

T R A N S P O W E R

Long-term demand forecast

September 2011

Substantive changes since the May 2011 draft are shown in red

© TRANSPOWER NEW ZEALAND LIMITED ALL RIGHTS RESERVED

This document is produced for external release. Its conclusions are based on the information currently available to Transpower and may change as further information becomes available either internally or externally.

Table of Contents

TABLE OF CONTENTS	1
EXECUTIVE SUMMARY	2
1 BACKGROUND	3
1.1 INTRODUCTION	3
1.2 RATIONALE FOR PRODUCING THE FORECAST	3
1.3 USE OF DEMAND FORECASTS IN GRID PLANNING	5
1.3.1 <i>Forecasts for grid modelling</i>	5
1.3.2 <i>Treatment of uncertainty</i>	6
1.4 DESIGN PRINCIPLES.....	6
1.5 DATAFLOW.....	7
1.6 PROCESS AND MILESTONES	9
2 METHODOLOGY	10
2.1 DEFINITIONS.....	10
2.2 MODELLING APPROACH.....	11
2.3 SENSITIVITIES AND SCENARIOS.....	12
3 RESULTS	13
3.1 NATIONAL PEAK FORECASTS	13
3.2 REGIONAL PEAK FORECASTS.....	15
APPENDIX A DETAILED METHODOLOGY AND INPUTS	20
A.1 HISTORICAL DEMAND DATA.....	20
A.2 STEP CHANGES	20
A.3 EFFECTS OF THE CHRISTCHURCH EARTHQUAKES	25
A.4 GDP AND POPULATION FORECASTS.....	25
A.5 THE ECONOMETRIC HLFM	26
A.6 THE ENDOGENOUS HLFM.....	28
A.7 THE AD HOC HLFM	30
A.8 THE MED-DERIVED HLFM.....	31
A.9 THE ALLOCATION METHODOLOGY	32
A.10 PLANNED IMPROVEMENTS	37
APPENDIX B DETAILED RESULTS	38
B.1 NATIONAL ENERGY	38
B.2 REGIONAL PEAKS	39
B.3 REGIONAL TROUGHS.....	47
APPENDIX C SENSITIVITIES AND SCENARIOS	48
C.1 SENSITIVITY TO RECENT CHANGES IN DEMAND	48
C.2 SENSITIVITY TO GDP FORECASTS	50
C.3 A ONE-YEAR BACKCAST	52
C.4 HEAT PUMPS	54
C.5 ELECTRIC VEHICLES	54
C.6 TIWAI DECOMMISSIONING	55
APPENDIX D VALIDATION	56
APPENDIX E TECHNICAL REVIEW BY NZIER	58
APPENDIX F TECHNICAL REVIEW BY PROFESSOR HYNDMAN	59

Executive Summary

This document presents our long-term demand forecast. **Some changes have been made to methodology and input assumptions since the May draft forecast, but the overall impact on results is minor. Canterbury forecasts for 2012-2015 are higher than in the draft.**

National and regional forecasts are included. Forecasts for individual grid exit points (GXPs) are currently in progress.

We have aimed to produce a forecast that:

- presents a credible range of uncertainty (i.e. covers the range of possible outcomes)
- is prepared and presented in a form that is suitable for grid planning activities
- is reasonably stable (i.e. it does not ‘flip-flop’ from year to year)
- accurately reflects seasonality.

The forecast achieves these goals using an ‘ensemble approach’ – i.e. it uses a combination of several prediction models, rather than trying to rely on a single “best” model.

National peak demand growth was relatively low in 2008, 2009 and 2010. **National peak demand would also have been low in 2011 to date, if it was not for the polar blast in August.** The forecast shows that moderate growth may continue – on the other hand, it is also possible that peak demand may quickly recover to pre-recession levels and continue to grow vigorously.

The new forecast of national peak demand is lower than in some other recent publications (such as the Electricity Commission’s Statement of Opportunities 2010 and our Annual Planning Reports 2010 and 2011). The SOO and APR forecasts were reasonable and consistent with the state of knowledge when they were prepared – but lower economic growth projections, and slow demand growth, have led to a downward correction.

For some regions (such as North Isthmus), prudent peak forecasts are lower than in recent Statements of Opportunities and Annual Planning Reports. For other regions (such as Nelson/Marlborough), high recent peaks and/or healthy economic growth projections have sustained forecasts at a pre-recession level.

The effects of the Christchurch earthquakes are considered (to the extent possible, given the current sketchy state of knowledge).

Heat pump, electric vehicle and early Tiwai closure scenarios are included.

Technical reviews carried out by NZIER and **Professor Hyndman of Monash University** are attached. The reviews are generally favourable, and include some recommendations for improvement.

The forecasting model is now publicly available. There is an email list for interested parties.

Comment on the approach, input assumptions and results is always welcome. Please provide any feedback to demandforecasting@transpower.co.nz.

1 Background

1.1 Introduction

This document presents a long-term demand forecast. It consists of:

- an explanation of the purpose of the forecast and how we plan to use it
- the design principles that have guided the analysis
- next steps and key milestones
- a high-level description of the forecast methodology (with more detail in Appendix A)
- an overview of the forecast results (with more detailed results in Appendix B, and sensitivities and alternative scenarios in Appendix C)
- a validation of forecasts against actual demand during winter 2011 to date (attached as Appendix D)
- a technical review carried out by NZIER (attached as Appendix E).
- a technical review carried out by Professor Hyndman of Monash University (attached as Appendix F).

National and regional forecasts are included. Forecasts for individual grid exit points (GXPs) are currently in progress.

1.2 Rationale for producing the forecast

In publishing this forecast, we seek to set out our current view on demand growth over the next 30 years (with a focus on the next decade).

We require long-term demand forecasts for a range of activities, such as:

- preparing the Annual Planning Report (APR)¹
- assessing the value of grid investments
- developing long term grid strategies such as those included in Transmission Tomorrow.²

In the last few years, we have taken the demand forecast from the Electricity Commission's Statement of Opportunities (SOO) as a base, and made modifications as necessary.

In some ways, this approach worked well. For instance, the process whereby the forecasts were modified based on customer feedback was effective – providing good information and encouraging a sense of engagement.

However, the approach also had its disadvantages – some are listed below.

- The forecasting cycle was very long. For example, the 2011 APR demand forecast was based on the 2010 SOO, which includes a demand forecast carried out in 2009. In preparing the forecast, the EC had access to data up to the end of 2008 – but because demand in 2008 was affected by a voluntary savings campaign, only data up to 2007 were used. It would be better to produce a forecast based on more up-to-date data.
- The forecast was quite unstable. Outputs varied substantially from year to year, in response to changes in input assumptions. Not only did predictions

¹ <http://www.gridnewzealand.co.nz/apr2011>

² <http://www.gridnewzealand.co.nz/transmissiontomorrow>

for individual GXPs vary, but regional and national trends also changed. This made effective grid planning difficult, and affected the credibility of the forecast.

- The treatment of seasonality was crude (for instance, it did not consider that the ratio of summer to winter demand at a GXP can change over time).

We have now developed a new purpose-built forecast, and seek feedback on the approach and results. We intend using the new forecast to help compile the APR 2012.

We are well placed to prepare demand forecasts for grid planning purposes:

- Our role as grid planner leads to specialised forecasting needs, as set out in Section 1.3. Many other agencies have needs which overlap with Transpower's, but no other organisation needs the full range, and no other organisation publishes a forecast that covers everything we require.
- Preparing the forecast inhouse allows a shorter forecasting cycle. We can now make changes quickly (as is often required – for instance, to use different regional or seasonal definitions).
- We have strong customer relationships which allow provision of accurate information on demand right down to the local level, but has a broad focus that allows us to take a long-term view of New Zealand's needs.

In some regards, the new forecasting methodology is similar to that used in previous years. For instance, we will continue to seek feedback from customers on individual-GXP forecasts, since this process regularly yields useful information.

However, we have sought to improve the approach where possible, following the design principles set out in Section 1.4.

1.3 Use of demand forecasts in grid planning

1.3.1 Forecasts for grid modelling

The main use of demand forecasts in grid planning is as inputs to models that simulate the operation of the power system, for example:

- in a detailed simulation model, typically with a short time step (e.g. DIgSILENT, PSS/E)
- in simulation models with longer time steps (e.g. GEM, SDDP).
- in economic models which calculate such things as the expected unserved energy resulting from a particular mode of failure.

The main need is for peak demand forecasts – since the capacity of the network can be tested by a high level of demand.

However, energy demand forecasts (in GWh) are also required for long term grid modelling.

Trough demand forecasts (i.e. predictions of how *low* demand may become) are also used – they can test the ability of the network to export surplus generation from a particular region or to manage voltage.

Forecasts at several spatial levels are required.

- National and regional demand forecasts can be useful – for instance, when constructing scenarios.
- Backbone studies use forecasts of demand at each GXP at island peak.
- Regional studies use forecasts of demand at each GXP at regional peak.
- Subregional studies use forecasts of demand at each of a small group of adjacent GXPs at their coincident peak.
- Local studies use forecasts of individual GXP peaks – for instance, to assess the need for new supply transformers.

Historically, the sum of annual regional peaks tends to exceed the annual national peak by about 5%.

