memo

**To:** NEPOOL Markets Committee

**From:** Markets Development

**Date:** October 9, 2018

**Subject:** Natural Gas Price Forecast Method for Energy Market Opportunity Costs

## Goal of this Memo

As part of the ISO's efforts to improve the treatment of energy market opportunity costs in advance of this coming winter, the ISO will employ a short-term gas price forecasting model. This memorandum explains how this forecasting model works. The forecasts generated by this model will be used to calculate opportunity cost adjustments to individual generators' reference prices in the energy market mitigation process.

Both electricity and gas price forecasts are a necessary input to determine fuel-related energy market opportunity costs for generation resources with inter-temporal production limitations. The ISO will procure short-term (6 day forward) proprietary forecasts of electricity prices from an external vendor. However, we have not identified a vendor of short-term gas prices (at New England delivery points) that can supply this information daily. Accordingly, to implement the opportunity cost enhancements in a timely manner for this coming winter, the ISO has developed the short-term gas price forecasting model described in this memo.

In sum, the ISO has determined that a neural network forecasting technique using past gas prices, past and short-term forecast weather data, and past and short-term forecast electricity prices will provide the most accurate results (among several alternative techniques studied.) When tested on winter 2017/2018 data, the average and median absolute errors of the day-ahead gas price forecast are $2.48/MMBtu (or 19.43%) and $0.83/MMBtu (12.51%), respectively. For the six-day forecast, the average and median absolute errors for the same period are $5.68/MMBtu (74.46%) and $3.79/MMBtu (59.05%), respectively. As is typical with time series forecasts for commodities, the model underestimates extreme spikes in prices.[1] Nonetheless, the model performs well overall and provides acceptable results for its intended purpose.

The remainder of this memo provides a brief description of the neural network technique employed by the model, an overview of the data used to generate the forecasts, and summary statistics and figures demonstrating the performance of the gas forecasting model. The appendix provides additional details about the model. To facilitate transparency, the appendix was designed to allow reproduction of the forecast given similar inputs.

The ISO is open to feedback on technical methods that may be used to improve the forecast.

[1] Note: the marginal hour sets the opportunity cost for a generator. As such, if the forecast underestimates an extreme price spike, this may not impact the opportunity cost of a generator, so long as the forecast predicts a sufficient portion of the spike.

**ISO New England Inc.**
One Sullivan Road
Holyoke, MA 01040-2841
413-535-4000



**iso-ne.com**
isonewswire.com
@isonewengland
iso-ne.com/isotogo
iso-ne.com/isoexpress

## Neural Networks: Brief Introduction

Neural networks were designed to simulate how human brains learn from information over time. More specifically, neural networks are composed of "neurons" which can "fire", depending on the values of predictive inputs and the past values of the target variable, natural gas price. The value from these neurons is weighted, and these weighted values are given to an "activation" function which determines if the neuron should fire. When neurons fire, this provides information to the model which is used to generate a forecast. The training process takes historical data and iterates through a sequence of possible weights on the neurons until the error of the network is minimized. Neural networks provide increased flexibility compared to traditional time series models, particularly for highly nonlinear or irregular (volatile) data.

## Gas Price Data

The gas price model forecasts the price of Algonquin pipeline natural gas (AGT-CG, non-G), which trades Monday through Friday. The source historical gas pricing data used to train the model is the volume-weighted day-ahead gas price index compiled by ICE and provided by ICE to the ISO.

The trading on a given day is for gas consumed on the next day, with the exception of Friday, which trades for gas consumed Saturday through Monday. Indeed, no observed trades (in the data from ICE) occur on Saturday or Sunday, so the Saturday price, determined on Friday, is the gas price for Saturday-Monday. (If there is a holiday, such as Christmas or Veterans Day, the gas price remains constant for an additional day as well.) Trading continues on Monday, with trades setting the price for Tuesday. When the model runs on a given morning, the gas price for that day and all previous days will be observed. So, on that given morning, the ISO will be forecasting gas prices for the next day and the following five days.

Following industry best practices, the ISO forecasts the natural logarithm of the ICE Algonquin gas price index, denoted as log(Price).

