Quantitative Modeling: Refining the Approach

Improving forecast results and model performance.

Morningstar Commodities Research
31 January 2019

Matthew Hong
Director of Research, Power and Gas
+1 312 244-7649
matthew.hong@morningstar.com

Michael O'Leary, Ph.D.
Senior Quantitative Analyst
+1 312 244-7488
michael.oleary@morningstar.com

Dan Grunwald
Associate, Power and Gas
+1 312 244-7135
daniel.grunwald@morningstar.com

Abhay Kumar Iyer
Software Engineer
+1 312 244-7423
abhaykumar.iyer@morningstar.com

Prashant Srivastava
Data Scientist
+91 22 61217107
prashant.srivastava@morningstar.com

Data Sources Used in This Publication
NOAA
PJM
ERCOT

Executive Summary

Over the past year, Morningstar has embarked on a project to use purely statistical methods to forecast load and day-ahead price in U.S. electricity markets. The latest phase of the project refined the vector autoregressive, or VAR, model methodology by adding functionality in the attribution engine and different switching regimes. We also developed a new modeling approach and tested various scenarios on a stand-alone and combined basis. The results from this phase of our testing are very promising, with the refinements improving the model's output. This white paper explores the improvements and summarizes the results of 135 test cases evaluated since the second white paper was published.

Project Recap

This iteration of Morningstar's first energy-based quantitative model explores the improvements proposed in the first quantitative modeling white paper, PJM Load Model, and the second quantitative modeling white paper, PJM Quant Model the Next Step. The first stage of the model involved using a statistical approach to provide a forward load forecast, leveraging multiple data inputs including prices, market fundamentals, and environmental conditions, and the second version of the model tested the scalability of this approach and its efficacy in different regions. This version of the model continues the work started in previous iterations with the goal of refining the approach. Similar to the last two white papers, the two regions used for testing were ERCOT Coastal and PJM PSEG Zones.

Improvements From Version 2

- Attribution Analysis
- Regime Switching
- Confidence Intervals
- Horse and Cart: A New Forecast Model

VAR and Attribution Analysis

An advantage of VAR models over more typical machine learning techniques, such as neural networks, is the ability to clearly identify the variables that are the biggest drivers of load. In general, a VAR process takes the following form:

$$\boldsymbol{y}_t = \boldsymbol{v} + \sum_{i=1}^{p} \boldsymbol{A}_i \boldsymbol{y}_{t-i} + \boldsymbol{u}_t,$$

where $\boldsymbol{v}$ is a fixed K-dimensional vector of intercept terms, $\boldsymbol{u}_{\mathrm{t}}$ is a K-dimensional white noise or innovation process and $\mathbf{A}_i$ are $K \times K$ coefficient matrices:

$$\boldsymbol{A}_i = \begin{bmatrix} \alpha_{11,i} & \cdots & \alpha_{1K,i} \\ \vdots & \ddots & \vdots \\ \alpha_{K1,i} & \cdots & \alpha_{KK,i} \end{bmatrix}.$$

In selecting the number (K) of features $(\boldsymbol{y}_t)$ for a VAR forecast, the number of lags ($\boldsymbol{p}$), which represents how far back in the time series the autocorrelations are included in the model, is also required. For example, what happened 24 hours ago ($\boldsymbol{t}$ -24) is expected to contribute to the forecast at time $\boldsymbol{t}$, both in terms of what the load was 24 hours ago as well as the other vectors (or features) included in the model. The extreme contributions at the 22rd and 23th lag shown in exhibits 1 and 2, which display the attribution for a 26-lag forecast model at time $\boldsymbol{t}$, confirm this intuition.

VAR Hybrid Model

Since the coefficients in the VAR model are based on longer-term trends, the forecast can be insensitive to sudden changes in the weather. To overcome that momentum in the VAR forecast, we estimate the first 24 hours of the forecast using the Horse (HOmochronic Recessed Smoothing Ensemble) methodology outlined at the end of the paper. This estimate is appended to the load history and brought into the VAR model along with the other weather variables. The VAR model then makes a 24-hour forecast and the process is repeated seven times for the seven-day forecast.

