TRANSPOWER PEAK DEMAND FORECAST UPDATES

1 Purpose

This document outlines improvements made to Transpower's peak demand forecast in response to feedback from the Ministry of Business Innovation and Employment (MBIE) and the New Zealand Institute of Economic Research (NZIER).

2 Background

During public consultation on the draft Electricity Demand and Generation Scenarios (EDGS), stakeholders indicated that MBIE should take a role in validating Transpower's peak demand forecast.

MBIE asked NZIER to undertake a high-level review of Transpower's forecasts and methodology which resulted in a number of suggested improvements. The improvements were ranked in order of importance and ease-of-implementation. MBIE, NZIER and Transpower subsequently agreed that Transpower should implement those improvements which would increase confidence in the forecasts without requiring fundamental changes in the model specifications and data requirements.

NZIER's memo stressed the value of the simplicity and transparency in the forecast methodology. We are mindful of the desirability of not overly complicating our approach.

3 Action points

We have agreed to make the following changes to our peak demand forecast model.

Model changes:

1. Investigate introducing a scheme of using individual model performance estimates to weight the ensemble results.
2. Investigate adding time-series terms to individual model regressions to deal with auto-correlation in model errors

Additional testing:

3. Implement in- and out-of-sample testing of individual forecasts and the ensemble.
4. Undertake testing of break-points in the short-term endogenous model.
5. Undertake testing of dropping data for dry years.
6. Undertake testing of using levels vs growth rates.
7. Undertake testing of log transformations and temperature correction for individual models.

Data consistency:

8. Communicate with MBIE to ensure that specific industrial customer forecasts are consistent with EDGS assumptions.

4 Progress

This section summarises actions taken in response to items in section three.

4.1 Ensemble weighting

Status: Investigated but not Implemented.

We have investigated weighting the ensemble by the quality of fit and are not satisfied that it provides an improvement over evenly-weighting the models.

The method of merging the ensemble provides some quality weighting through the variance in the model fit. At this stage we consider this an appropriate trade-off between accuracy and transparency.

4.2 Time series terms

Status: Implemented

All models other than the short-term endogenous model are now regression models with AR(1) errors, i.e.:

$$y_t = \beta_0 + \beta X_t + \epsilon_t$$

Where :

- y_t is the response variable at time *t*
- X_t are the predictor variables at time *t*
- β_0 is a constant term
- β are the regression coefficients associated with the predictors in X

And ε_t is the t^{th} element of an autoregressive sequence of errors, such that:

$$\epsilon_t = \theta\epsilon_{t-1} + \omega_t$$

Where

- θ is the auto-regression coefficient in the time-series of errors
- ω is a white-noise term

We have found that including two autoregressive terms in the error model provides a more robust approach for a short-term endogenous model.. Overall we feel that this model preforms better than our original short term endogenous model with its breakpoint – see our comments further below. Hence, we have respecified the short-term endogenous model to use an additional lag in the error rather than a break-point, giving:

$$\epsilon_t = \theta_1 \epsilon_{t-1} + \theta_2 \epsilon_{t-2} + \omega_t$$

Each individual model is more fully specified in Appendix A.1

Results:

Autocorrelation in residuals

The autocorrelation in residuals significantly reduced as shown in the graph below.