There is a need for seasonal forecasts. Traditionally winter peak poses the most challenge to the power system, but in some regions of New Zealand, summer peak is actually higher. Also, in summer, some generating plants are unavailable (for instance, due to cooling or maintenance requirements), and transmission line ratings are reduced (lines are not allowed to carry as much power as in winter, to prevent them from sagging).

Now that shoulder line ratings have been introduced,³ separate shoulder demand forecasts are also required.

We expect further demands on the forecast as variable line ratings are introduced, but these are not covered here.

Forecasts must extend to 2027 for the APR, and in practice we predict demand out to 2040. It cannot be overemphasised that long-term demand forecasting is speculative. It is impossible to predict demand 20-30 years in the future with any certainty. Nevertheless, long-term forecasts are required.

³ <http://www.gridnewzealand.co.nz/shoulder-rating>

1.3.2 Treatment of uncertainty

An essential part of the forecasting process is to explicitly account for the uncertainty in future demand. It is not useful to seek a single “most accurate” point estimate. It is more important to describe the range of plausible outcomes, and to take this range into account when assessing investment needs.

In grid planning, one approach is to carry out a deterministic analysis, using a single demand forecast. Typically a prudent forecast would be used – i.e. one that is high on the distribution of uncertainty. This is the approach that has traditionally been taken in the preparation of the APR. Previous APRs have used a demand forecast that was set to the P90 level (i.e. the 90th percentile of the distribution of possibilities) for the first five years, and thereafter slowed to the expected growth rate. The rationale was:

- for the first five years, it would not be prudent to use an expected forecast. This reflects the asymmetry of risk (it is worse for a grid planner to underestimate transmission needs than to overestimate them) and using an expected forecast would often lead to bad outcomes. It is more optimal to use a P90 forecast
- it is not necessary to use a P90 forecast beyond the first five years. It is reasonable to plan on slower demand growth - if demand growth turns out to be rapid, and the P90 scenario becomes more plausible, then there will still be time to adjust the forecast upward and implement solutions.

Another approach is to use stochastic analysis, incorporating the uncertainty of the demand forecast. In its simplest form, this might entail using three scenarios (high, medium and low demand) and assigning them appropriate weights. A larger number of scenarios can be used if appropriate.

The forecasts in this document include expected, P90 and prudent projections.⁴ The prudent forecast is equal to the P90 for the first five years (until 2017, inclusive) and grows at the same rate as the expected forecast thereafter. We consider this an appropriate basis for use in future APRs.

Sometimes a load duration curve or load probability curve is used in stochastic analysis – measuring not only the highest, lowest or mean demand, but the entire cumulative distribution of demand over some period. This document does not include projected load duration curves, but it is anticipated that they will be incorporated in the 2012 forecast. In the interim, simple load duration curve projections can be produced by scaling an historical reference curve.

1.4 Design principles

The design principles that we seek to follow are set out below:

represents the range of uncertainty – the forecast must describe the range of possible outcomes, rather than having an unrealistic expectation of pinpoint accuracy

fit for purpose – the forecast must be prepared and presented in a form that is suitable for grid planning activities

stable – the forecast should not change unduly from year to year, and should not be excessively sensitive to any one input parameter

seasonal – the forecast must accurately represent seasonal trends, to the extent possible given available information (this follows from the ‘fit for purpose’ principle).

⁴ Throughout this document, the words ‘forecast’ and ‘projection’ are used interchangeably.

We have used an **ensemble approach**. The use of an ensemble – a suite of models used to get multiple viewpoints on an issue – is innovative in the long-term electricity demand forecasting arena, but is common practice in other disciplines such as meteorology.⁵

The basic principle of ensemble forecasting is that it is not possible to produce a single correct model, because not only are the model parameters not known, but the true model structure is uncertain. In many contexts it has been shown to be more accurate to produce several models with different structures, and combine the results. The simplest way to combine the results is averaging – taking the mean of all the model outputs to produce a point estimate – but it is better to use the ensemble of models to produce a distribution of uncertainty. The ensemble approach supports the stability objective (because each member of the ensemble uses different inputs, so the ensemble as a whole is not overly sensitive to a change in any one input)

The ensemble approach also supports the representing-the-range-of-uncertainty objective (because it allows for uncertainty, not only in the inputs fed into the demand forecasting model, but also as to which model is ‘correct’).

We have prioritised **independent summer and winter forecasts**. In earlier forecasts, summer and winter projections were typically linked – it was assumed, for instance, that the ratio of summer peak to winter peak would remain constant for each grid exit point. The new approach produces independent (but consistent) forecasts of summer and winter peak.

Following the introduction of transmission line shoulder ratings, a **shoulder forecast** has also been added. As the enhanced line ratings system develops further, it may be useful to produce an even more detailed representation of seasonality.

In terms of process, we place high importance on **industry engagement**. We will seek feedback on national and regional forecasts on an annual basis, and will consider responses carefully. The main respondents are expected to be generator-retailers, major electricity consumers, network companies and regulators – but comments from other parties will be gratefully received.

We will discuss GXP-level forecasts with customers directly, recognising that they have the most knowledge when it comes to potential demand changes at the local level.

In order for the industry to engage effectively, there will need to be **transparency** about how the forecast is produced. We will aim to provide as much information as possible about the process (while protecting the confidentiality of any parties who provide sensitive information).

1.5 Dataflow

The flow of data through this process is shown overleaf.

⁵ Some good references on the ensemble approach:

http://en.wikipedia.org/wiki/Ensemble_forecasting

<http://faculty.fuqua.duke.edu/~clemen/bio/Published%20Papers/13.CombiningReview-Clemen-IJOF-89.pdf>

(a little dated, but arguably still the best review of the field)

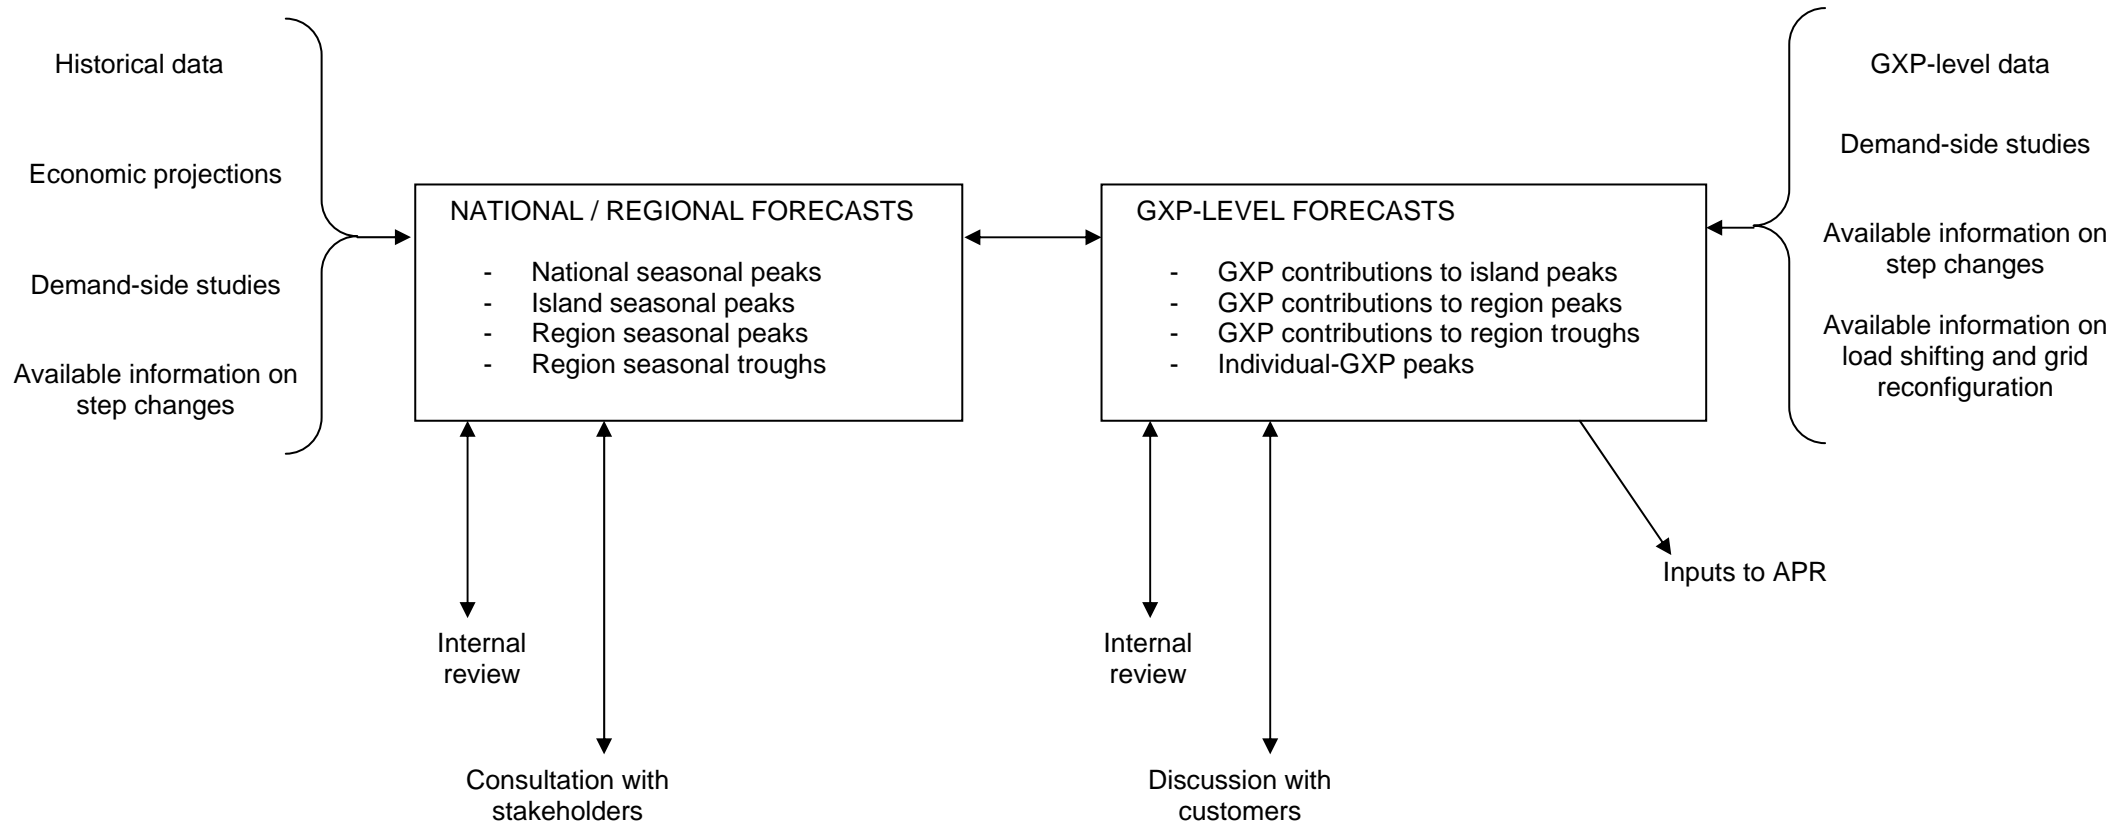
<http://www.jstor.org/pss/3008764> (the paper that started it all in 1969)

