. Prices revert gradually toward the annual trajectory while sampling month-specific historical volatility, ensuring realistic fluctuations without allowing runaway drift. In the market-based pathway, monthly futures prices are mapped directly into an hourly time series, preserving the forward curve structure without adding short-term noise.

To reflect geographic differences in gas supply costs, the model applies regional basis adjustments that vary seasonally and by pipeline. Each regional series is produced by modifying the underlying Henry Hub path with a zone-specific mean basis and month-by-month seasonal multipliers, resulting in realistic spatial price separation across the system to ensure that fuel-cost dynamics are consistently represented across the entire forecasting stack.

Renewable Generation Component Forecast

The renewable generation model begins by calculating hourly capacity factors from the most recent full year of generation data, capturing how production typically varies by month and time of day. Future capacity additions are then layered onto this framework by applying the historical patterns to projected build-outs, ensuring that growth in installed wind and solar translates directly into realistic hourly output expectations.

To maintain spatial accuracy, the model distributes total renewable capacity across zones using historical shares, reflecting where resources are concentrated and how they contribute to zonal dynamics. Behind-the-meter solar is excluded from capacity-factor calculations to prevent double-counting and maintain consistency with load forecasts. The result is a zone-specific, hourly renewable dataset that aligns with both historical behavior and planned future development.

Figure 3: Daily Renewable Shape by Zone

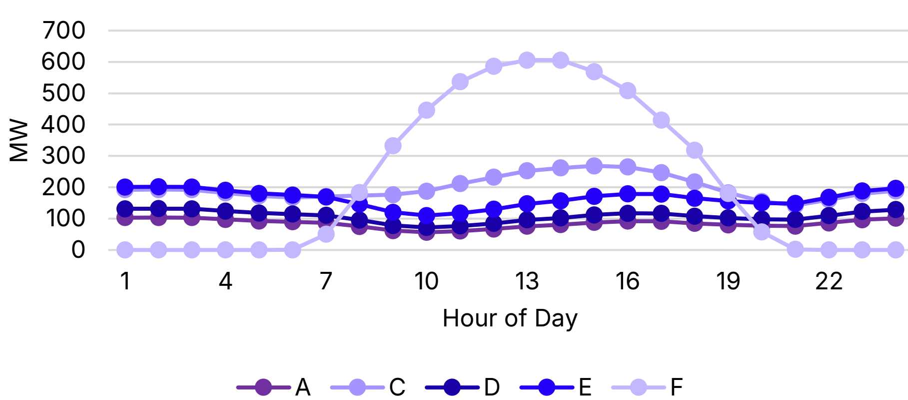

Dispatch-driven Price Forecast

Our economic-dispatch engine simulates hourly power-system operations by combining load, renewable output, fuel prices, and an evolving generation fleet into a unified merit-order framework. For each hour, the model calculates net demand in each zone, incorporates available external imports, and evaluates the marginal operating cost of every generator based on heat rate, fuel price, and variable O&M. Generators are then stacked from lowest to highest marginal cost, with must-run resources dispatched first, followed by economic units until system demand is met. The marginal unit sets the system clearing price, with scarcity adders applied in periods when net load approaches or exceeds available capacity. This same structure applies across historical, current, and future fleets, allowing the model to reflect retirements, new builds, and transmission upgrades such as CHPE.

Figure 4: Import and Export Flows for 2024

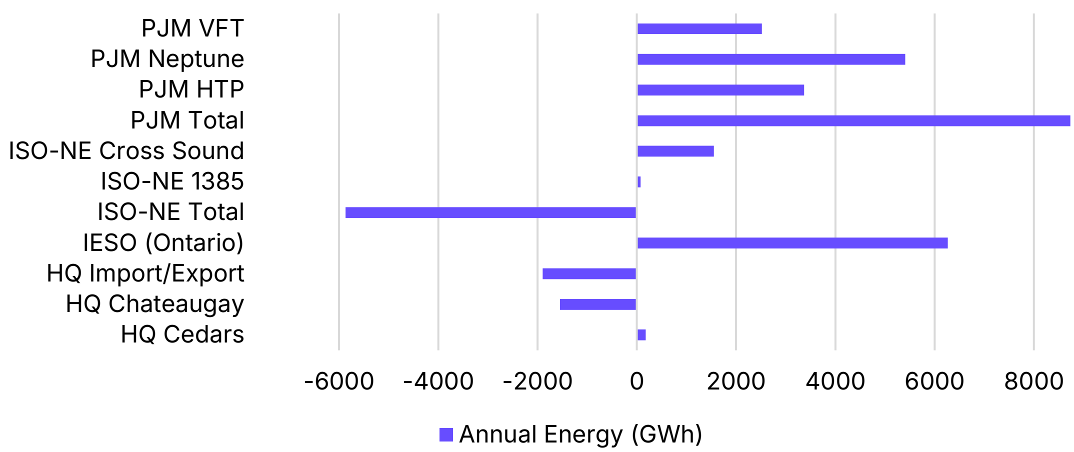

Once the system lambda is determined, the model derives zonal prices by applying congestion factors and historical basis adjustments to capture structural differences across zones. External imports follow a seasonal and time-of-day profile calibrated from historical flow data, ensuring realistic variation in available transfer capability. The full time-series dispatch routine processes each year's fleet independently and outputs hourly price signals, zonal generation levels, scarcity events, and congestion effects.