PJM PSEG Attribution Analysis

The image shows three temperature attributes colored according to their relative impact based on the color-coded key to the right. The x-axis shows the impact of historic data by lag period from left to right using the same color key. Exhibit 1 shows the attribution output from the VAR approach, which has been the principal methodology pursued by the team. This heat map represents the load attribution for hour-ending 1 on May 1, 2018. The chart reflects the variables used in the model to forecast load. As additional variables are identified and included, the attribution engine can be run to show users each respective variable impact on the forecast output. As an example, and in the interest of readability, temperature and derived temperature inputs were used in PSEG.

Exhibit 1 PJM PSEG Attribution Analysis May 1, 2018, Hour Ending 1

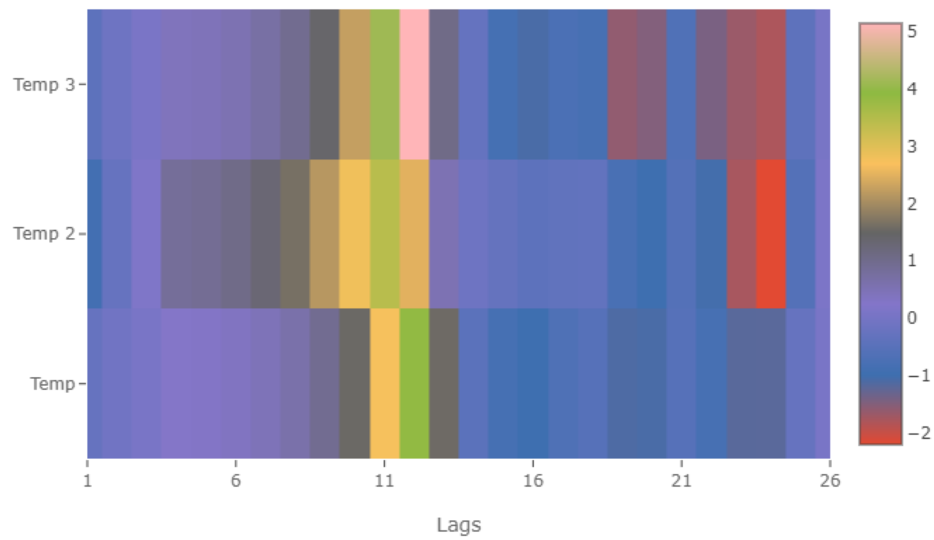

Source: Morningstar.

ERCOT Coastal Attribution Analysis

Similar to the attribution in PJM PSEG, the same approach was applied to ERCOT Coastal region. Exhibit 2 shows the attribution analysis done for hour-ending 1 for the June 1, 2018, forecast. Like the cases in PJM, additional variables can be added to expand the attribution heat map. Analyzing the attribution output for this displayed hour, we can see the positive and negative impact temperature has on load in the 10th and 23rd lag. The impact of cooling degree-days on load is also neutral on the June 1 forecast, which likely means that temperatures did not hit a threshold requiring additional cooling demand. Although the impact of several of these variables may be obvious, the intuition and conventional wisdom can be confirmed in a rigorous statistical manner through attribution.