<http://www.nber.org/papers/w3965> (probably the longest running ensemble forecast)

<http://www.metoffice.gov.uk/research/areas/data-assimilation-and-ensembles/ensemble-forecasting/MOGREPS>

(an application in meteorology)

DATAFLOW DIAGRAM



1.6 Process and milestones

We have:

- prepared draft national and regional forecasts and published them for comment
- published the demand forecasting model code, inputs and outputs at <http://code.google.com/p/ledfm> – both to provide transparency, and to allow other agencies to use the model if they so wish
- set up a mailing list for interested parties (if you wish to be added to this list, please email demandforecasting@transpower.co.nz)
- commissioned reviews of the forecasting methodology by NZIER and Professor Hyndman of Monash University
- finalised the national and regional forecasts and published them (in this document)
- largely completed preparation of individual-GXP forecasts, including discussion with customers.

The APR 2012, incorporating the new demand forecasts, will be published in the autumn of 2012.

We will update the forecast in the winter of 2012 and seek industry feedback on the revised forecast.

If time permits, we will also institute an annual validation report that compares predictions from previous years with actual demand. Validation is an important part of the forecasting cycle and can be used to drive improvements in the methodology.

2 Methodology

2.1 Definitions

The forecast covers **annual / seasonal, national / island / regional, P90 / prudent / expected half-hourly peaks** for 2011-2040.

The forecast also covers **seasonal, night / day, regional troughs**.

Calendar years are used throughout. The forecast is based on **demand at GXP level, net of embedded generation**.⁶

Seasons	Consistent with transmission line ratings. ⁷ Winter is 7 a.m. 10 May – 7 a.m. 20 Oct. Shoulder is 7 a.m. 20 Oct – 7 a.m. 1 Dec and 7 a.m. 15 Mar – 7 a.m. 10 May. Summer is 1 Jan – 7 a.m. 14 Mar and 7 a.m. 1 Dec – 31 Dec.
Night and day	Consistent with transmission line ratings. Night is 9 pm to 7 am. Day is 7 am to 9 pm.
National peak	Maximum national demand in a year or season.
Island peak	Maximum island demand in a year or season (not to be confused with island contribution to national peak).
Regional peak	Maximum regional demand in a year or season (not to be confused with regional contribution to national or island peak). Regions are consistent with the APR 2011 ⁸ – note that these differ in minor respects from the SOO 2010 and APR 2010.
Regional trough	Minimum regional demand in a year or season. Can be important when studying export constraints out of a region.
P90 forecast	For peak, the 90 th percentile of the distribution of outcomes (10% POE). For trough, the 10 th percentile of the distribution (90% POE).
Expected forecast	The mean (not median) of the distribution of outcomes.
Prudent forecast	Equal to the P90 forecast for the first five years (until 2017 inclusive). Thereafter, grows at the same growth rate as the expected forecast. This is what has traditionally been used in the APR.
Half-hourly peak	Also known as 30 minute peak. The mean demand in the trading period with maximum mean demand (as opposed to instantaneous peak).
Demand at GXP level	Demand including distribution losses but excluding transmission losses.
Net of embedded generation	Gross demand minus the output of embedded generation. When a GXP is injecting power onto the grid, the demand 'net of embedded generation' is a negative number. Note that this is a new approach – in previous APR and SOO demand forecasts, a net injection was treated as a zero rather than a negative value. The new approach is a more appropriate basis for modelling. The downside is that it makes it harder to compare results with previous forecasts – an adjustment is needed.

⁶ Also known as (net) grid offtake.

⁷ The enhanced line ratings system is still developing – see <http://www.gridnewzealand.co.nz/innovative-technology>.

⁸ <http://www.gridnewzealand.co.nz/apr2011>

To be clear – this means that if the forecast is used in a supply/demand model, embedded generation (such as Glenbrook) should not be modelled on the supply side.

Do make sure to model relatively minor grid-connected plants like Argyle, Kapuni and Whareroa.

Waipori is a corner case – the grid-connected part of Waipori generation should be modelled, but not the embedded part.

2.2 Modelling approach

An ensemble of ‘high level’ models is used to produce a top-down national forecast. This forecast is then allocated down to the regional level. The emphasis is on modelling uncertainty wherever it occurs – yielding a distribution of outcomes, rather than a point estimate.

The full process is as follows:

1. collate input data, including:
 - historical demand (Appendix A.1)
 - potential step changes (A.2 for the overall approach and A.3 for the effects of the Christchurch earthquakes – **both now updated**)
 - forecasts of population and GDP (A.4)
2. develop a suite of high level forecasting models (HLFMs). Each HLFM produces randomised forecasts of annual national peak and energy for 2011-2040.

Currently there are four HLFMs – in no particular order:

- the econometric model, which uses forecasts of population and GDP (A.5)
- the endogenous model, a “ruler”-type regression approach (A.6)
- the ad hoc model, based on expert judgement of Transpower staff (A.7)
- the MED-derived model, which is adapted from projections in the Ministry of Economic Development’s Energy Outlook (A.8).

Together, these HLFMs form the ensemble. **(Peak-to-energy ratios have been reassessed in all four HLFMs, as recommended by several reviewers. Changes are documented in the relevant sections of the Appendix.)**

3. develop an allocation methodology (AM) which allocates the national forecasts down to regions (A.9). The AM produces all the outputs required:
 - seasonal mean demand
 - annual / seasonal, national / island / regional half-hourly peaks
 - seasonal, night / day, regional troughs.

In other words, the AM bridges the gap between the quantities that the HLFMs produce and the quantities that the user needs.

4. repeat the following process N times (for some large N):
 - randomise key inputs (step changes, GDP and population)
 - select a random HLFM and use it to produce one randomization of the national forecasts, based on the randomised inputs
 - use the AM to produce one randomization of all the outputs required, based on the national forecasts *and* the matching randomised inputs
 - record the results.
5. for each output of interest (e.g. Auckland regional peak in winter 2017), collate the distribution of that output across the N randomizations.
6. Use this distribution to determine the P90, expected, and prudent predictions.

For the avoidance of doubt, we can access the full distribution of any output if this is required for simulation work – we are not limited to using P90/expected/prudent summaries only.

Some key requirements are that:

1. the ensemble of HLFMs must:
 - be easy to understand
 - produce a sufficient spread of outcomes to represent the true uncertainty
 - use diverse inputs, so that the ensemble as a whole is not too sensitive to changes in any one input
2. the AM must:
 - produce the full range of outputs required
 - produce plausible seasonal and regional patterns
 - model the reasonable range of uncertainty at each stage
 - again, avoid undue sensitivity to changes in any one input.

A key influence in identifying the key requirements and developing the approach was a review of the Electricity Commission's demand forecast carried out by Statistics Research Associates and published in 2010.⁹ The reviewer, Dr Peter Thompson, recommended the use of "a suite of competing forecasting models" (our ensemble). The review included a list of other recommendations, some of which we have already taken up, and others of which will guide future development.¹⁰

The development work carried out by Brian Kirtlan of the Electricity Authority (formerly Electricity Commission) over recent years is also gratefully acknowledged.

A list of planned improvements to the methodology, largely based on the NZIER and Hyndman reviews, is included as Appendix A.10.

