AnalytEcon

ECAST

A Brief Guide

July 2014

1. INTRODUCTION

ECast is a statistical forecasting model for the Australian National Electricity Market (NEM). ECast forecasts half hourly price and demand and provides a detailed assessment of price risks over a 48-hour horizon for the mainland NEM regions. This risk assessment includes:

- Forecast probabilities of extreme high and low price events; and
- Confidence bounds about demand and price forecasts that the current volatility in the spot market.

ECast makes use of recent advances in information processing and robust methods of model selection and estimation to create a unique forecasting platform that evolves with changes in the market. The forecasting model has been explicitly structured operate with and without a carbon tax and is operating in post tax mode.

This document is an introductory guide to the forecast methodology, the structure of model and an explanation of the model output. The guide is intended for use with an example of the complete ECast model output provided in the format of an Excel workbook.

2. GENERAL APPROACH

Under certain demand and supply conditions, often brought on by weather conditions, electricity prices in the NEM take on extremely high and/or very low (negative) values. While these events are infrequent, the extremely long tails of the electricity price distribution represent a material source of risk for NEM market participants and other traders. The overriding objective of ECast is to provide a forward assessment of these risks.

ECast treats exceptionally high and low prices as extreme events that form their own price distributions. The balance of prices forms a central regime, again with its own price distribution. Together these three regimes form a mixed distribution that is flexible enough to capture the full shape of the price distribution and sufficiently structured to forecast prices along with their associated confidence bounds.

Isolating and then modelling extreme prices allows for a robust assessment of the likely outcomes of an extreme price event and a balanced assessment of their contribution to downside and upside market risk. The outcome is a better and more complete picture of the range of likely market outcomes over the next 48 hours.

ECast works by defining price thresholds to establish the boundary between regular and extreme prices. These thresholds define the three price regimes: ‘high’, ’central’, and ‘low’.

Associated with each price regime is a forecast of:

- The probability that a price will fall within that regime, i.e. the probability that the price will be in the high, central or low price range;
- The price, given that the prices falls within the high, central or low price regime (a conditional price).

The overall price forecast is constructed as a weighted average. The conditional price forecasts in each regime are weighted by the forecast probability that each regime will occur and then summed. While this might seem like a long way to construct a price forecast, treating the price distribution as composite gives a more accurate a robust representation of the forward price distribution.

In addition, the forecast probabilities that a high or low extreme will occur can also serve as a warning or alert that there are elevated upside or downside price risks. Alerts can be set when the probability of an event exceeds a user specified critical value. In ECast alerts are benchmarked on the basis of historical predictions of high and low price events. In contrast to the model estimation where the price thresholds are fixed, the alerts can be set for and benchmarked against different price threshold of interest.

In addition the probability of high or low price event can to profile downside and upside price risk by calculating a conditional price given a specified event probability is exceeded. For example, the daily average and given the probability of a high price event is greater than 25, 50 or 75 per cent. These conditional means can then be compared to the corresponding unconditional prices for the market as a whole. These profiles discussed in more detail in the in the section covering model outputs.

While ECast demand forecasts make use of the same information processing, model selection and estimation methods, the overall approach is more conventional. This is simply a reflection of the much less extreme shape of the electricity demand distribution.

2.1. THE ECAST METHODOLOGY

ECast price forecasts are done in six stages:

- Prices are converted to current dollars using the CPI and then transformed to reduce the skewness of the price distribution by taking the cube root. In contrast to natural logarithm the range of cube root includes negative prices.
- Threshold prices are set to establish the high and low price regimes, these regimes differ prior to and after the introduction of the carbon tax. The repeal of the tax would be treated as a reversion to the pre tax distribution.

- The probability that prices will, over the next 48 hours, be in the high or low price regime is forecast explicitly (the balance is the probability of being in the central price regime) is derived.
- The forecast price is then estimated in each regime, the forecast is the:
  - Mean or expected price in the central regime; and
  - Median price in the high and low price regimes. The median is conservative forecast as it will be below the expected price in the high price regime and above in the low price regime. However, the forecasts will be more robust in the sense that the majority of forecast errors will be smaller, with a relatively small number of very large errors.
- The overall price forecast is the sum of the regime prices, weighted by the regime probabilities, which is then transformed back to level terms.
- The daily price forecast is the simple average of the half hourly forecasts.

Electricity demand forecast is estimated unconditionally in level terms.