Exhibit 2 ERCOT Coastal Load Attribution Analysis - June 1, 2018 Hour Ending 1

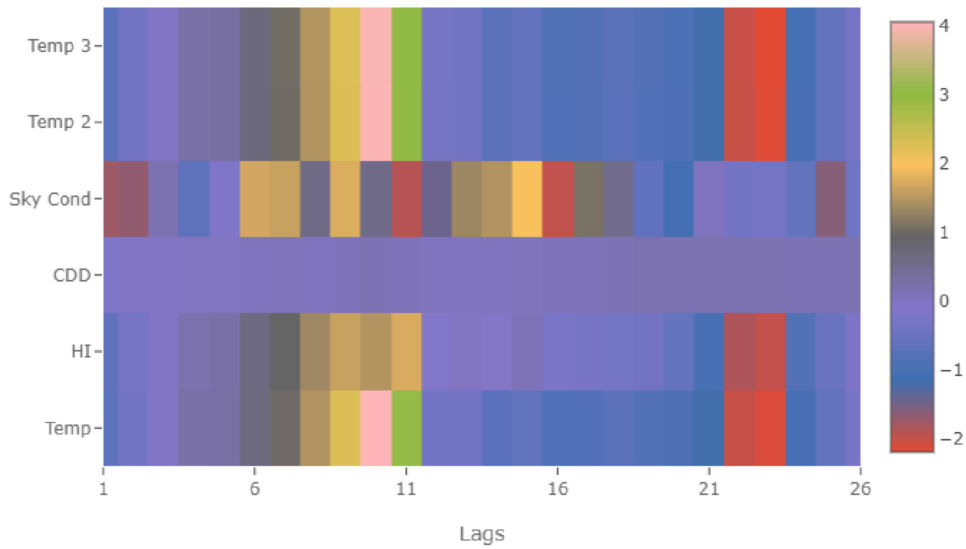

Source: Morningstar.

The ability to calculate the load contribution each variable has on a specific forecasting hour is an advantage of using a VAR approach over other statistical methodologies where the drivers are often impossible to decipher. The ability to determine which variables drive load at different times of the year not only helps statistically confirm existing biases, but also narrows down future models and runs to the variables that are most impactful. Future iterations of the attribution analysis will look to identify and visually show drivers of change between model runs to define the source of changes to the load forecast.

Switching Regimes

Back-testing for the last two white papers used undifferentiated data to generate the forecast, but in this version, two switching regimes were tested. The first switching regime set out to create different models by season, and the second switching regime looked to differentiate between weekends/holidays and weekdays. Electricity demand differs significantly between high-demand periods often seen in winter and summer, and lower-demand periods in the shoulder months. It is also well-known that load profiles change between weekdays, when demand is typically higher, and weekends/holidays, when load is lower. The regime switching tests not only verified this difference using a purely statistical approach, but also improved the forecast models output when back-testing.

Seasonal Switching

To construct the seasonal switching model, a rolling three months was used instead of a fixed seasonal block. The switching helped improve the model's accuracy significantly. For example, the model's output without any regime switching component in ERCOT Coastal saw a mean absolute percent error, or MAPE, of 25% for the nine-month testing period across the four forecasting periods (24-hour, 48-hour, 120-hour, and 168-hour). The MAPE is the moving average of the absolute percent error between the actual load and the forecast load. Thus, the 1-day MAPE represents the average absolute percent error over 24 hours. This compares to an 11% MAPE when a seasonal switch was used to generate the forecast. Exhibit 3 presents the MAPE difference between the base case, which is a model with no regime switching, and the scenario where a seasonal switching regime was introduced.

Exhibit 3 MAPE Difference Between Annual and Seasonal Data—ERCOT Coastal

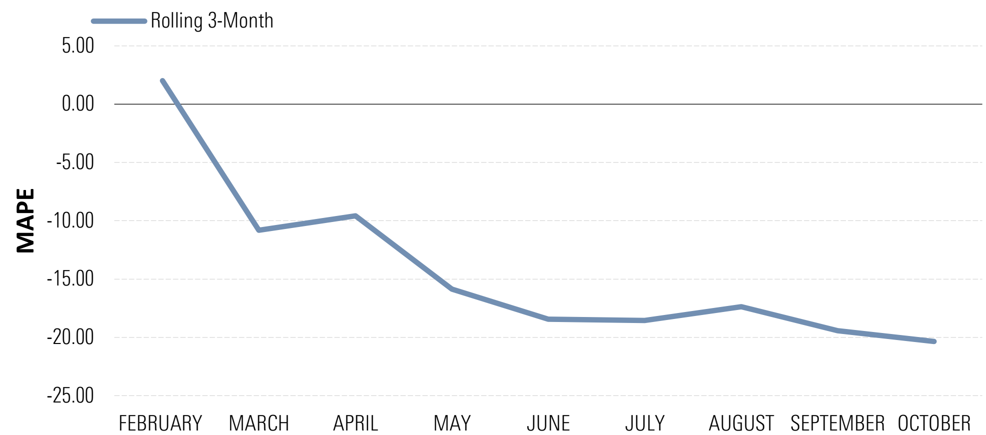

Source: Morningstar.