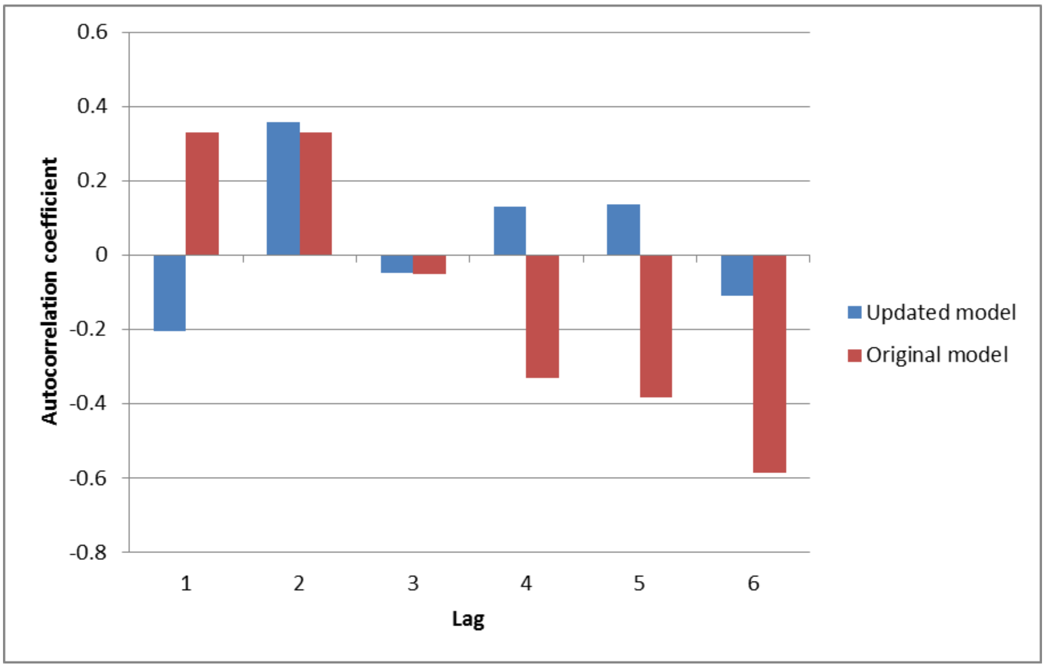

Figure 1: Comparison of autocorrelation in residuals from ensemble model (Auckland winter peak)

Level forecast

The level of the first forecast year is more consistent with the level of the last data year as a result of the correction by the autoregressive term. In general, this has tended to translate the forecast down over the whole forecast horizon. In some regions, there has also been an impact on the growth-rate, but the strongest effect has been on the starting point of the forecast.

The graph below illustrates the first year change for the Auckland winter peak forecast. The effect in Auckland is stronger than in other regions as there has been little or no growth in peak demand in the past several years despite population and GDP growth and a cold winter in 2015.

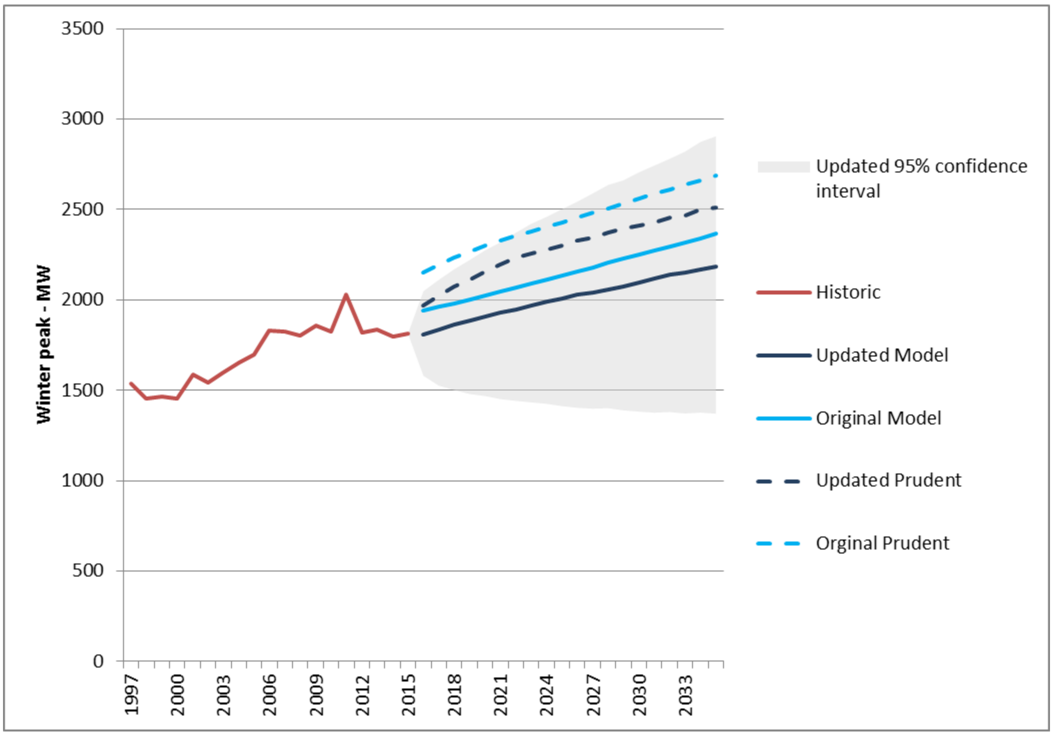

Figure 2: Comparison of forecast levels (Auckland winter peak)