Machine-Learning-driven Price Forecast

Our ML price-forecasting framework uses a regularized regression approach to learn relationships between historical market conditions and observed energy prices. The model blends standard time-based indicators with engineered features that capture load dynamics, temperature effects, and known seasonal patterns. To represent these recurring patterns explicitly, the framework employs Fourier features that encode annual, weekly, and daily cycles, enabling the model to recognize recurring variations such as summer peaks, weekday - weekend differences, and intraday ramps. Historical data - load, temperature, and observed LBMPs - forms the training set from which a Ridge regression model learns the underlying structure while avoiding overfitting through regularization.

Once trained, the model applies the same feature construction to future conditions, generating forward-looking price predictions based on projected load, temperature, and fuel-price trends.

Blended Result

The blended-forecasting module merges the strengths of the machine-learning model and the economic-dispatch stack model by assigning time-varying weights based on forecast horizon. In the near term the ML model receives more weight because it best captures recent behavioral patterns and short-run dynamics. Beyond that window, the weighting gradually shifts toward the fundamentals-driven stack model, reaching a full transition by 2040 to reflect the greater reliability of structural modeling over longer horizons. The system generates both forecasts independently, calculates the number of days ahead for each hour, applies the horizon-based weighting function, and produces a single blended price series. This approach ensures smooth, intuitive transitions between short-term signal-driven accuracy and long-term fundamental consistency.

Figure 5: Zone J Annualized Prices

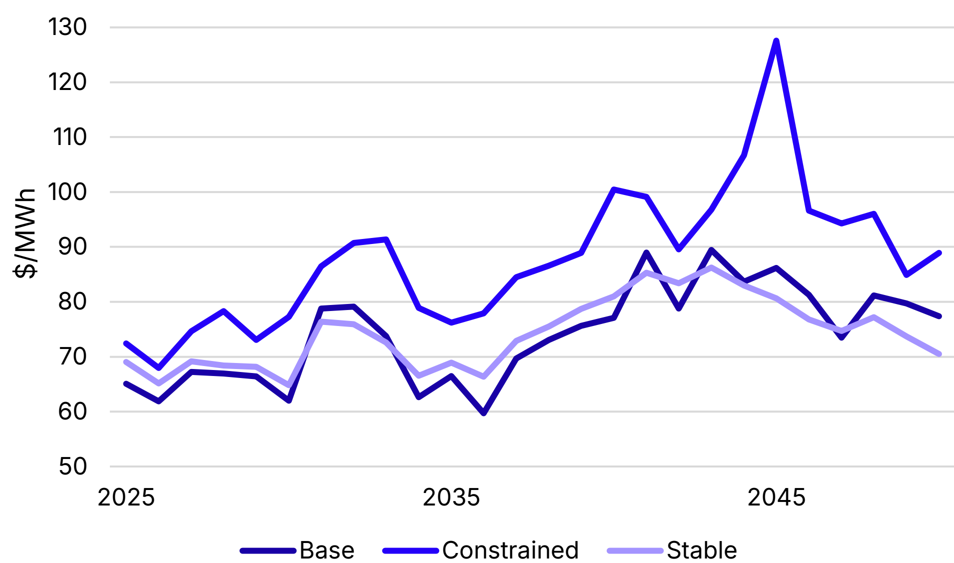

Configuration and Parameters

The forecasting platform is highly customizable, with all major behavioral assumptions, scenario definitions, and model parameters controlled through configuration settings. Each component - load, gas, temperature, renewables, and the economic-dispatch stack - relies on anchors to define long-term trajectories, scenario variants, and structural constraints. Adjustable parameters include annual energy targets, seasonal peak magnitudes, capacity expansions, fuel-price paths, and climate assumptions. The configuration layer acts as the bridge between model mechanics and market outlook: it governs the shape of inputs while keeping the underlying algorithms stable and transparent.

We model three scenarios, stable, base, and constrained, to reflect scarcity pricing parameters, congestion methodology, zonal separation factors, transmission limits under baseline and alternative cases, import limits, and the commissioning of major upgrades such as CHPE. Adjusting these values allows users to reflect policy changes, new transmission projects, or alternative market conditions.

Dimension	Stable	Base	Constrained
Gas Price	Settled Futures	Baseline	Aggressive
Load Growth	Moderate	Baseline	Aggressive
Renewables	Aggressive	Baseline	Moderate
Temperature	Cooler	Baseline	Hotter
Transmission	Improved	Baseline	Baseline

Contact:

Malachi Coleman - Associate
malachi@noreva.ai

Commercial
sales@noreva.ai

Noreva
675 Third Ave, 31st Floor
NY, NY 10017

noreva