Weekday/Weekend Regime Switching

On top of the decision to differentiate between seasons by employing a rolling three-month window, a second switching regime was designed to differentiate between weekends and weekdays. The demand profile of the grid differs significantly between weekdays and weekends, and the profile of individual users in any specific geography drives those differences. This difference makes the need for two distinct intraweek regimes important. In cases where no differentiation was made, weekend load forecasts consistently overshot actual load. Exhibit 4 looks at monthly average MAPEs between the rolling three-month model and the rolling three-month model with weekend/weekday differentiation.

Exhibit 4 Error Difference Between Annual and Seasonal Data With Weekend Differentiation—ERCOT Coastal

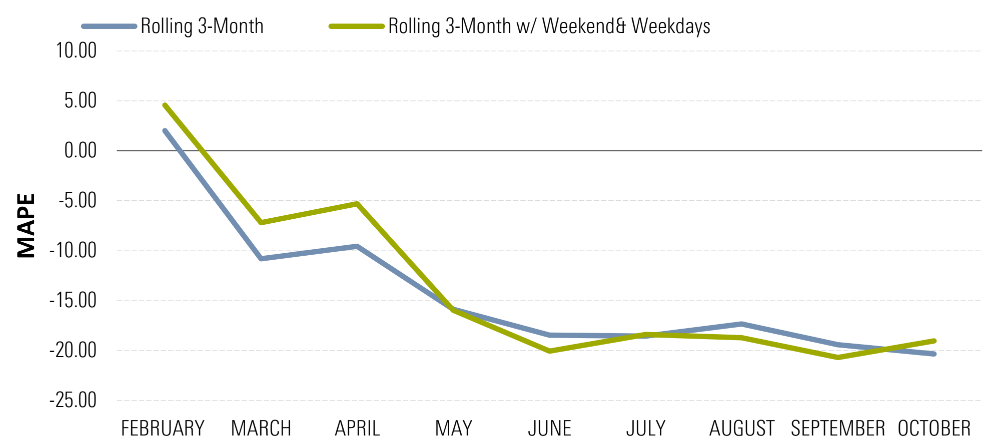

Source: Morningstar.

In the summer months and the fall shoulder, the intraweek switching regime improves the MAPEs by a few percentage points versus the seasonal-switching-only forecast. However, the performance when back-testing in the winter months shows the seasonal-only model performing better than the model where both switching regimes are used. The use of more history to train the model and/or the introduction of defined seasonal blocks may improve the model's accuracy in the winter. Deeper scenarios involving additional history and variables will be tested in the next phase to determine the optimal mix to train the switching component of the VAR model.

Horse and Cart Model

In addition to refinements in the VAR forecast model, we introduced a new model that is equally scalable and shows very promising results. The horse and cart model (HOmochronic Recessed Smoothing Ensemble with Corrective AutoRegression Transfer function) is a regression model that draws inspiration from linear control systems to predict load. The 95% confidence interval shown in exhibits 5 and 6 is calculated based on the distribution of historical loads within ±3 degrees Fahrenheit to provide context for the forecast.

Exhibits 5 and 6 show a seven-day forecast using the horse and cart methodology in both ERCOT and PJM, respectively, with confidence intervals calculated for the week of July 15, 2018. Starting with the ERCOT results, the Morningstar forecast MAPE against actual load for those days were 1.6%, 1.7%, 2.6%, and 2.2% for the 24-hour, 48-hour, 120-hour, and 168-hour forecast, respectively. Additionally, both the forecast and actual loads fall within the confidence intervals calculated for the respective period.