The gas-price forecasting model uses slightly different inputs based on the number of days ahead the forecast is being prepared for:

- The day-ahead gas price forecast model uses three days of gas prices for prediction: the price for today, the price yesterday, and the price the day before yesterday.
- For the two-day forecast, the neural network uses the price forecast for tomorrow, the price today, and the price yesterday.
- For the three-day forecast, the neural network uses the price of the current day, the price forecast for the day after the current day, and the price forecast for two days from the current day.
- And so on, for each of the four-day through the six-day forecasts.

### Weather Data – Mean Temperature

Mean daily temperature is the only weather variable used by the neural network. The ISO tested the predictive power of heating degree days, cooling degree days, dew point, and wind speed but found them to be insignificant predictors of gas price, at least during three winter months (December, January, and February), and after controlling for mean temperature.

The day-ahead (DA) gas price forecast model includes four days of mean temperature data: the forecasted mean temperature for tomorrow, the forecasted mean temperature today, and the observed mean temperature for the previous two days. For the two-day forecast, the mean temperature today, tomorrow, and in two days are all forecasts, while the mean temperature yesterday is observed. And so forth.

The mean temperature forecast comes from the ISO's composite weather forecast used to forecast system load.

### Daily Electricity Price Data – Mean Hub LMP

The neural network also uses forecast and actual mean daily electricity hub locational marginal price (LMP) as a predictor for natural gas price. As noted, the ISO will obtain proprietary day-ahead through six-day LMP forecasts from an outside vendor. Much like the mean temperature data, the neural network uses four days of mean LMP data.

For the day-ahead model, the neural network uses the outside vendor's day-ahead LMP forecast for the forecasted day, yesterday's day-ahead prices for today for the current day's average LMP, and the observed LMP for yesterday and the day before. For the two-day forecast, the neural network will use the vendor's LMP forecast for tomorrow and two days from now, yesterday's day-ahead forecast for today, and the observed mean LMP for yesterday. And so forth.

### Neural Network Performance

Table 1 provides the mean absolute error (MAE) for the gas price forecast and other comparable forecasts. The errors for mean temperature forecasts are included, as well, to indicate their variance overtime. The column labeled as a "naïve model" assumes each forecast is simply the last observed price: the day-ahead forecast is the current day's price, the two-day forecast is also the current day's gas price, etc. Naïve models constructed this way are a standard initial benchmark for assessing a forecasting model's performance.

Compared to the naïve model, the neural network is superior overall, but has a particular advantage at predicting large spikes in the gas price. See Figure 1 in the Appendix. This advantage, however, declines overtime. As the neural network predicts further into the future, its inputs become more prone to error, which reduces the accuracy of its forecast.

More specifically, because the forecast relies on previous values of gas price forecast to predict future values, errors in the forecast for the day-ahead price compound the error in the two-day forecast. This error is further compounded in the three-day forecast, etc. Columns 5 and 6 of Table 1 show how the errors in the inputs increase overtime, as well.

To obtain the model performance results shown in Table 1, the ISO use the standard "split-sample technique". This technique uses several years of historical data to train, or "fit", the forecasting model. Then the model is applied, as a forecast, to predict the actual prices – day by day – observed during a 'test sample' period (approximately 15% of the seven years of data available). The accuracy of the forecasting process, as applied to that winter 2017/2018 test sample data, is summarized in Tables 1 and 2 below.

Given these caveats, the neural network method appears to perform acceptably well and better than presently available alternatives.

Table 1: Mean Absolute Error

| **Days Ahead** | **Neural Network** ($/MMBtu) | **Naïve Model** ($/MMBtu) | **Vendor Electricity LMP**[2] ($/MWh) | **Mean Temperature** (°F) |
|---|---|---|---|---|
| 1 | $2.29 | $3.33 | $19.01 | 1.35°F |
| 2 | $3.10 | $4.72 | $24.96 | 1.41°F |
| 3 | $3.58 | $5.29 | $28.02 | 1.70°F |
| 4 | $4.60 | $5.79 | $31.88 | 2.28°F |
| 5 | $5.51 | $5.96 | $36.06 | 3.28°F |
| 6 | $5.93 | $6.37 | $38.80 | 3.83°F |

Table 2: Mean Absolute Percent Error

| **Days Ahead** | **Neural Network** | **Naïve Model** | **Vendor Electricity LMP** | **Mean Temperature** |
|---|---|---|---|---|
| 1 | 19.43% | 26.89% | 20.29% | 7.10% |
| 2 | 29.04% | 42.83% | 29.78% | 7.48% |
| 3 | 38.45% | 53.64% | 29.27% | 8.19% |
| 4 | 53.60% | 63.47% | 32.22% | 10.15% |
| 5 | 66.12% | 68.84% | 38.44% | 14.97% |
| 6 | 74.46% | 81.21% | 45.53% | 18.21% |

[2] The Mean Absolute Error data for the vendor's electricity LMP forecast is for their hourly predictions between 12/25/2017 and 1/20/2018.