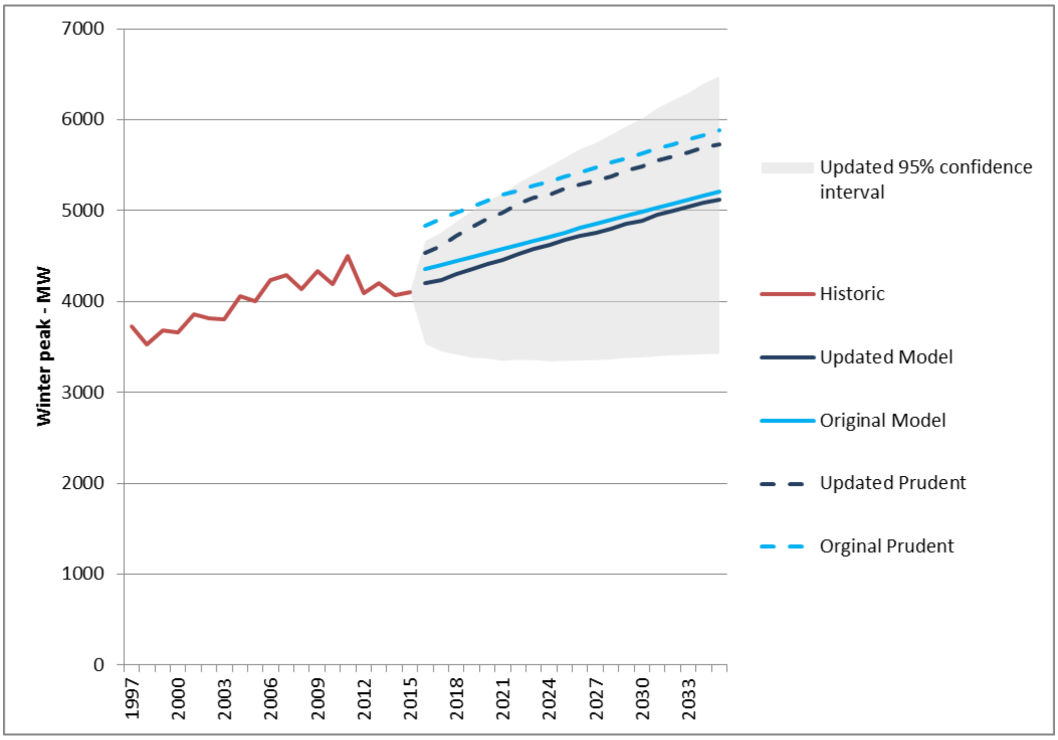

Figure 3: Comparison of forecast levels (North Island winter peak)

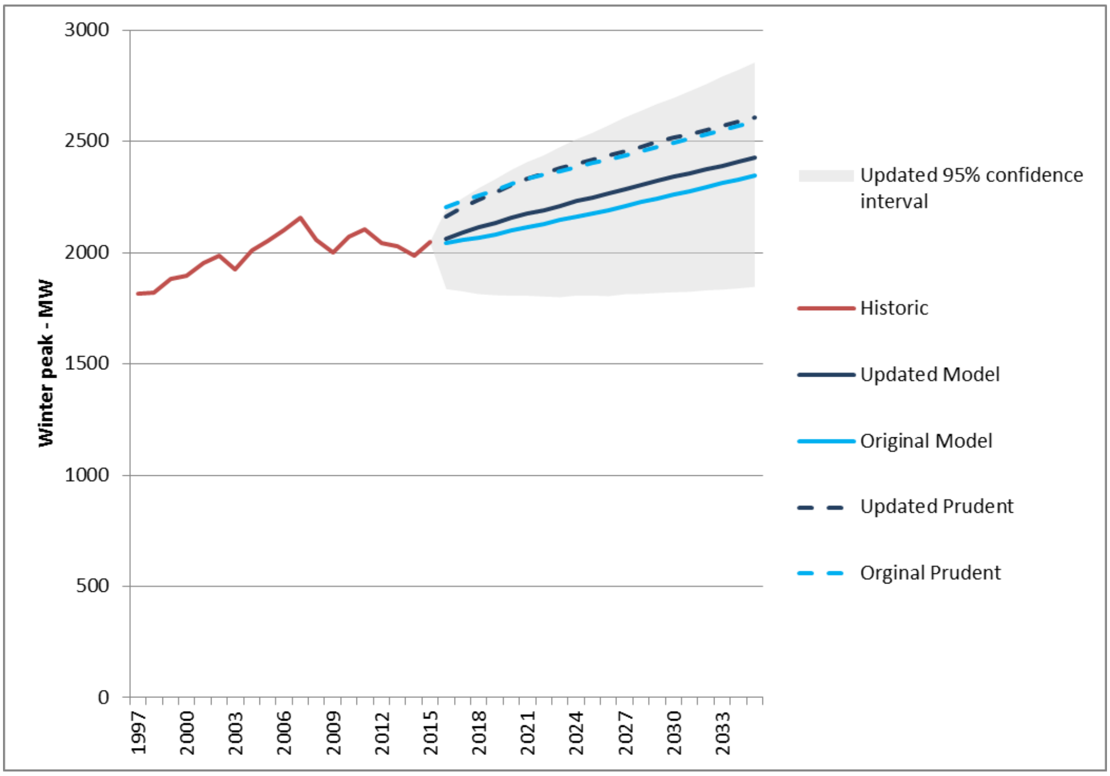

Figure 4: Comparison of forecast levels (South Island winter peak)