2.3 Sensitivities and scenarios

Three sensitivities, intended to test the stability of the forecast, are included:

- Appendix C.1 demonstrates that the forecast is not unduly sensitive to recent changes in (national or regional) peak demand
- C.2 demonstrates that the forecast is not unduly sensitive to changes in (national or regional) GDP forecasts
- C.3 presents a one-year backcast, demonstrating that if the new methodology had been carried out a year ago, the results would not have been unduly different from the present forecast.

We will consider variations on the base forecast as required for a particular application. However, three particular scenarios are often discussed and it seems worthwhile to include them at this point:

- Appendix C.4 models an increase in the penetration and use of heat pumps
- C.5 models the effect of electric vehicles
- C.6 models the effect of decommissioning of the NZAS aluminium smelter at Tiwai Point in the 2020s. (We have no particular reason to expect this scenario, but have not ruled it out as impossible.)

⁹ <http://www.ea.govt.nz/document/13381/download/industry/modelling/demand-forecasting/demand-forecast-review-documents/>

¹⁰ Though any deficiencies in the work should be attributed to Transpower rather than SRA.


3 Results

3.1 National peak forecasts

The national peak forecast is shown in Figure 1. Figure 2 demonstrates how the expected and P90 forecasts were produced, by combining the four HLFMs.

For completeness, the national energy forecast is included in Appendix B.1.

Figure 1: Annual national peak forecast (updated)




The distribution of plausible outcomes is quite broad (“thin-tailed”) – some extreme scenarios are well above the P90 level.

In other words, demand could potentially grow very fast or not at all – but a more moderate scenario is more likely.

There has been very little change in annual national forecasts since the draft report.


Figure 2: Four HLFMs and the resulting ensemble forecast (updated)



The ensemble P90 forecast is greater than the MED and endogenous P90 forecasts, but less than the ad hoc and econometric P90 forecasts.

For comparison, some other recent forecasts of national peak demand are shown in Figure 3. (All forecasts shown are prudent, i.e. P90 for the first five years, slowing to the expected growth rate thereafter. Small corrections have been made for different treatment of embedded generation.)

Figure 3: Various prudent forecasts of national peak demand (updated)



The earlier forecasts were a reasonable reflection of the state of knowledge at the time. However, the new forecast is lower, based on several years of low peak demand, and lower projected economic growth.

The low peaks in 2008, 2009 and 2010 deserve some explanation. Demand growth over this period was well below the usual 1.5-2% range. This was driven by a range of factors, including:

- the 2008 dry winter and associated high prices
- reduced demand at the Tiwai smelter, following the 2008 transformer failure
- the recession
- new embedded generation offsetting demand growth
- weather conditions
- higher residential electricity prices
- increased uptake of electricity efficiency measures
- increased response to spot prices by major industrial consumers
- the use of load control.


It is hard to determine which of these factors are most important.

The forecast acknowledges that peak demand may rebound quickly from the 2008-2010 slump, but indicates that, in the next few years, peaks are unlikely to rise to the levels that were contemplated before the recession.

3.2 Regional peak forecasts


Some example forecasts are shown here, to illustrate the approach. Peak forecasts for all regions are provided in Appendix B.2; trough forecasts in B.3. **Except where noted, changes in forecasts since the May draft are relatively minor.**

Figure 4: Annual regional peak forecast – Auckland (updated)




The Auckland forecast is for reasonably strong demand growth, driven by high GDP and population growth projections.

Figure 5: Annual regional peak forecast – Taranaki (updated)




The Taranaki forecast also predicts high growth in the next few years, but this is driven by step changes rather than economic projections.

Figure 6: Annual regional peak forecast – Hawkes Bay (updated)



The Hawkes Bay forecast is for much slower growth – because no major step changes are predicted, projected GDP growth is low, and demand growth has been slow in the last few years.

Figure 7: Annual regional peak forecast – Canterbury (updated)



The Canterbury forecast is driven by the effects of the Christchurch earthquakes, offset in the mid term by strong agricultural growth.

The new forecast is higher than the May draft, for the next few years (see Appendix A.3).

For comparison, some recent forecasts of regional peak demand are shown below. (All forecasts shown are prudent, i.e. P90 for the first five years, slowing to the expected growth rate thereafter. Small corrections have been made for different treatment of embedded generation.)

Figure 8: Various forecasts of North Isthmus peak demand (updated)




Figure 9: Various forecasts of Auckland peak demand (updated)




Figure 10: Various forecasts of Bay of Plenty peak demand (updated)




Figure 11: Various forecasts of Wellington peak demand (updated)




Figure 12: Various forecasts of Nelson/Marlborough peak demand (updated)





Figure 13: Various forecasts of Canterbury peak demand (updated)



Appendix A Detailed methodology and inputs

A.1 Historical demand data

The main historical demand dataset used in this forecast was extracted from the Electricity Authority's Centralised Dataset¹¹, and covers the period from 1997 to 2010 inclusive. Demand is measured at GXP level, net of embedded generation. Net injection is treated as negative demand.

The query and dataset will be provided on request. **One correction has been made since the draft version of this report – Moturoa demand (in Taranaki) had been omitted for the second half of 2010, and is now correctly included.**

The econometric HLFM was fitted using a longer-term demand dataset (A.6).

A.2 Step changes

Step changes in the last few years have included load reductions at some major North Island pulp mills, various new loads in Taranaki, West Coast mining development, dairy development in South Canterbury, and the construction of the White Hill wind farm. The AM knows to separate these changes out when it assesses and projects underlying growth in the relevant regions.

Possibly the largest step change in the next few years will be load growth at the NZAS aluminium smelter at Tiwai. The Lower South Island Reliability Transmission Investment Proposal included a high projection of Tiwai load growth ("LSI Tiwai projection"), with smelter load eventually increasing to 690 MW. In this forecast, we do not assume that smelter load will follow the LSI Tiwai projection. Rather, we assume Tiwai demand to be uniformly distributed, with lower bound of 620 MW and 90th percentile equal to the LSI Tiwai projection.

A list of other potential step changes in the next decade is shown overleaf. The list includes new loads, changes at existing sites, and new embedded generation. These step changes are used in the AM and some (not all) HLFMs.

Note that the effects of centralised demand-side initiatives such as the potential Upper North Island DSM trial are not included.

Step changes are randomised in the forecasting process. In any particular randomization, a given change can occur on time, late or not at all. Each batch of related step changes is assigned a probability of occurrence, and each individual step change is assigned 'earliest possible date' and 'maximum deferral' parameters – the model will randomly choose a commissioning date in this range.

Beyond 2015 and especially beyond 2020, there will be more step changes that we cannot even anticipate yet. The model attempts to allow for these by adding step changes at random. In each year beyond 2020, the following step changes are created and distributed randomly between regions:

- about three new small loads (5 MW each)
- about a 25% chance of a new large load (with size lognormally distributed, usually averaging 50 MW, but no more than 20% of the region's 2010 peak)
- about a 50% chance of a new embedded generator (with size lognormally distributed, usually averaging 10 MW, but no more than 20% of the region's 2010 peak).

Generic step changes are also added between 2015 and 2020, but with lower probability.

¹¹ <http://www.ea.govt.nz/industry/modelling/cds/>

A list of possible step changes in the next decade is shown below. These are used in the AM, and some (not all) of the HLFMs. The figures shown are very approximate, and any feedback enabling us to refine these assumptions would be most welcome.

(*) Starred load factors are used without modification at national peak, but reduced at regional peak (on the basis that the regional peak is unlikely to occur when the relevant embedded generator is at medium to high output).

Nature	GXP	Region	Earliest possible year	MW	Probability	Could be up to this many years late	Load factor – winter peak	Load factor - summer peak	Mean load factor – year round
Meremere load shift to Huntly	MER	Auckland	2013	-10	1	1	0.9	0.7	0.8
	HLY	Waikato	2013	10					
Te Uku wind farm	TWH	Waikato	2011	-64	1	1	(*) 0.4	(*) 0.4	0.4
Misc load	WHU	Waikato	2012	5	0.7	3	1	1	1
Misc load	KIN	Bay_of_Plenty	2011	4	0.7	3	1	1	1
Geothermal cogen	KAW	Bay_of_Plenty	2011	-15	0.7	2	0.5	0.5	0.5
Te Rere Hau wind farm	LTN	Central_Districts	2011	-15	1	1	(*) 0.4	(*) 0.4	0.4
Kupe gas processing	HWA	Taranaki	2011	12	1	0	0.9	0.9	0.9
Ahuroa gas	SFD	Taranaki	2011	4	1	1	0.8	0.8	0.8
	SFD	Taranaki	2012	4					