Exhibit 5 ERCOT Coastal Seven-Day Horse and Cart Forecast—July 15, 2018

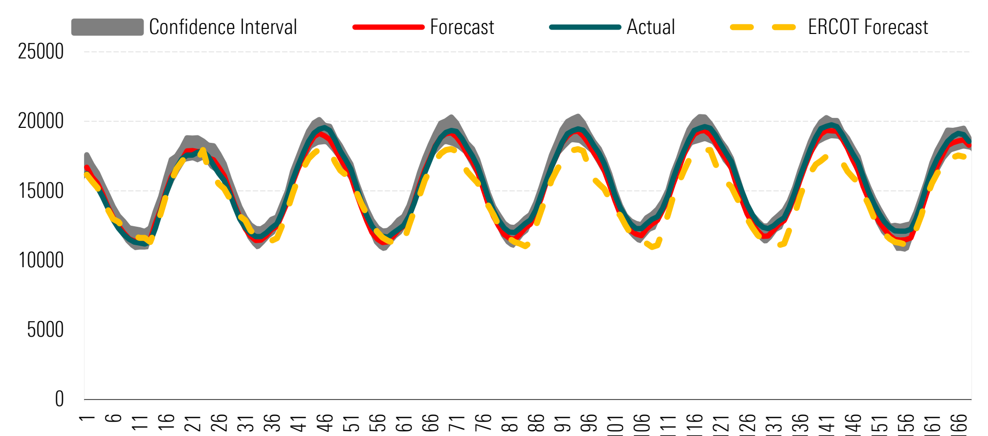

Source: Morningstar.

Moving to the PJM PSEG results for the same week (July 15), the Morningstar forecast had a 4.6%, 3.9%, 3.0%, and 4.8% MAPE in the 24-hour, 48-hour, 120-hour, and 168-hour forecasts, respectively. PJM's forecast had a 4.0%, 4.7%, 4.6% and 6.9% MAPE for the same period. Exhibit 6 charts both Morningstar's and PJM's seven-day forecast against actual results. Confidence intervals against

Morningstar's forecast are also included in the chart. Over this period, Morningstar's forecast not only outperformed PJM's own internal load forecast for PSEG but also captured the load shape more accurately.

Exhibit 6 PJM PSEG Seven-Day Horse and Cart Forecast—July 15, 2018

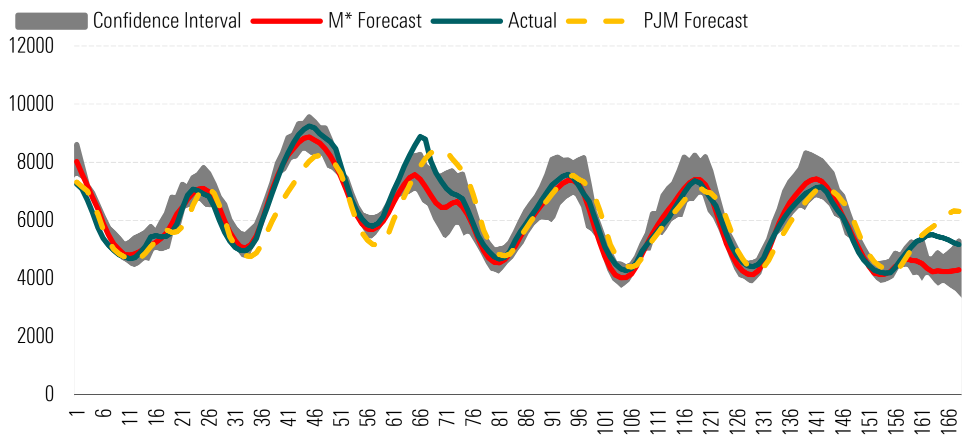

Source: Morningstar, PJM

Horse and Cart Monthly MAPE Averages

Several scenarios were run using the horse and cart methodology to test the model's output across different parts of the year. Exhibit 7 shows the average MAPEs by months for the four forecast periods. As can be seen in the chart, short-term forecasts (24-hour and 48-hour) appear to perform slightly better than the longer forecasts (120-hour and 168-hour), but all four periods show very positive results. The outlier month for the tests in the 120-hour and 168-hour forecast was October, which still averaged only a 5.6% and 5.0% MAPE. When all the MAPEs are averaged across the back-tested nine-month period for all forecasts, the MAPE average comes in at 3.0%.