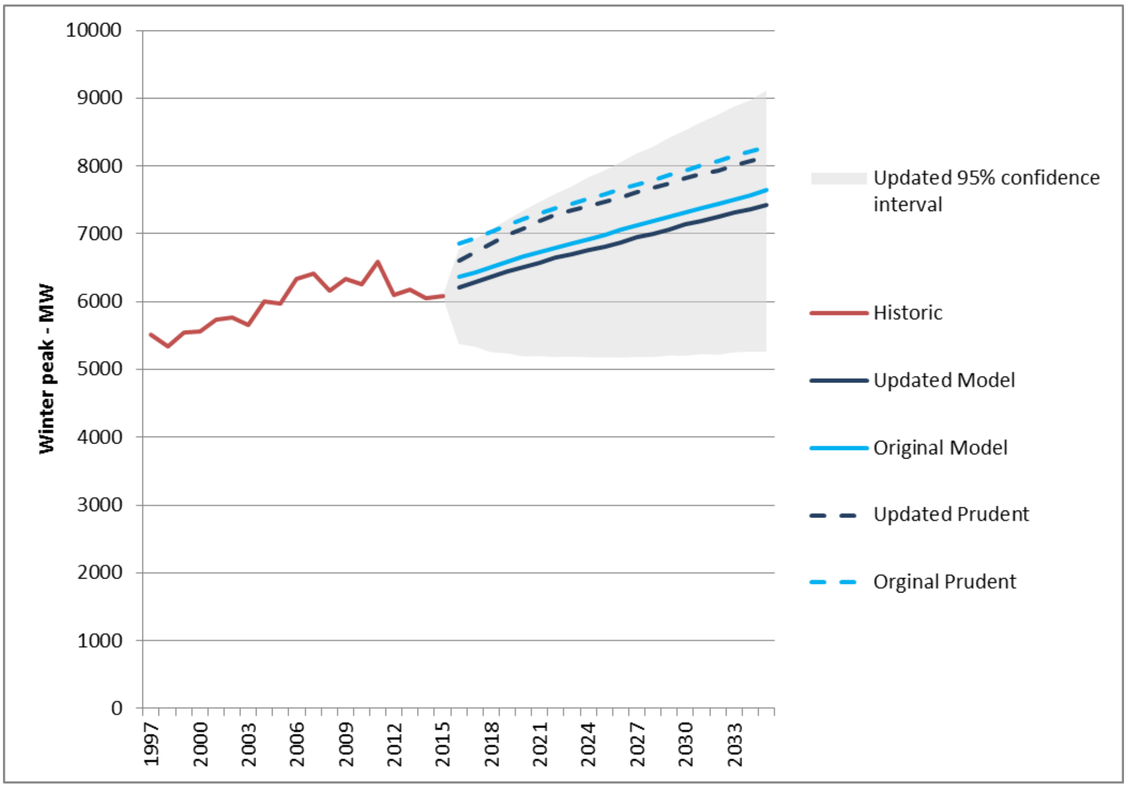

Figure 5: Comparison of forecast levels (National winter peak)

4.3 In and out-sample testing

Status: Implemented.

The user can now select a proportion of the historical data to use as the training sample. The rest will be withheld and used for out-of-sample testing

N-ahead forecasting

We tested the model by withholding 1 to 5 observations and forecasting the last remaining years of data using a model fitted to the earlier subset.

Results

We have considered several sets of individual model specifications and calculated and out-of-sample error for each. Several observations arose from this:

1. Differencing tended to reduce out-of-sample error and the difference between in and out-of-sample error
2. Logging the response variable tended to increase out-of-sample error as there is little evidence of exponential growth in the last several years of observed demand

Table 1: Example set of individual model specifications

Run	Differenced exogenous model	AR terms in error model	Log endogenous models	Log Exogenous model	AR model type
1	no	1	0	0	regression with AR errors
2	no	1	0	1	regression with AR errors
3	no	1	1	0	regression with AR errors
4	no	1	1	1	regression with AR errors
5	yes	1	0	0	regression with AR errors
6	yes	1	0	1	regression with AR errors
7	yes	1	1	0	regression with AR errors
8	yes	2	1	1	regression with AR errors
9	no	2	0	0	regression with AR errors
10	no	2	0	1	regression with AR errors
11	no	2	1	0	regression with AR errors
12	no	2	1	1	regression with AR errors

13	yes	2	0	0	regression with AR errors
14	yes	2	0	1	regression with AR errors
15	yes	2	1	0	regression with AR errors
16	yes	2	1	1	regression with AR errors

There was not one specification that had the lowest out-of-sample mean absolute percentage error across all regions. However, this test did allow us to pick a specification that performed well across most regions. The graph below shows the out-of-sample mean absolute percentage error for the ensemble model in Auckland.