Nature	GXP	Region	Earliest possible year	MW	Probability	Could be up to this many years late	Load factor – winter peak	Load factor - summer peak	Mean load factor – year round
Dairy - crops - industrials	MST	Wellington	2011	7	0.5	2	0.5	0.8	0.5
New activity at Pike River	ATU	West Coast	2013	10	0.5	3	0.9	0.9	0.7
Amethyst hydro	HKK	West Coast	2014	-6	0.7	2	0.5	0.5	0.5
Westland Dairy	HKK	West Coast	2014	4	0.7	2	0.3	0.9	0.5
Mining	ORO	West Coast	2013	5	0.7	1	0.9	0.9	0.7
Mining	DOB	West Coast	2015	4	0.7	2	1	1	1
Mining	DOB	West Coast	2018	4	0.7	3	1	1	1
Dairy processing	HOR	Canterbury	2011	2	0.7	2	0.3	0.9	0.5
	HOR	Canterbury	2014	4					
	HOR	Canterbury	2016	4					
Dairy processing	HOR	Canterbury	2012	5	0.7	2	0.3	0.9	0.5
	HOR	Canterbury	2014	3					
Irrigation	CUL	Canterbury	2012	3	0.7	2	0	0.9	0.4
	CUL	Canterbury	2016	3					

Nature	GXP	Region	Earliest possible year	MW	Probability	Could be up to this many years late	Load factor – winter peak	Load factor - summer peak	Mean load factor – year round
Irrigation	ASB	Canterbury	2012	4	0.7	2	0	0.9	0.4
Irrigation	SBK	Canterbury	2012	3	0.7	2	0	0.9	0.4
Dairy processing	SPN	Canterbury	2013	10	0.5	1	0.3	0.9	0.5
Irrigation	WPR	Canterbury	2017	10	0.7	3	0	0.9	0.4
Dairy processing	STU	South Canterbury	2011	3	0.7	2	0.3	0.9	0.5
Irrigation & industrial	TIM	South Canterbury	2011	3	0.7	2	0.5	0.8	0.5
	TIM	South Canterbury	2013	3					
Irrigation	BPT	South Canterbury	2013	7	0.7	2	0	0.9	0.4
Irrigation	BPD	South Canterbury	2014	9	0.7	3	0	0.9	0.4
Dairy processing	STU	South Canterbury	2014	5	0.7	2	0.3	0.9	0.5
	STU	South Canterbury	2017	4					
Holcim cement	OAM	South Canterbury	2014	13	0.7	3	0.8	0.8	0.8
	WPT	West Coast	2014	-5					
Associated with Holcim	OAM	South Canterbury	<i>(Linear increase - from nil in 2011 to 10 MW in 2020)</i>						

Nature	GXP	Region	Earliest possible year	MW	Probability	Could be up to this many years late	Load factor – winter peak	Load factor - summer peak	Mean load factor – year round
Irrigation	STU	South Canterbury	2017	35	0.7	2	0	0.9	0.4
	STU	South Canterbury	2018	10					
Dairy processing	TMK	South Canterbury	2018	5	0.7	2	0.3	0.9	0.5
Misc dairy development	STU	South Canterbury	<i>(Linear increase - from nil in 2011 to 10 MW in 2020)</i>		0.7	2	0.3	0.9	0.5
Misc dairy development	OAM	South Canterbury	<i>(Linear increase - from nil in 2011 to 6 MW in 2016)</i>		0.7	2	0.3	0.9	0.5
Mahinerangi 1 wind farm	HWB	Otago/Southland	2011	-36	1	1	(*) 0.4	(*) 0.4	0.4
Dairy processing	GOR	Otago/Southland	2012	7	0.7	2	0.3	0.9	0.5
Solid Energy briquettes	GOR	Otago/Southland	2014	20	0.7	1	0.75	0.75	0.75
Solid Energy lignite to urea	GOR	Otago/Southland	2016	20	0.5	3	0.9	0.9	0.8
Dairy processing	EDN	Otago/Southland	2016	4	0.7	2	0.3	0.9	0.5

A.3 Effects of the Christchurch earthquakes

The devastating earthquakes of 4 September 2010, 22 February 2011 and 13 June 2011 are not long behind us. It is still difficult to determine how the rebuild will proceed.

We have taken an initial view on the potential effects of the earthquake on electricity demand, and will continue to update this view on an ongoing basis.

In the short term, it is clear that Christchurch demand is substantially reduced. We expect that load will return to more normal levels over the next few years (though some specific areas may never return to pre-earthquake projections).

We speculate that relocation of businesses out of Christchurch may lead to demand increases in some other centres. But hard data about relocation are not yet available.

We have adjusted the forecast accordingly, and will continue to review the situation.

In the final version of this forecast we have halved the assumed reduction in demand in Canterbury. This adjustment was made following the high peaks seen in the Orion area during the August 2011 polar blast (Appendix D). It appears that the region can still experience high demand peaks in very adverse weather conditions, notwithstanding the effects of the earthquakes.

A.4 GDP and population forecasts

Regional GDP and population forecasts are key inputs of the econometric HLFM (Appendix A.5) and the AM (A.9).

The base forecasts come from NZIER and Statistics New Zealand respectively. However there is considerable uncertainty about both quantities, and it is important to factor this into the forecast.

In each of the N iterations of the forecast process, regional GDP and population are randomised around the base values.

First, population is randomised – it can be set to Statistics New Zealand's low national forecast, high national forecast, or (most likely) somewhere in between.

Second, GDP is randomised. Four different sources of uncertainty are modelled:

- national interdecadal variation. This is assumed to be correlated with population growth (so a "high population" scenario cannot be a "low GDP" scenario);
- national AR (2) MA (1) noise;
- regional interdecadal variation; and
- regional white noise.

This is a complicated procedure but it does produce a pattern of variation very similar to what has been observed in the past.

Note that the effects of the Christchurch earthquakes are not factored into the Canterbury regional forecasts of population or GDP. Rather, these effects are modelled separately (as described in Appendix A.3).

A.5 The econometric HLFM

The econometric model is based on the observed relationship between annual national electricity demand less Tiwai (GWh), national population and national GDP since 1975.

This relationship changed around 1990, when electricity intensity started to decline (i.e. given the observed rates of GDP and population growth, demand growth was lower than it would have been in the 1970s or 1980s).

A regression model that incorporates GDP, population and the decline in electricity intensity since 1990 can explain almost all the historical variability in annual national demand.

In projecting electricity demand forwards, we assume that population and GDP growth will be associated with demand growth – but we also consider that the relationship between econometric variables and demand may change over time. Electricity intensity may decline further, or may begin to increase again.

The resulting HLFM, then, is:

$$E_y = \exp(-2.695 + 0.5 \log(P_y) + 0.463 \log(G_y) + I_y + \text{white noise } (\sigma = 0.012)) + T_y$$

where:

E_y is annual national energy demand (GWh) in year y ,

P_y and G_y are (randomised) projections of population and GDP (95/96 \$M) respectively,

I_y is a randomised intensity effect, equalling $\sum_{i \leq y} \Delta_i$, where $\Delta_i = \Delta_1$ during the 2010s, $\Delta_i = \Delta_2$ during the 2020s, and $\Delta_i = \Delta_3$ thereafter, and $\Delta_1 \dots \Delta_3$ are iid $N(0, 0.01)$,

and T_y is projected Tiwai energy demand (GWh) in year y .

The coefficients of the E_y equation are based on a regression fit.

The coefficient of P_y was held fixed at 0.5. (Because population and GDP are highly correlated, there is a wide range of (P_y, G_y) pairs with high historical explanatory power. It is impossible to tell which is 'correct'. Forcing $P_y = 0.5$ is intuitively pleasing and gives population and GDP roughly equal weight. The Electricity Commission achieved a similar effect by splitting demand into residential – driven largely by population – and commercial – driven largely by GDP. We do not do this because we can get the same effect using a single model, and because we do not trust the historical data relating to the breakdown of demand into residential, commercial and industrial sectors.)

Note that step changes other than Tiwai growth are not incorporated into this HLFM – the assumption is that the dynamics around new loads and new embedded generation are broadly captured by the long-term econometric relationship.

An additional step is required to convert the resulting energy forecast (GWh) into a peak forecast (MW).¹² Peak is more variable than energy, so additional variation is injected at this point:

$$D_y = E2P_y (E_y - T_y) * (1 + \text{white noise } (\sigma = 0.01)) + \text{Tiwai peak demand}$$

where D_y is annual national peak demand (MW) in year y , and $E2P_y$ is an energy-to-peak ratio, normally distributed with mean 0.178 and standard deviation increasing linearly from nil in 2011 to 0.01 in 2040. This variation in peak growth relative to

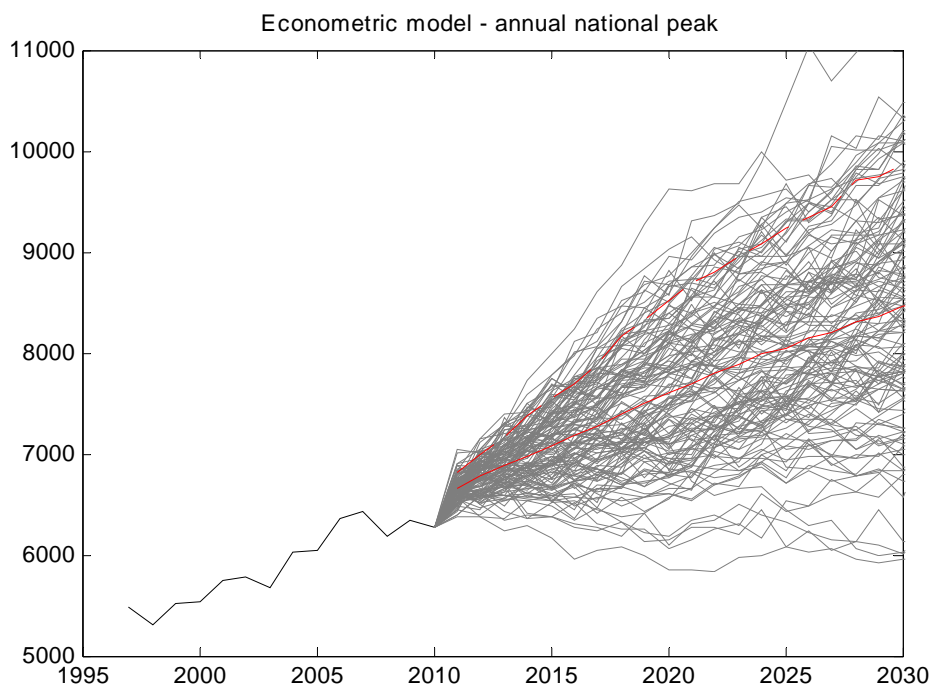
¹² It would have been preferable to forecast peak demand directly, rather than forecasting energy demand and then converting energy to peak. However, there is not a long time series of national peak demand data. We have tried and failed to construct such a time series from available records.

energy growth was introduced in response to advice by reviewers. It certainly seems reasonable to consider that peak and energy growth could decouple, particularly in the longer term.