Exhibit 7 ERCOT Coastal Monthly Average MAPE

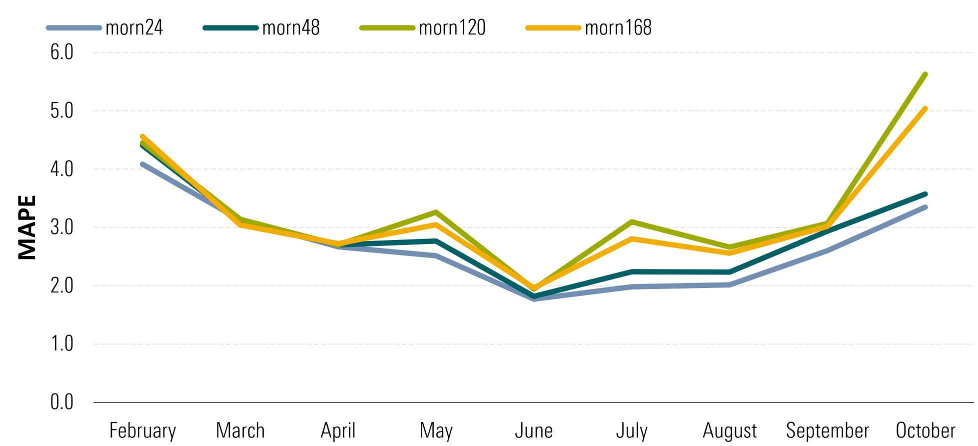

Source: Morningstar.

PJM PSEG's forecast tells a very similar story, with the monthly average MAPEs coming back in a relatively tight 3.5% range between 3.6% and 7.1%, as can be seen in Exhibit 8. Like the ERCOT Coastal results, the shorter-term forecasts performed better than the longer-term forecast in aggregate, with October results being an outlier. The better performance may be attributed to more accurate 24-hour and 48-hour forecasts, but additional testing is required to determine if that is the case. The average monthly MAPE for the test runs across the four forecast periods was 5.2%, which is also a very positive result for this methodology.

Exhibit 8 PSEG Monthly Average MAPE

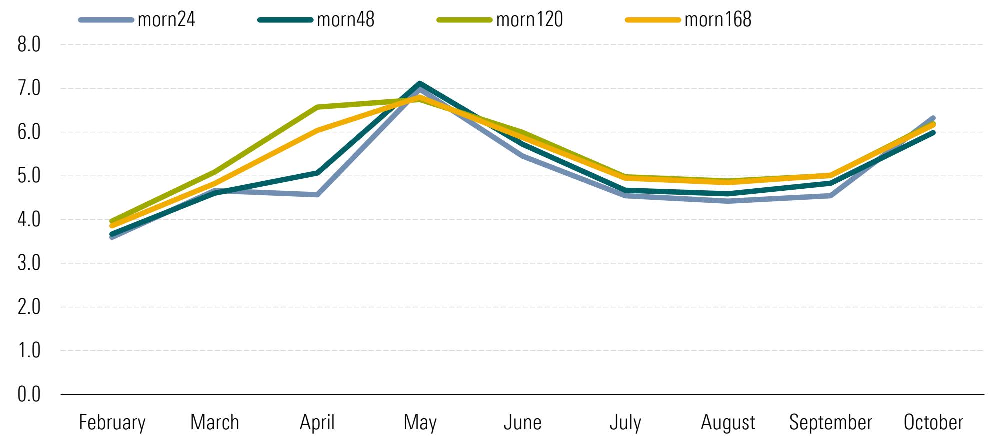

Source: Morningstar.

The availability of PJM's seven-day forecast allowed us to benchmark our forecast results against PJM's internal model, and the results are competitive. Exhibit 9 displays the MAPE difference between Morningstar's and PJM's internal forecasts, with negative numbers representing cases where Morningstar's MAPE was less than PJM's. The Morningstar model performs better in the summer months, with the remaining months within a few percentage points. Over the nine-month testing period, the Morningstar forecast came within 0.75% of the PJM internal forecast across the four forecasting periods.