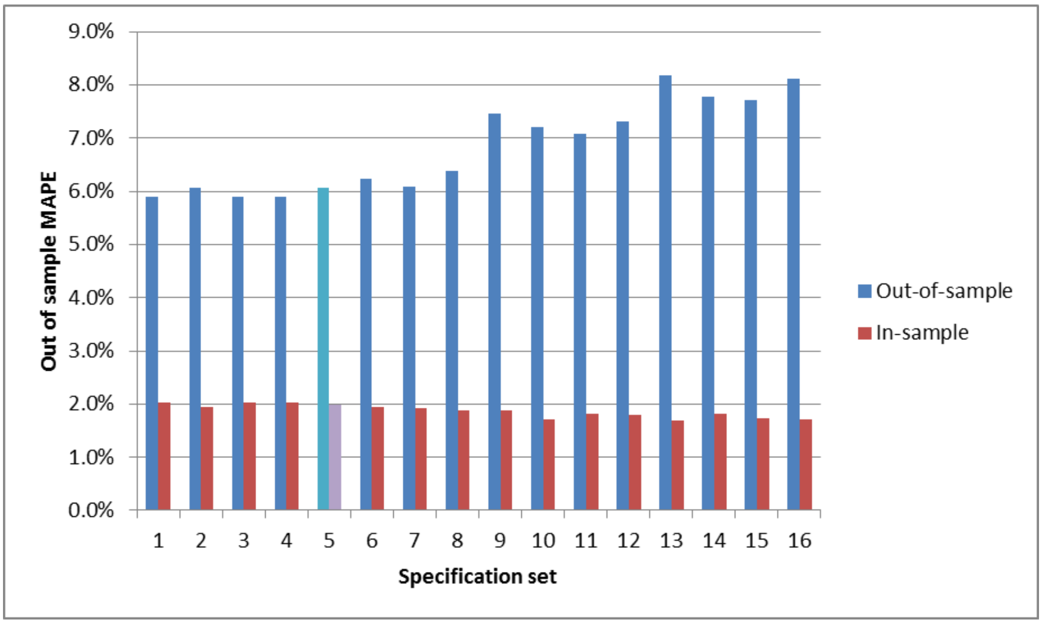

Figure 6: Comparison of forecast levels (South Island winter peak)

The bar highlighted is the set we have chosen to include as it performed well in most the regions and is easily interpreted.

The chosen run corresponds to:

- No log transformations
- Differencing in all models other than the MBIE model
- One AR term in error model other than short-term endogenous model, which has two.
- Regression model with autoregressive errors

This specification also had among the lowest increase between in-sample and out-of-sample error over many regions.

The graphs below show winter peak forecasts to 2020 fitted to data truncated at 2010 to 2015 to show the effect of additional years of data on the long-term trend. The shaded regions indicate years where at least one forecast has had data held back.

There remains a tendency – particularly in Auckland – to over-estimate the most recent years relative to observed peaks, which we will continue to monitor.

The error in 2015 is driven primarily by the temperature component of the individual models. The 2015 winter was cold by historical standards, but we have not seen a corresponding increase in peak winter load. The relationship between temperature and peak appears stronger in the North Island so the out-of-sample forecasts for the South Island tend to perform better.

It is unclear if this is a one-off result but we intend to continue to monitor the performance of the model. We will consider if improvements can be made, such as in our use of a temperature term, as part of our next demand forecasting work cycle.