The constant of 0.178 (formerly 0.18) is the mean of annual ratios of non-Tiwai peak to non-Tiwai energy in 2006, 2007, 2009 and 2010 (2008 was not used as it was affected by a savings campaign).

In order to fit this HLFM, it was necessary to construct a time series of annual national demand (GWh) dating back to 1975. Recent data were sourced from the Electricity Authority’s Centralised Dataset; older data from the Ministry of Economic Development’s Energy Datafile. However, the Energy Datafile figures measure a quite different quantity from the Centralised Dataset series, and a calibration factor was applied to bring the two series onto the same (approximate) scale.

The plot below (updated) shows a sample of the projections produced by the econometric HLFM. The plotted quantity is D_y , annual national peak demand. The two red lines are the mean and P90 of the projections.



This HLFM produces a wide range of outcomes, largely as a result of the randomization of population, GDP and energy intensity.

NZIER queried the electricity intensity effect: “In the draft forecast report, the effects of energy intensity are assumed to have a normal distribution with a mean of zero. Shouldn’t the mean of this effect be negative rather than zero in order to indicate that there is downward tendency in electricity consumed per unit of GDP? We have indeed observed that since 1990, the electricity intensity has been declining.”

We agree with NZIER that the econometric model should predict an expected decline in energy intensity. In fact this is already the case. In this regard the formulation is perhaps misleading – while the I parameter does have a mean of zero, the choice of P and G coefficients and the future distributions of covariates tend to lead to electricity intensity decreasing over time.

The expected peak demand growth rate – about 1.25% over 20 years – is well below the expected GDP growth rate – more like 1.9%. (In per capita terms, these rates become 0.5% and 1.2% respectively.) There are some realizations where peak demand growth is higher than GDP growth, but they are a minority.

So the status quo seems appropriate in this regard.

A.6 The endogenous HLFM

The endogenous model does not use econometric predictors, and is based on the historical trend of national peak demand growth since 1997.

A simple approach would be to fit a linear regression on peak demand versus time. However, this might give misleading results because:

- the historical trend is influenced by savings campaigns in dry years
- the historical trend is influenced by changes in industrial loads and embedded generation, which may not continue at the same rate
- potential future step changes should be considered.

We therefore fit the regression, but:

- omit savings campaign years
- separate out identifiable industrial loads and embedded generators before fitting the regression, and project these components of demand separately
- adjust the results for potential step changes.

On first principles one would expect that the effect of weather (or temperature) on demand would be significant and should be modelled. This is the approach that was taken in the demand forecast for Transpower's Annual Security Assessment.¹³ An earlier version of the endogenous model attempted to correct the historical demand series for the effects of temperature (on a by-half-hour-by-region basis), fit the regression model on temperature-corrected peaks, project it forward, and then reinject the temperature-based variability. However, this fairly complicated process:

- had only a modest effect on the results of the endogenous HLFM
- had very little effect on the overall results of the ensemble
- raised some questions about the effectiveness of the temperature correction process in regions where peak demand is heavily affected by load control (e.g. Canterbury).

Accordingly, temperature effects are now not included in the model.

There is considerable doubt as to whether the recent downturn in peak demand (2008-2010) will be sustained, or whether demand will rebound quickly to the level that would have been expected before the recession. Peak demand in 2010 was particularly low (relative to expectations), and if the model is fitted with 2010 included, this tends to bring down projected growth. Therefore, two cases are considered:

- a low probability case with 2010 removed, leading to higher growth; and
- a higher probability case with 2010 left in, leading to lower growth.

Because the true trend is uncertain, each randomization takes a sample from the confidence envelope of the regression parameters and uses this to project forwards.

The resulting HLFM, then, is:

$$D_y = a_i + b_i(y - \bar{y}) + \text{white noise}(\sigma) + \text{randomised industrial and embedded} \\ + \text{step changes}$$

where:

$$D_y \text{ is annual national peak demand (MW) in year } y, \\ a_i \sim N(\hat{\alpha}, \sigma_{\hat{\alpha}}) \text{ and } b_i \sim N(\hat{\beta}, \sigma_{\hat{\beta}}),$$

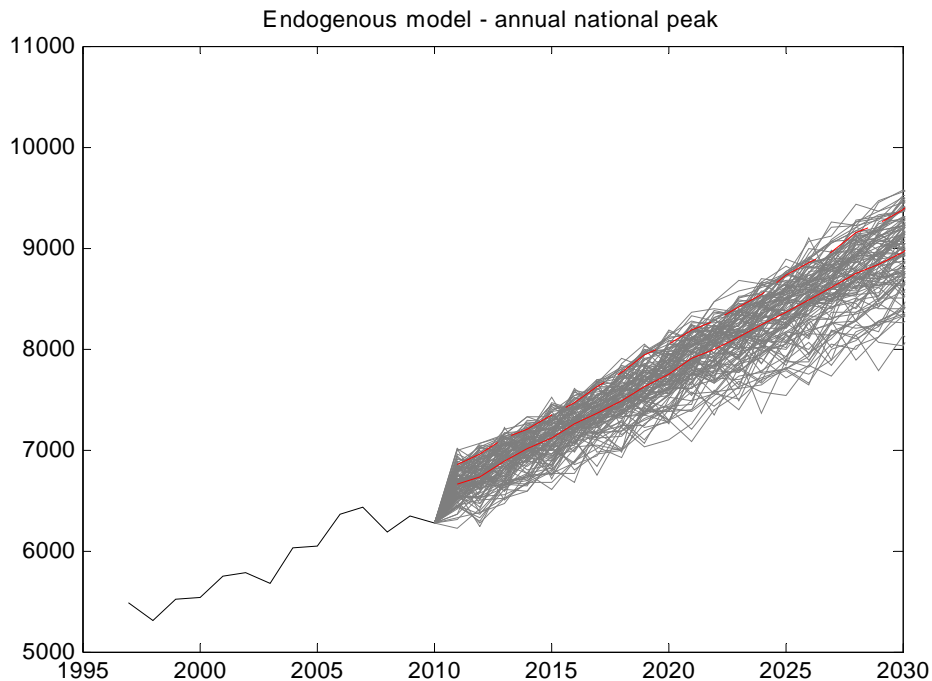
¹³ http://www.systemoperator.co.nz/f4571.43932449/Security_of_supply_medium_term_demand_forecast_in_2010.pdf

and with probability 0.5, $(\hat{\alpha}, \sigma_{\hat{\alpha}}, \hat{\beta}, \sigma_{\hat{\beta}}, \bar{y}, \sigma) = (4972, 29, 93, 6.7, 2003.4, 134)$
 (based on a fit including 2010)

or with probability 0.5, $(\hat{\alpha}, \sigma_{\hat{\alpha}}, \hat{\beta}, \sigma_{\hat{\beta}}, \bar{y}, \sigma) = (4928, 21, 102, 5.5, 2002.7, 117)$
 (based on a fit excluding 2010)

An accompanying energy forecast (GWh) is produced similarly, though with less white noise. **Since the draft forecast, we have increased the magnitude of this white noise. We considered allowing the peak-to-energy ratio to drift over time, but decided not to – because the model results are so tightly distributed (due to the assumption of linear growth) that even if we allow for peak/energy drift, the distribution of forecasts will still be closely ‘packed’ relative to the other HLFMs.**

The plot below (updated) shows a sample of the projections produced by the endogenous HLFM. The plotted quantity is D_y , annual national peak demand. The two red lines are the mean and P90 of the projections.



This HLFM produces a much narrower range of outcomes than the econometric HLFM.

A.7 The ad hoc HLFM

The ad hoc model is based unapologetically on judgements made by our staff, and has no theoretical backing.¹⁴

It focuses on how the country will emerge from the recession and the associated dip in electricity consumption since 2008. Three possible (not exhaustive) world views are put forward:

- high growth (20% weight) – underlying peak demand (less Tiwai) will quickly return to a “2007 peak plus 2.5% p.a.” path, slowing to 2% p.a. after 2020
- medium growth (60% weight) – underlying peak demand (less Tiwai) will continue on a “2009 peak plus 2% p.a.” path, slowing to 1.5% p.a. after 2020
- low growth (20% weight) – underlying peak demand (less Tiwai) will languish on a “2009 peak + 0.5% p.a.” path, accelerating to 1% p.a. after 2015.