The historical forecast errors in prices and demand is also modelled to construct confidence bounds about the forecasts that vary with market conditions such time of day and recent levels of price volatility. It also facilitates the generation of Monte Carlo style price outcomes from the ECast system. Monte Carlo outcomes can be combined with longer-term projection of average electricity prices and demand for use in project valuation or evaluating marketing strategies that depend on tails of the electricity price distribution.

Price thresholds

Selecting appropriate price threshold requires a balance between capturing extreme price risk and retaining a sufficient number of observed events to produce reliable statistical estimates of the forecast probabilities and prices. The shape of the price distribution has a large influence on what constitutes and extreme price event. The upper tail of the distribution of spot electricity prices is much longer than the lower tail. The 99^{th} percentile may represent considerably more upside price risk than downside represented by the 1^{st} percentile (the percentiles that bound the highest and lowest one per cent of prices). This would favour the selection of more extreme bound for the low price threshold.

The introduction of the carbon tax in July 2012 has also had a significant impact on the distribution of spot market prices in NEM. The 99^{th}, median (50^{th}) and first percentiles of the historical price distributions prior to and after the introduction of the tax are shown for each NEM region in Table 2-1. The extent to which the distributions are skewed to the right is clear, with the median being much closer to the 0.5^{th} than the 99^{th} percentile.

After the introduction of the carbon tax there has been a reasonably consistent increase in the lower end and middle of the price distribution (as indicated by first percentile and median of the price distribution in each of the regions). There is considerable variation in the pass through of the tax between NEM regions. Therefore, the low price regime thresholds were set at the first percentile in the pre and post tax periods in each region. Given the carbon tax is repealed, the lower tail would initially be assumed to revert to its pre-tax thresholds.

The pass through of the tax in the upper tails of the distributions is not obvious. This could be due, in part, because the pass through of the tax is small component of the total price and therefore has less of an influence on high-end spot market price outcomes. In any case, with the limited mount of post tax price data and the long upper tail of the price distribution it does not appear that pre and post tax thresholds can reliably be identified. Therefore, the threshold for the high price regime is set for the 99th percentile over the entire historical period.

The price thresholds evolve with time and given the carbon tax is repealed the pass through back into spot market prices will be reviewed on an ongoing basis and the thresholds potentially revised. Current thresholds are reported as a part of the forecast output.

Table 2-1. The 99th, 50th and 0.5th percentiles of the pre and carbon tax price distributions, as of June 2014.

State	Pre Carbon Tax			Carbon Tax		
	99%	Median	1%	99%	Median	1%
NSW	$182	$29	$14	$93	$52	$41
QLD	$181	$26	$13	$219	$53	$32
SA	$187	$34	$10	$203	$53	$27
VIC	$143	$30	$10	$110	$49	$35

Statistical techniques

The explanatory variables to predict prices and demand include:

- Past levels of prices and demand in each of the NEM regions;
- Forecast weather conditions in each capital city;
- Fixed factors, such as time of day, day of the week and season; and
- Fixed factors for the introduction and repeal to the carbon tax.

The unique feature of ECast is the way this basic information is processed and forward prices are modelled. ECast extracts trends in the level and volatility of prices and demand at different time scales using discrete time series decomposition. The working hypothesis is that changes in, for example, past prices over the last hour, hour or several hours have different information content with respect to forward state of the market. Short term trends may be more or less persistent that longer term trends. Further, the relevance of this content will change over different forecasts horizons. That is, the trends most relevant to 12 versus a 48 hour horizon are likely to differ.

As there are a very large number of explanatory variables that might be included a robust method of model selection is used to choose a model with the best reliable fit from a large number of potential models. Reliability is gained by not allowing the model to be over fitted to the historical data. Adding additional variables will give a better model fit but this will eventually come at the expense of the precision with which model parameters are estimated and the reliability of the forecast.

Robust methods of estimation are also used when modelling the tails of the price distributions. The result is an overall model that generates reliable forecasts given the prior information available. The forecasting system evolves with changes in the market as the entire process of model selection and estimation is redone each time the model is updated.

ECast only makes use of information made at the time of the forecast. Separate models as are optimised for each forecast horizon. This also serves to prevent forecasts errors from compounding which can happen when the forecast in the next period depends on the forecast in the last period. There are currently two models, one with a 24 hour horizon and a 48 hour horizon[1].

The distribution of the historical forecast errors in prices and demands is modelled as a function of fixed factors and measures of current price volatility. These measures of volatility are at different time scales, that is, separating short form longer term sources of variability prices as for example, half hourly, hourly and longer. A general additive modelling approach is used to estimate the historical error distribution parameters and generate the confidence bounds.