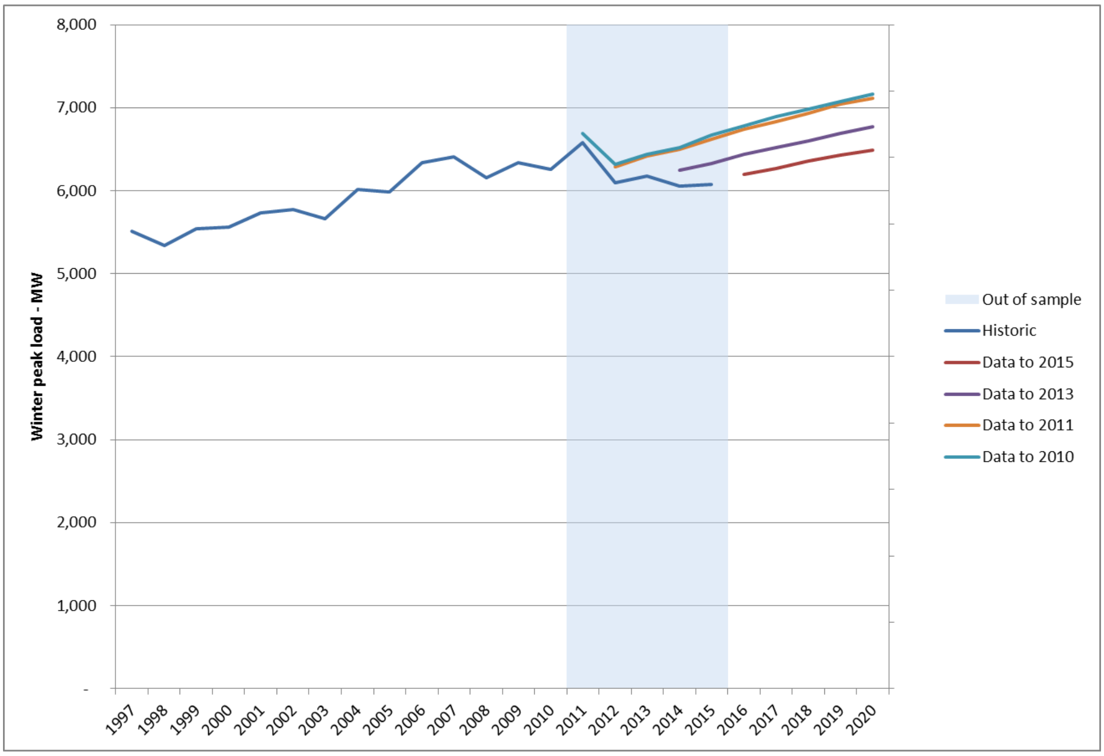

Figure 3: National peak out of sample forecasts

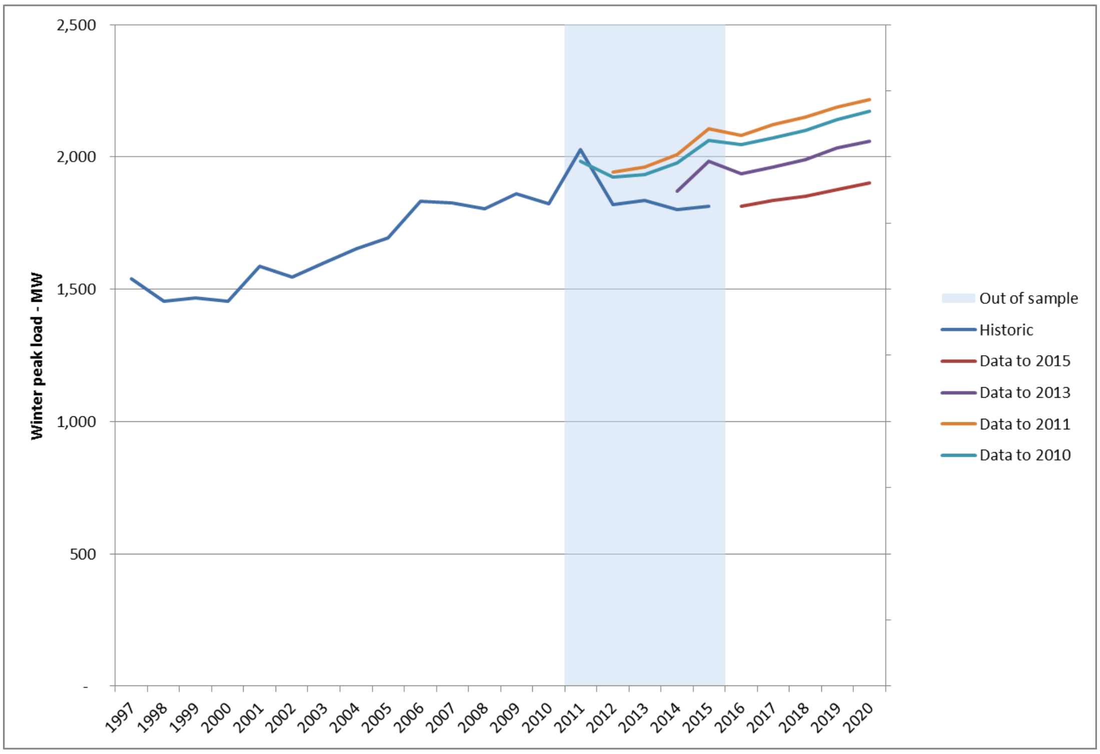

Figure 4: Auckland peak out of sample forecasts

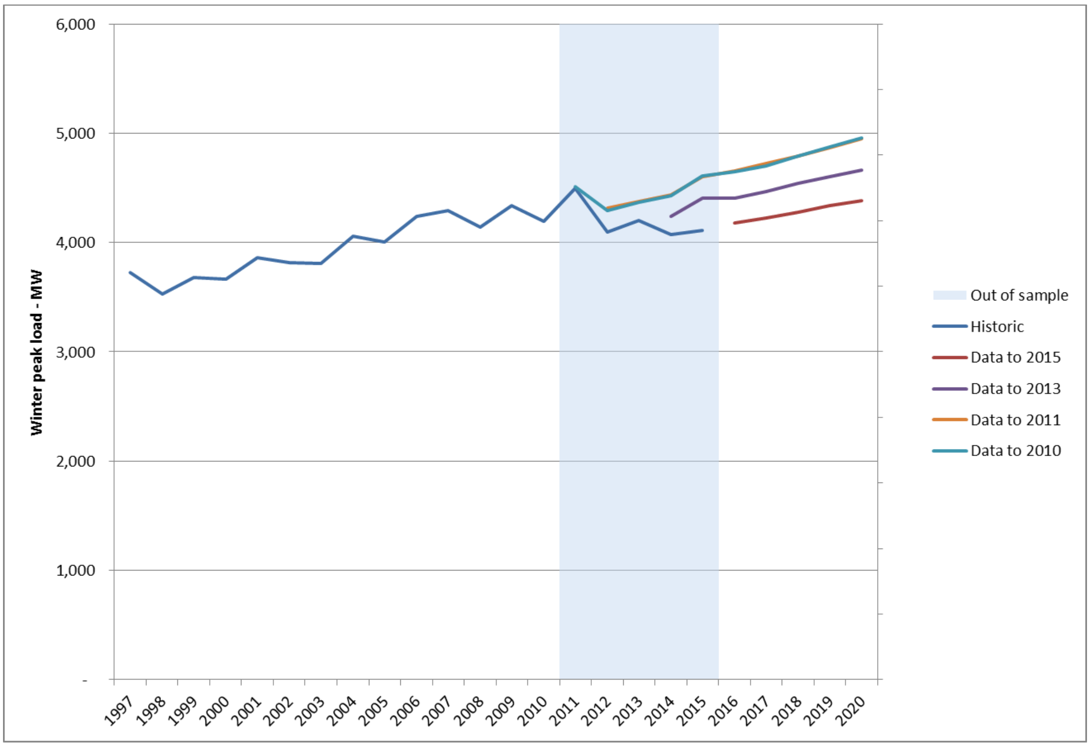

Figure 5:– North Island peak out of sample forecasts

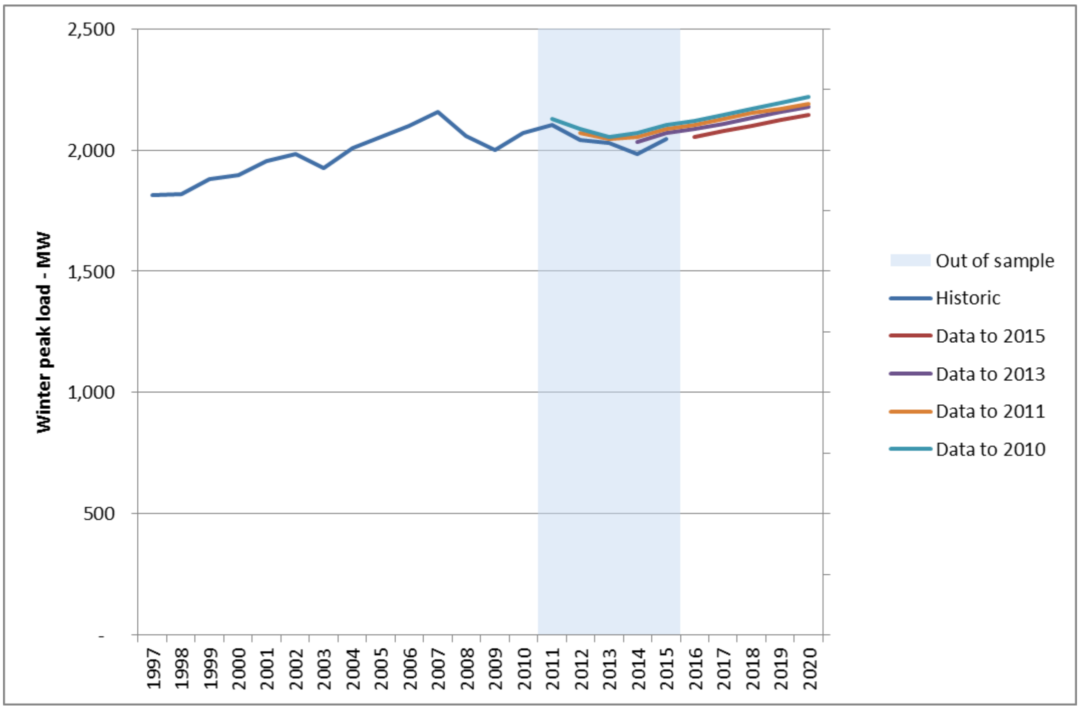

Figure 6: South Island peak out of sample forecasts

4.4 Dropped data for dry years

Status: Implemented

All data years are now included in all models.

The 2001 and 2003 data-points stand out in some regions but not others. For simplicity, we have decided to include this data in the model rather than conduct a test for each region that decides whether or not they are genuine outliers for each region.

4.5 Testing break-points

Status: Not applicable

We have re-specified the endogenous short-term model the same as the long-term endogenous model but with an additional AR(2) error term. The break-point has been removed. The break-point provided some forecast value if the observed peaks flattened but produced unrealistically positive/negative growth in several regions where the last few years have dropped significantly.

A second autoregressive term tends to put more weight on the recent years of low growth as appropriate to the data. Therefore we consider it a more consistent and transparent approach to a short-term model.

4.6 Levels vs growth rates

Status: Implemented.

All models other than the MBIE-based model are now differenced once in order to obtain more stationary inputs.

4.7 Log transformations

Status: Implemented