Note that step changes other than Tiwai growth are not incorporated into this HLFM.

For each randomization, the HLFM first picks a world view at random and establishes “underlying” peak demand growth U_y , then:

$$D_y = U_y * (1 + \text{white noise } (\sigma_D)) + \text{Tiwai contribution to peak}$$

$$E_y = U_y / 0.178 * (1 + \text{white noise } (\sigma_E)) + T_y$$

where:

- D_y is annual national peak demand (MW),
- E_y is annual total energy demand (GWh),
- T_y is projected Tiwai energy demand (GWh) in year y ,
- σ_D increases linearly from 0.02 in 2011 to 0.04 in 2030,
- σ_E increases linearly from 0.01 in 2011 to 0.02 in 2030.


The rationale for the peak-to-energy ratio of 0.178 is set out in A.5.

The white noise standard deviations are loosely based on recent historical data. The modelled increase in these standard deviations over time results in a more plausible (less trimodal) distribution of outcomes.

This plot (updated) shows a sample of the projections produced by the ad hoc HLFM.

The plotted quantity is D_y , annual national peak demand. The two red lines are the mean and P90 of the projections.

The outcomes of this HLFM simply reflect the input assumptions – demand growth may be fast, slow or somewhere in between.



¹⁴ This model should be considered as a placeholder. For future years, the intention is to carry out a formal Delphi forecast – that is, to use a structured communication technique to elicit expert opinions on demand growth. The exercise might be limited to Transpower staff, but ideally would include experts from other organisations as well.

A.8 The MED-derived HLFM

The MED-derived model is based on projections from the Ministry of Economic Development’s Energy Outlook (www.med.govt.nz/energyoutlook). We consider it important to incorporate MED’s views into the forecasting process.

However, some modifications are necessary because MED’s forecasts:

- are expressed in terms of annual national electrical energy consumption (PJ or GWh), rather than peak demand
- are expected forecasts that do not consider random year-to-year variation.

The process is to:

- randomly pick one of MED’s projections (with probabilities arbitrarily set to high-growth 20%, base case 60%, low-growth 20%)
- extract the projected growth rate (expressed relative to the 2005-09 average)
- use this growth rate to project total energy demand from historical actuals
- multiply by an expected-peak-to-expected-energy ratio, which may change randomly over time
- allow for year-to-year variation in both energy and peak.

Note that step changes are not incorporated into this HLFM, except in so far as they are included in MED’s forecast.

The resulting HLFM, then, is:

$$D_y = U_y * P2E_y * (1 + \text{white noise } (\sigma = 0.02))$$

$$E_y = U_y * (\text{hours per year} / 1000) * (1 + \text{white noise } (\sigma = 0.015))$$

where:

D_y is annual national peak demand (MW) in year y ,

E_y is annual total energy demand (GWh),

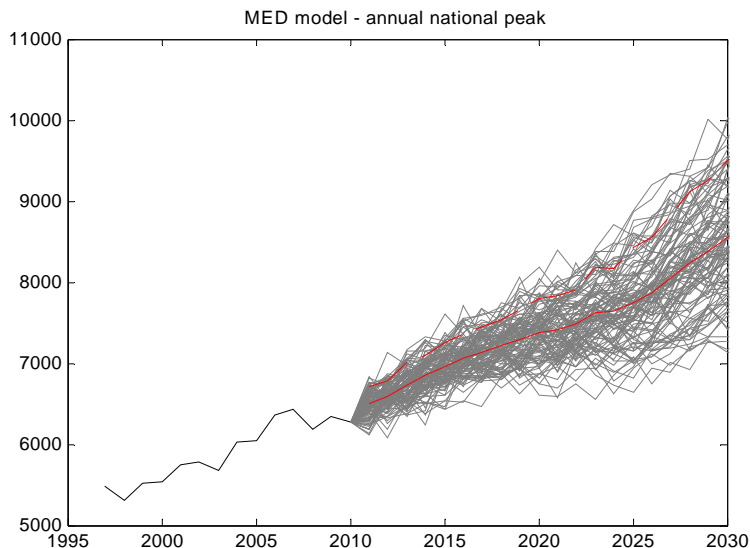
U_y = mean national energy demand over 2005-09 x MED growth rate,

$P2E_y = 1.47 + \max(0.2, (y - 2010) / 30) * \text{an } N(0,0.1) \text{ variate.}$

The white noise standard deviations are based on recent historical data, as is the base peak-to-mean-demand ratio of 1.47. The parameters relating to future variation in the peak-to-mean-demand ratio are purely speculative. **Reviewers agree it is important to model the possibility of the peak-to-mean-demand ratio changing.**

This plot (updated) shows a sample of the projections produced by the MED HLFM. The plotted quantity is D_y , annual national peak demand. The two red lines are the mean and P90 of the projections.

This HLFM predicts somewhat restrained growth over the next few years, relative to the other models.



A.9 The allocation methodology

The job of the AM (allocation methodology) is to convert national forecasts into all the outputs required:

- seasonal mean demand
- annual / seasonal, national / island / regional half-hourly peaks
- seasonal, night / day, regional half-hourly troughs.

Unavoidably this process is fiddly. However, it plays a key role in the forecast – bridging the gap between the quantities that the HLFMs produce and the quantities that the user needs.

The AM is run once for each of the N iterations of the main loop (see Section 2.2). It is based on randomizations of annual national energy, annual national peak, GDP, population, and step changes. Further randomness is incorporated at various stages of the AM. The cumulative effect is to model the reasonable range of uncertainty.

The AM algorithm is not set out in full here as it is lengthy, but the high level steps are described below. More detail will be provided on request. Please note that the choice of parameters such as white noise standard deviations is judgement-based, rather than resulting from some kind of model fit.

The **first** step is to produce projections of national seasonal peaks. (The seasons used are defined in Section 2.1). The HLFMs provide annual peaks, which by assumption occur in winter; we need to produce shoulder and summer peaks. We use the ratio of winter peak to another season's peak from a randomly chosen recent year, except:

- where we know a major load has had an unusual seasonal pattern in recent years (e.g. Tiwai demand has changed over the seasons since the 2008 transformer failure; in future years, we would expect Tiwai load to be more or less flat)
- where a future step change is expected to be seasonal (for instance, irrigation, dairy processing, heat pump cooling load)
- an allowance is made for the summer/winter ratio to slowly drift over time, for some as-yet-unknown reason.¹⁵

The **second** step is to break down national projections into the contributions of individual components¹⁶ – where a component might be *industrial* (e.g. Tiwai), *embedded* (e.g. White Hill wind), *new* (e.g. Solid Energy briquette plant), or *residual* (e.g. total Otago/Southland demand minus the three components listed above). The approach depends on the type of component:

- for most industrial and embedded components, the contribution to national peak in each future year is set to the value observed in some randomly chosen recent year¹⁷
- there are some exceptions, where we consider that recent history is not representative of the plant's potential contribution to national peak – in these cases, a different value has been supplied. An example is Norske Skog Tasman's plant at Kawerau.

¹⁵ In recent years, there has been no clear trend of national summer demand increasing relative to winter, or vice versa. Still, this may change in future. The variability in the summer/winter ratio is assumed to start at nil (in 2011) and increase to a standard deviation of 5% (in 2040). The shoulder/winter ratio is assumed to vary just half as much.

¹⁶ As opposed to the contributions of individual components to *regional* or *island* peaks.

¹⁷ A little additional variation is added, since the future contribution to national peak could be outside the historical range – this is multiplicative noise with standard deviation 3%.

- for new components, the contribution to national peak is supplied (see Appendix A.2). Some random variation is added.
- for residual components, the contribution to national peak in each future year is first set to the value in some randomly chosen recent year, and then a growth factor is applied.

The growth factor is based on a combination of predictors:¹⁸

- recent peak growth in the region (if contribution to national peak in 2008, 2009 and 2010 was substantially greater than in 2003, 2004 and 2005 then this suggests growth may continue to be fast)
- the regional population projection (relative to the average over 2006-2010)
- the regional GDP projection (relative to the average over 2006-2010)
- a constant growth scenario – indicating the possibility that any region can have at least moderate growth, even if all the above factors seem to indicate otherwise.

These predictors are combined to produce a growth factor for each region. For the start of the model horizon, each of the four predictors is given equal weight; but the weight on the ‘recent peak growth’ predictor reduces linearly to 0 by 2022, and the weight on the ‘constant growth scenario’ predictor reduces linearly to 0 by 2032. In other words, in the long term, population and GDP projections are expected to drive organic growth.

Long-term random drift and short-term white noise are applied to the growth factors. The short-term white noise is multiplicative, with a standard deviation increasing from 4% in 2012 to 20% in 2040. The long-term drift is also multiplicative, starting at 1 and changes linearly over the model horizon, converging on a value with a standard deviation of 10% in 2040. This process is applied for peak demand in each season, and also for annual energy demand (GWh).