## Appendix

The first section of this appendix describes the neural network in detail and the  the structure of the model. The second section provides graphical evidence of the neural network's performance. The final section discusses the performance of a third forecasting technique, an Autoregressive Integreted Moving Average model with exogenous inputs (ARIMAX) which the ISO also evaluated.

## Appendix A - Neural Network Model Details

The neural network uses electricity LMP and mean temperature as inputs and previous natural gas prices as feedback to the model. The model includes three lags for both the inputs and the feedback, while the inputs also include data from the forecasted day. For example, the model includes forecasted mean temperature for the forecasted day. The ISO decided to use three lags after testing the performance of the neural network with fewer lags, which resulted in greater prediction error, and more lags, which failed to improve the performance of the network. This analysis was confirmed by testing the statistical significance of the correlation of the lag variables (*viz*. the partial autocorrelation function) at higher lags.

The ISO's neural network model uses two so-called "hidden" layers ("Hidden layers" is a term-of-art in neural network model development). The choice of hidden layers contributes to how well the model fits the training data. Too many hidden layers will cause the model to over-fit itself to the training data, while an insufficient number of hidden layers will lead to underperformance. To determine the optimal number of hidden layers, the ISO created and tested models with one through ten hidden layers, and found that the model which most accurately predicted the out-of-sample winter 2017/2018 data had two hidden layers.

Multiple training functions are available for neural networks. The ISO used Bayesian Regularization Backpropagation (BRB), which takes longer to run than other methods but is more effective for highly nonlinear data.

To train the neural network, the ISO has weather data, gas prices, and electricity LMPs going back to December 1, 2011. Because this model is concerned with gas prices during winter months, only data from December, January, and February is used. Given that data from winter 2017/2018 is used to test the model, only data from December 1, 2011 to February 28, 2017 is used to train the model, for a total of six winters.

Of the six winters of data available, 85% of the data is used for training and the remaining 15% is used for validation.[3] Validation is used to ensure the weights measured during training will generalize to data not used for training. The data chosen for validation/training is randomized, but, following best practices, the ISO chose to divide the training and validation data into blocks; that is, all of the validation data consists of consecutive days.

Recall that the gas price on Saturday is the same as the gas price on Sunday and Monday. As a result, the day-ahead or two-day forecast created on Saturday for Sunday and Monday, respectively, should be the observed price on Saturday. Similarly, for the day-ahead forecast created on Sunday, the observed price

[3] No data is reserved for testing during design as we keep all of winter 2017/2018 for that purpose.

on Sunday should be used as the forecast price for Monday rather than a forecast generated by the neural network. This logic holds for other holidays as well.

The neural network forecasts are iterative: the day-ahead forecast is created first so that it can be used as feedback for the two-day forecast, followed by the two-day forecast so that it can be used as feedback for the three-day forecast, etc. This is equivalent to multi-step forecasts using a “closed” neural network.

## Appendix B - Neural Network Performance

Figures 1 through 6 below provide the observed gas price, the neural net’s forecast prices, and the neural net’s error for the day-ahead gas price forecast through the six-day gas price forecast. The forecasts provide visual affirmation of the statistics in Table 1: the day-ahead forecast predicts spikes in price with greater accuracy than a naïve model, but this advantage fades as the forecast date becomes further in the future from the current date. Note, the ISO also tested the neural network on data from winter 2016/2017 and obtained similar results.