We have tested log transformations for all models and found that it provides little benefit in fit and results in unrealistically high growth-rates in some regions. Therefore we have left the response and predictor variables untransformed in all models for ease of interpretation.

4.8 Temperature correction

Status: implemented.

Each model can now use the temperature variable for the winter forecast. The AIC for a fitted model with and without the temperature variable is calculated. If the AIC comparison suggests that temperature is significant enough to warrant an additional variable in the model, the model with temperature is used. Otherwise. the temperature variable is discarded.

Results:

Temperature is found to be more significant in North Island and in cities.

4.9 Industrial data

Status: We will discuss this with MBIE to ensure some alignment in assumptions.

5 Future Work

We consider the forecasts in their present state to be fit-for-purpose assuming there is not a wholesale uptake of distributed generation and battery storage.

However, we are mindful that our long-term (or even medium-term) forecasts of grid supplied demand maybe heavily influenced by updated forecasts of solar photovoltaics and battery storage.

With this in mind we intend to do further work this year on developing an approach to adjust our current forecasts to reflect the possible uptake of these technologies.

At this stage, we intend to develop future forecasts using a two-stage process. The first stage would forecast grid demand in a business as usual manner, using our standard approach. In effect, we would assume that there is no rapid uptake of new distributed technologies. The second stage would take the first stage forecast and estimate grid demand with various levels of new photovoltaic distributed generation and storage. We are currently considering the details of how to implement the second stage.