3. ECAST INTERFACE AND OUTPUT

ECast data acquisition and checking, model updating and forecast delivery is automated Forecasts are scheduled through a simple graphical interface. Once the forecast is completed,

[1] This could be refined optimising models are greater number of time scales, as for example, 6, 12, 24, 18 and 48 hours.

ECast automatically generates a report in the form of an Excel spreadsheet that is uploaded to an FTP site. The Excel format allows convenient access to the forecasts and historical model output. However, there are no internal formulas in the spreadsheet.

There are two reports with output tabs for each NEM region.

- A long report generated, typically when the model is updated, that includes information on the historical performance of the model and a graphical analysis of the regime probabilities and prices for each NEM region:
- A short report that contains only the forecasts and recent model performance measures.

In each report there is a Notes tab in which provides the:

- Time of the forecast was generated;
- Weather forecasts for the capital cities;
- Time when the NEM data was updated;
- Last time the model was updated; and
- Current high and low price thresholds.

3.1. HISTORICAL FORECASTING PERFORMANCE (TAB)

The historical forecast performance tab contains:

- Summary measures of forecast accuracy for the probabilities of high and low price events, prices and demand;
- A detailed breakdown of forecast accuracy by time period and season for prices and demand.

The historical output tab also includes a price risk assessment based on the probabilities high and low price events. This includes summary measures of the accuracy the model can predict an event with a specified threshold. The events specified for the:

- Maximum daily price;
- Average daily price; and
- Minimum daily price.

The selected price and probability thresholds can easily be modified to meet a users needs. This risk assessment is discussed in greater detail in the final section of the ECast Guide.

There is also a comparison of average prices that are conditional on an event probability exceeding a specified threshold. That is, given a threshold probability is exceeded, what is the expected spot market price.   These can be compared to the corresponding unconditional (overall) average price. This is intended to demonstrate how the event probabilities can be used to construct a profile of upside or downside risk. Conditional prices across a range of event probabilities are used to create a risk profile. The profile constructed in the historical performance tab is based on three threshold probabilities, 25, 50 and 75 per cent.

Plots

High and low price event probabilities are plotted along with the corresponding daily maximum, average and minimum prices for the past year. Prices are scales to facilitate comparison. The scaling is done by first taking the cube root of price and then scaling it as a percentage of the observed price range, the formula for the latter is:

$$\frac{P_T - P_{\text{Min}}}{P_{\text{Max}} - P_{\text{Min}}}$$

This puts prices and on the same zero to one interval as the probabilities. The actual price range can vary substantially between NEM regions is given for each graph. The plots are model generated. However, event probabilities and prices are given for the past years are provided in tabular format in a separate tab.

3.1.1. Summary Measures of Forecast Accuracy

The most widely used measure of historical forecast accuracy is R-Square. In context, this is the percentage of total variation in prices or demand about their respective means, successfully predicted by the model. As it is a scale free measure it can be used to compare forecast accuracy of prices and demands with different means and levels of variability, as for example between NEM regions. An R-Square of zero implies that tem model forecasts no better than forecasting price and demand will take on their historical mean. An R-Square of 100% is a perfect forecast. The R-Square values are report for the forecast high and low price event probabilities, prices and demand.

One issue with the R-Square is that a few very large errors, which are likely to occur when there are large price spikes, can make a very large contribution to the variance about the mean and substantially reduce the percentage of variation explained by the model. A robust measure of R-Square is also provided for prices. The measure is constructed by comparing the rank order of the forecast to the rank order of actual price outcomes. A rank R-Square of 100% indicates that order of the forecasts from low to high is exactly the same as the order of the price outcomes. A rank R-Square of zero implies is not relationship between the forecast and actual price outcome. However, it is implicit in the use of such a measure that very large errors, relative to the average price, do occur.

The mean absolute percentage error (MAPE) is another summary measure that, in contrast to R-Square, provides a relative measure of the magnitude of the forecast errors. It is less sensitive to influence of a few very large errors due to price spikes and is commonly used with electricity price forecasts due. However, that same level of relative error can have substantially different levels of absolute error when means of variables being compared are different, as for example, when comparing the MAPE for demand during peak and off peak periods.

In ECast MAPE is used:

- As an overall measure of price and demand forecasts accuracy; and
- The breakdown of forecast error levels by time period and season.

However, in the case of prices the MAPE is adjusted to take into account two important characteristics of the price data:

- Prices can be near or in rare instances equal to zero; and
- The long upper tail of the price distribution is extremely long.