Last, a technical note on the statisitcal error process related to Figures 6, 7, and 8. Figures 7 and 8 provide the error autocorrelations and input-error autocorrelations for various lags. Note that none of the autocorrelations/correlations are statistically different from zero at the 95% significance level. Given Figures 7 and 8, the apparent error autocorrelations in Figure 6 for the six-day forecast is likely attributable to the large lag between the last observed six-day error and the forecast date. For example, on a given date, the last observable six-day error is for the six-day gas price forecast for the price today, made six days ago. Thus, there is a twelve day gap between the day the forecast was made for the last observed six-day error and the forecast day for the six-day forecast being made today. This gap limits the predictive value of the last observed errors as an iterative input into the forecasting process.

Figure 1: Neural Network Day-Ahead Performance, Winter 2017/2018

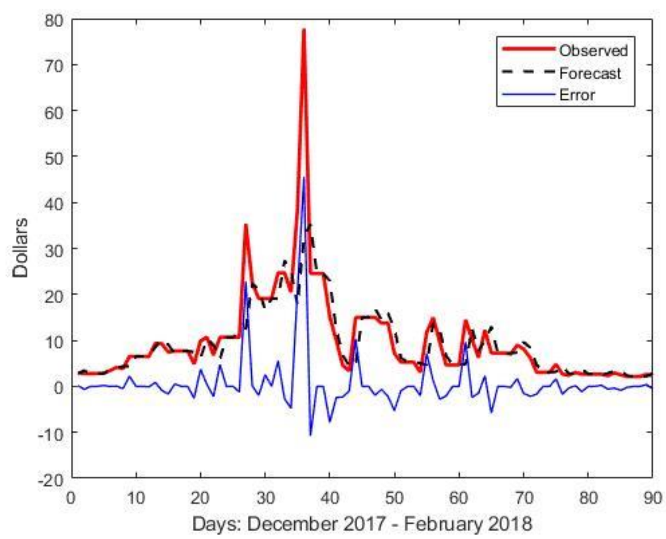


Figure 2: Neural Network Two-Day Performance, Winter 2017/2018

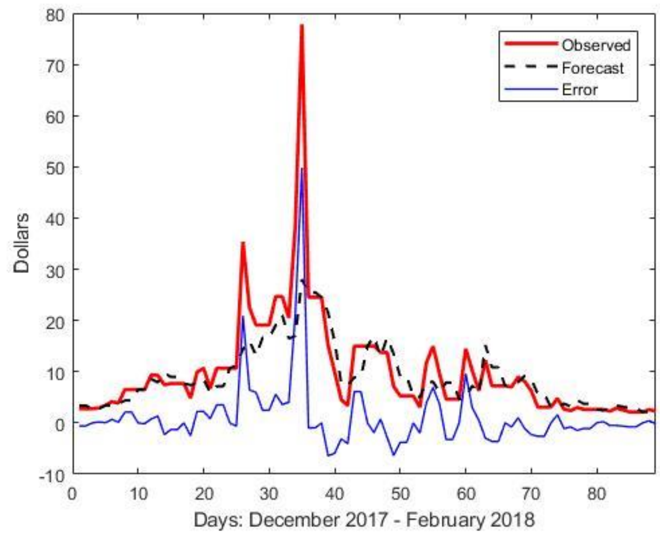


Figure 3: Neural Network Three-Day Performance, Winter 2017/2018

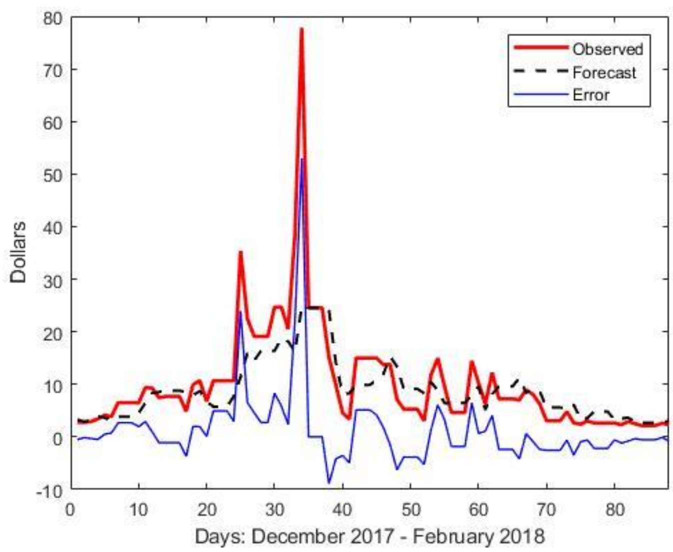


Figure 4: Neural Network Four-Day Performance, Winter 2017/2018

**ISO New England Inc.**
One Sullivan Road
Holyoke, MA 01040-2841
413-535-4000



**iso-ne.com**
isonewswire.com
@isonewengland
iso-ne.com/isotogo
iso-ne.com/isoexpress

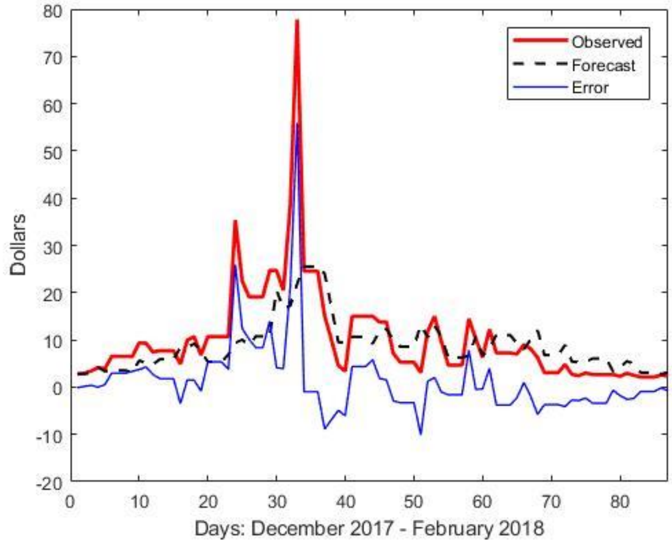


Figure 5: Neural Network Five-Day Performance, Winter 2017/2018

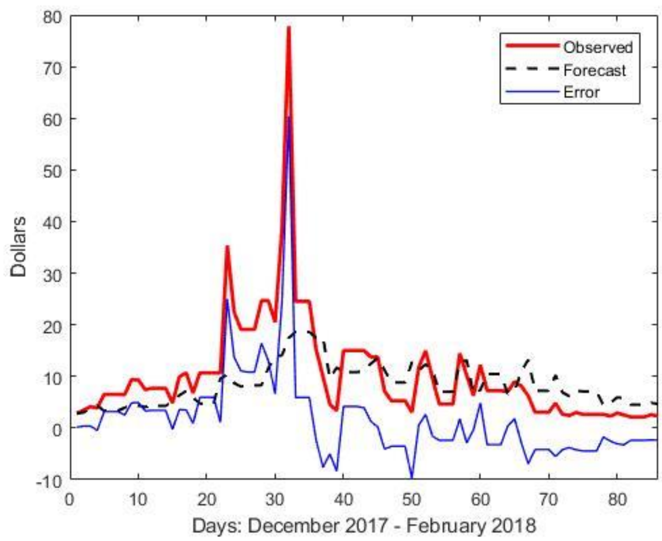


Figure 6: Neural Network Six-Day Performance, Winter 2017/2018

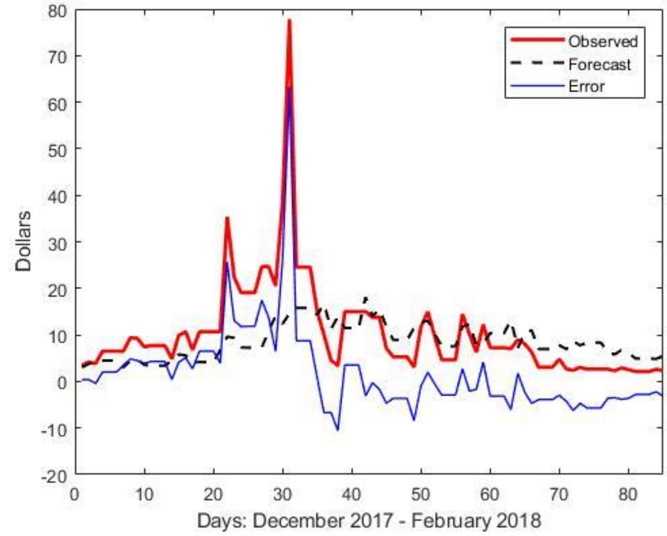


Figure 7: Neural Network Error Autocorrelation, Winter 2017/2018

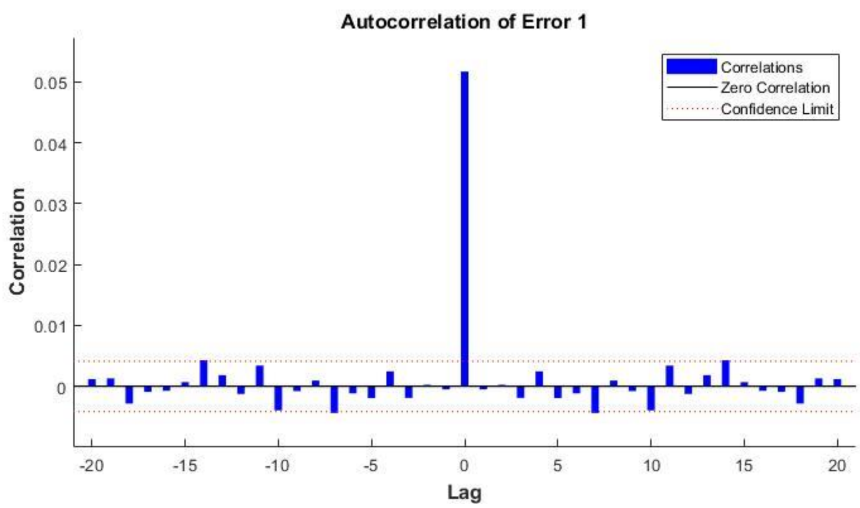


Figure 8: Neural Network Input, Error Correlation

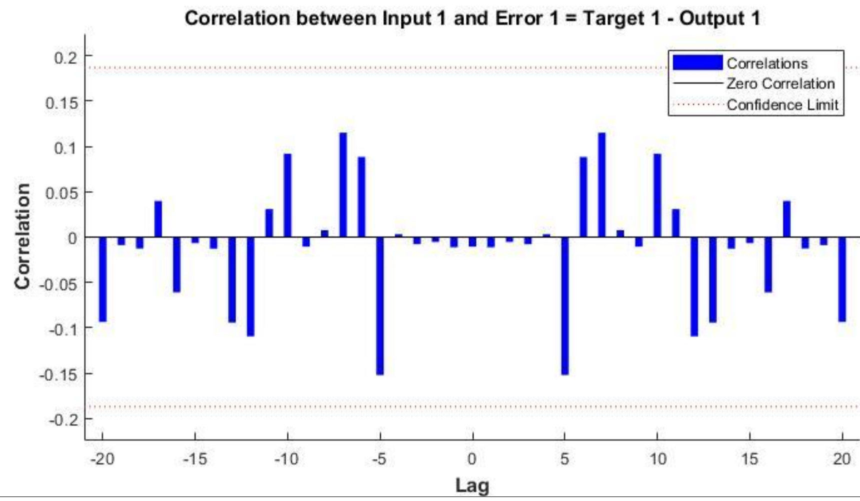


## Appendix C- ARIMAX Alternative

In addition to the neural network, the ISO also created an autoregressive integrated moving average model with exogenous inputs (ARIMAX) as an alternative. Like the neural network, the ARIMAX model included mean temperature and electricity LMPs as inputs, and, like the neural network, the ARIMAX model included three lags for each of these inputs and feedbacks. See Table 3 for a comparison of the mean absolute errors for the neural network and ARIMAX models.

Table 3: Mean Absolute Error – ARIMAX and Neural Network

| **Days Ahead** | **Neural Network** ($/MMBtu) | **Naïve Model** ($/MMBtu) | **ARIMAX** ($/MMBtu) |
|---|---|---|---|
| 1 | $2.29 | $3.33 | $2.62 |
| 2 | $3.10 | $4.72 | $3.15 |
| 3 | $3.58 | $5.29 | $4.03 |
| 4 | $4.60 | $5.79 | $5.16 |
| 5 | $5.51 | $5.96 | $6.23 |
| 6 | $5.93 | $6.37 | $6.53 |

The ISO also considered a forecast composed of a weighted average of the neural network and the ARIMAX model. This composite forecast, however, failed to provide superior predictions, compared to the neural network alone, when tested on the winter 2017/2018 data.