Exhibit 9 PSEG Monthly Average MAPE Difference With PJM Seven-Day Forecast

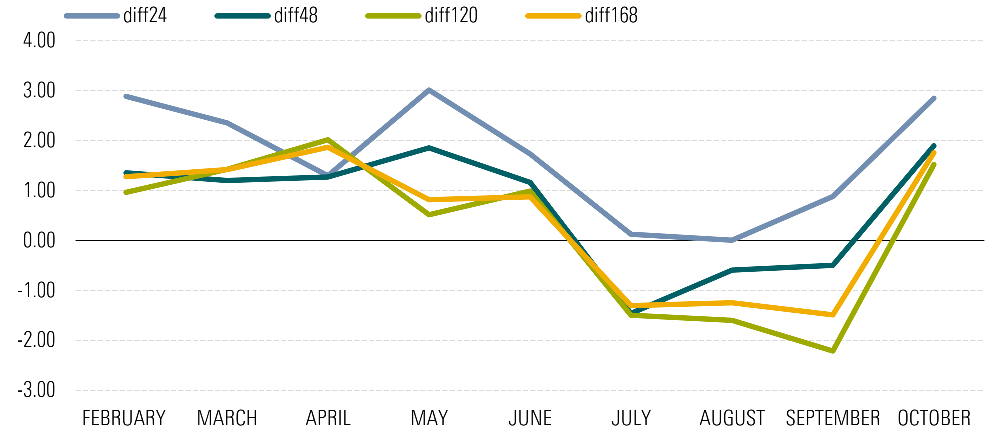

Source: Morningstar.

The results from the horse and cart model are very promising in both PJM and ERCOT. The next phase will look to develop this model's methodology with further tests and deployment into a pseudo-live environment, where it can be tested as a trading and analytics tool. Some of the additional tests that will be conducted on this model include incorporating more history and expanding the number of weather variables used to generate the forecast.

Conclusion

The results from this iteration show considerable improvement from phase two tests conducted earlier in the fall of 2018. The test cases from both PJM and ERCOT show a marked improvement in the primary VAR model and from the introduction of the horse and cart model. In this phase, over 135 cases were tested, and several model improvements were made. The improvements include a defined switching regime, the inclusion of confidence intervals, and the creation of an attribution engine. The next phase of the project will focus on providing the output of both models, the VAR and horse and cart, on a daily basis. We will also refine the model usability with a focus on the end user interface.

The following features will be included in the next iteration of Morningstar's load and price forecast model:

- Delta attribution for the VAR model
- Deploying the model into a pseudo-live environment with daily runs
- Improving the attribution visualizations
- Further refining the VAR and horse and cart models
- Creating a day-ahead price forecast using the horse and cart model

Feedback

We actively seek feedback on model development from clients and prospects and would be interested in discussing the value of this work with commercial trading and analytics as well as to compare our results and approach with internal and external alternatives used by the industry. Please contact Matt Hong or Michael O'Leary for additional comments or questions.

Productization

Our experience in managing data suggests the model results could be made available to clients either as several data feeds or an analytic platform where customers can interact with the model directly. We welcome feedback and suggestions regarding the optimal mode to engage our customers with our model results.

About Morningstar® Commodities Research™

Morningstar Commodities Research provides independent, fundamental research differentiated by a consistent focus on the competitive dynamics in worldwide commodities markets. This joint effort between Morningstar's Research and Commodities & Energy groups leverages the expertise of Morningstar's 23 energy, utilities, basic materials, and commodities analysts as well as Morningstar's extensive data platform. Morningstar Commodities Research initially will focus on North American power and natural gas markets with plans to expand coverage of other markets worldwide.

Morningstar, Inc. is a leading provider of independent investment research in North America, Europe, Australia, and Asia. The company offers an extensive line of products and services for individuals, financial advisors, and institutions. Morningstar's Commodities & Energy group provides superior quality market data and analytical products for energy data management systems, financial and agricultural data management, historical analysis, trading, risk management, and forecasting.

For More Information

+1 800 546-9646 North America
+44 20 3194 1455 Europe
commoditydata-sales@morningstar.com

MORNINGSTAR®

22 West Washington Street
Chicago, IL 60602 USA