The standard way to calculate a MAPE is to take the mean of the individual percentage errors. However as prices approach zero the parentage error approaches infinity. The alternative is to calculate the ratio of means, which is more robust. This percentage error is the ratio of the forecast error to the mean price.   In ECast the MAPE for prices is the ratio of the error to its corresponding half hourly historical mean. The use of half hourly means facilitates the comparison of forecast errors at different times of the day.

The extremely long tails of the price distribution mace it difficult to give a good picture of forecast accuracy with any single measure. MAPE can easily be recast in terms of percentiles to characterise the distribution of price forecasting errors. IN ECast, the focus is on the upper tail of the forecast errors as described by the median, 75th and 95th per centile of absolute percentage error in the price forecast.

3.2. FORECAST OUTPUT (TAB)

The Excel forecast tab is the same in the long and short reports. The forecasts are made 24 and 48 hours ahead at regional level, and include:

- Half hourly and daily prices and demand and the 10th and 90th confidence bounds;
- Half hourly and daily maximum probabilities that a high or low price event will occur;

- The forecast tab includes daily summary information for the previous seven days (midnight to midnight). This includes daily maximum, average and minimum prices, total demand and he previous week forecasts of the maximum event probabilities;

Half hour prices and demand are also provided for the previous week along with:

- The mean absolute percentage error for each day; and
- The mean absolute percentage error half hourly time period.

A traffic light colouring scheme is used to indicate when the errors are outside of the normal range. This typically happens when there are extreme price events. Ideally, they will be associated with elevated forecast event probabilities. However, unanticipated events occur and these will be associated with large errors and low event probabilities.

4. ECAST EVENT PROBABILITIES

While the high and low regime probabilities pertain to a specific price threshold, they can also be interpreted as an alert or warning of a other impending high or low price events of interest. The extent to which these alerts can predict period of increased price volatility may allow market participants more effectively manage downside and upside price. An example is constructed in the historical ECast model output, showing the risk exposure to:

- Daily maximum prices at or above $1,000;
- Daily average prices at or above $; and
- Daily minimum prices at or below (pre tax) and (post tax).

Trade-offs of setting a critical value

An alert is triggered when the high and low event probabilities exceed a critical value at any point during the next 24 to 48 hours. There are four possible outcomes of predicting or issuing an alert for a discrete event as set out below.

Table 4-1. Forecast outcomes

Status	Event Occurred	No Event
Alert Issued	True Positive	False Positive
Alert Not issued	False Negative	True Negative

The likelihood of these outcomes is dependent not only on how well the model forecasts the critical probability but also on value at which the alert is generated. A lower critical value will increase the number of positive predictions, true and false. A higher critical value will increase the number of negative predictions, true and false.

The ideal critical value balances the risks of correctly or incorrectly responding to an event. The benefits of either true positive or true negative prediction arise from:

- Taking action to capitalise on or mitigate the outcome of a price event;
- Avoiding the costs of taking action to capitalise or mitigate the outcome of a price event that does not occur.

The costs of either a false negative or false positive are just the mirror image of the benefits, arising from incurring the costs of:

- Taking action to capitalise of mitigate on the outcome of a price event when it does not occur;
- Not taking action to capitalise or mitigate an outcome when an event occurs.

When the balance between the benefits and costs of reacting to an alert are even, a neutral critical value, 50 per cent, is appropriate. When the benefits outweigh the costs a lower critical value is warranted. When the costs outweigh the benefits a higher critical value is warranted.

In the example provided in the historical ECast output an alert is deemed made if the maximum daily probability of either a high or low price event exceeds:

- 30 per cent;
- 50 per cent; and
- 70 per cent.

There is another factor that mitigates the costs of a false alert. While prices may not exceed or fall below specific price points, prices may still tend to be considerably above or below when an alert is issued. The alert may still serve to identify elevated levels of risk.

Reporting

For each alert probability and price threshold the following information is reported in the Historical tab for each NEM region:

- The average number of times that the event has occurred per year and per season;

- The percentage of events that were predicted;
- The average number false alerts per year and per season;
- The percentage of events that were predicted without precedent, that is, there was no event in the preceding 24 hours.[2]

Historical event probabilities and predictions can be used to test trading and other response strategies to manage spot market risks.

[2] Given that there was an event in the previous 24 to 48 hours is would seem safe to presume there is an elevated level of price risk and the value of an alert may be small. In contrast, an alert that captures an unprecedented event may have a much greater value.