As indicated above, there are certain areas of our modelling that we intend to review further in our next demand forecasting cycle. No forecasting methodology is perfect and we recognise the importance of constantly reviewing the performance of our forecasts in relation to new information.

A.1 Model specifications

A.1.1 Endogenous long-term

Response variable: peak demand

Predictor variable(s): Year, temperature[1]

Differenced: Once

Log transformations: peak demand

Error Model: AR(1)

A.1.2 Endogenous short-term

Response variable: peak demand

Predictor variable(s): Year, temperature[1]

Differenced: Once

Log transformations: peak demand

Error Model: AR(2)

A.1.3 MBIE long-term

Response variable: peak demand

Predictor variable(s): MBIE energy forecast/data, temperature[1]

Differenced: None

Log transformations: None

Error Model: AR(1)

A.1.4 Exogenous long-term

Response variable: peak demand

Predictor variable(s): GDP, temperature[1,2]

Differenced: Once

Log transformations: None

[1] The model selects whether or not to include temperature by comparing AIC of the fit with and without temperature included.
[2] We have re-specified the exogenous model as GDP only. We found little fit benefit in including the additional term and several regions returned illogical coefficients if allowed to fit to both. We consider consistency of having the same variable in all models greater than the benefit of the extra term.

Error Model: AR(1)