An example of the results is shown in the plot on page 36.

The **third** step is to move from each component’s contribution to national peak to its contribution to island and regional peak. Again, the approach depends on the type of component:

- for most industrial and embedded components, the contribution to island and regional peak in each future year is set to the value in some randomly chosen recent year¹⁹
- there are some exceptions, such as Norske Skog Tasman’s plant at Kawerau and Carter Holt Harvey’s plant at Kinleith, where we consider that recent history is not representative of the plant’s potential contribution to regional peak – in these cases, a different value has been supplied
- for new components, the contribution to island and regional peak is supplied (see Appendix A.2)
- for residual components, the contribution to island peak is set to the contribution to national peak multiplied by a scaling factor.²⁰ The scaling

¹⁸ The use of multiple predictors of regional growth may seem unnecessarily complicated. However it fulfils an important function, by making the forecast robust to changes in the value of any one predictor. For instance, altering the GDP forecast for a region does not drastically alter the forecast of regional peak. See Appendix C for sensitivities that demonstrate robustness.

¹⁹ A little additional variation is then added, since the future contribution to peak could fall outside the historical range – this is multiplicative noise with standard deviation 4%.

factor is equal to the ratio of (contribution to regional peak) to (contribution to national peak), in the relevant season of some recent year. If this ratio is extreme then it is constrained to fall within reasonable bounds. The same approach is used for regional peak.

This process is applied for peak demand in each season. Component contributions can then be added up to produce regional peaks. Of course, this results in island peaks whose sum is greater than the national peak, and regional peaks whose sum is greater again.

Seasonal peaks can then be converted into annual peaks (trivially – the annual peak for a given region is equal to the maximum of the seasonal peaks in a given year).

The last step is to predict regional troughs (by season, for both night and day). This step is of secondary importance and an abbreviated methodology is used. Demand is broken down into components and the contribution of each component to regional trough is projected, based on recent history and supplied information.

We have spent some time checking the results of the AM for reasonableness. This has included comparing them with previous years and with recent historical values.


Generally our conclusion is that the AM produces a reasonable range of uncertainty (for instance, the combination of multiple sources of white noise does not produce unrealistically high variability at the regional level).

²⁰ Again, a little additional variation is then added – multiplicative noise with standard deviation 1% (island) or 2% (region).

For interest, a complete list of industrial and embedded components is provided below. In addition, for each region, there is a new component (covering step changes) and a residual component (covering everything else).

Component	Region	Type	Notes
Glenbrook (NZ Steel)	Auckland	Industrial	Recent peak contribution may not be representative of possible future values – projections are supplied (revised since draft report)
Pacific Steel	Auckland	Industrial	
Kawerau (Norske Skog)	Bay of Plenty	Industrial	Recent peak contribution may not be representative of possible future values – projections are supplied (revised since draft report)
Kinleith (CHH)	Bay of Plenty	Industrial	
Karioi (Winstone PI)	Central Districts	Industrial	
Whirinaki (PanPac)	Hawkes Bay	Industrial	
Marsden Point (NZ Refining)	North Isthmus	Industrial	
Tiwai (NZAS)	Otago/Southland	Industrial	Treated as a special case
South Canterbury recent dairy/irrigation growth	South Canterbury	'Industrial'	This has been split into a separate component to give a more accurate picture of other load growth in the region
Kupe	Taranaki	Industrial	Projections are supplied
Motunui (Methanex)	Taranaki	Industrial	Projections are supplied
Te Rapa	Waikato	Industrial	
Oceana Gold	West Coast	Industrial	Projections are supplied
Pike River	West Coast	Industrial	Not yet clear what ongoing mining activity will take place at Pike River
Highbank	Canterbury	Embedded	
Rotokawa	Central Districts	Embedded	
Tararua Wind 1 & 2	Central Districts	Embedded	Not TWF 3, which is grid connected
White Hill	Otago/Southland	Embedded	Projections are supplied
Mangahewa	Taranaki	Embedded	Projections are supplied

Figure 14: How the allocation methodology breaks down national peak demand into individual components: one example



Note: This graph shows just one possible trajectory of future demand – the full forecast is made up of many randomizations like this, each of which will have a different pattern.

A.10 Planned improvements

When time permits, we intend to:

- further review the treatment of energy-to-peak ratios in all four HLFMs (as suggested by both reviewers and Brian Kirtlan of the Electricity Authority)
- document the allocation methodology in more detail (as urged by all reviewers)
- replace the judgement-based HLFM with an application of the Delphi method (Hyndman, bullet point 14)
- refit the endogenous HLFM, carefully considering how to treat the anomalous 2009, 2010 and 2011 peaks
- review temperature correction of demand and consider including this process in the endogenous model, if it is deemed to produce sensible results and add value (Hyndman bp 13)
- refit the econometric HLFM, modelling per-capita energy demand, and using per-capita GDP as a covariate (Hyndman bp 9)
- consider using cooling/heating degree days as covariates in the econometric HLFM (Hyndman bp 10)
- review the various constants in the HLFMs and the allocation methodology, seeking to estimate parameters rather than arbitrarily picking them whenever possible (Hyndman bps 16, 19)
- do more extensive backcasting (Hyndman bp 18)
- institute an annual validation report (Appendix D)
- extend the forecast to support variable line ratings.


Appendix B Detailed results

B.1 National energy

Our main interest is in peak; however, an energy forecast has been prepared and is provided here for completeness.

The national forecast is shown in Figure 15, expressed in terms of mean demand (MW) rather than energy (GWh) (the two are identical down to a scaling factor, though note this scaling factor differs slightly between normal and leap years).

Figure 15: Annual national mean demand forecast (updated)



Regional energy forecasts have also been prepared, and will be provided on request.

B.2 Regional peaks

The following plots show P90, prudent and expected annual peaks for every region.

Recall the prudent peak is P90 for the first five years, and grows at the same rate as the expected forecast thereafter.

Except where noted, the change in forecasts since the May draft is relatively minor.

Figure 16: Annual regional peak forecast – North Isthmus (updated)




Figure 17: Annual regional peak forecast – Auckland (updated)




Figure 18: Annual regional peak forecast – Waikato (updated)




Figure 19: Annual regional peak forecast – Bay of Plenty (updated)




Figure 20: Annual regional peak forecast – Central Districts (updated)




Figure 21: Annual regional peak forecast – Hawkes Bay (updated)




Figure 22: Annual regional peak forecast – Taranaki (updated)




Figure 23: Annual regional peak forecast – Wellington (updated)




Figure 24: Annual regional peak forecast – Nelson/Marlborough (updated)






Figure 25: Annual regional peak forecast – West Coast (updated)



The new forecast is higher than the May draft, due to possible new mining loads.

Figure 26: Annual regional peak forecast – Canterbury (updated)



The new forecast is higher than the May draft; the allowance for earthquake effects has been reduced, and there is possible new dairy processing load.

Figure 27: Annual regional peak forecast – South Canterbury (updated)





Figure 28: Annual regional peak forecast – Otago/Southland (updated)



The following plots show seasonal forecasts for Canterbury and South Canterbury – two regions where strong growth in summer-based agricultural load is expected.

Figure 29: Seasonal prudent forecasts – Canterbury (updated)





Figure 30: Seasonal prudent forecasts – South Canterbury (updated)



B.3 Regional troughs

For each region, trough forecasts have been prepared (separately for day and night in each season).

These forecasts are of limited general interest, but some examples are shown below for completeness.

Figure 31: Expected trough forecasts – West Coast




Figure 32: Expected trough forecasts – Taranaki



Appendix C Sensitivities and scenarios

C.1 Sensitivity to recent changes in demand

It is intended that the forecast should not be unduly sensitive to changes in any one input. This section demonstrates that:


- if recent national peak demand had been higher, the national peak forecast would not have been drastically higher
- if recent regional peak demand had been higher in, say, Auckland, the regional peak forecast would not have been drastically higher.

Two sensitivity runs have been carried out:

- “high recent national peaks” – inflating national peak demand by 100 MW in 2008, 200 MW in 2009, and 300 MW in 2010
- “high recent Auckland peaks” – inflating Auckland regional peaks, and Auckland contribution to national peaks, by 30 MW in 2008, 60 MW in 2009, and 90 MW in 2010.²¹

The results are shown below. All forecasts shown are prudent (P90 for the first five years, growing at the same rate as the expected forecast thereafter).

Figure 33: National peak forecasts – “high recent national peaks” vs base case




The forecast is higher, because:

- higher recent demand increases the projections of the endogenous HLFM
- although we cannot know what the ad hoc HLFM would have predicted on the basis of these data, we assume it would have been somewhat higher.

²¹ Other regions’ contributions to national peak are derated by the same amount, so as to leave total national peaks unchanged.

However the difference between base case and sensitivity is moderate, because the econometric and MED HLFMs are not affected. If these HLFMs were not used, the gap would be much wider.

Figure 34: Auckland regional forecasts – “high recent Auckland peaks” vs base case



The forecast is higher, because the AM takes recent growth into account when projecting regional peaks.

However, the difference between base case and sensitivity is moderate. If the AM was purely driven by the trend over the last few years, the gap would be much wider.

Our conclusion is that the sensitivity of the forecast to recent peaks is reasonable.

C.2 Sensitivity to GDP forecasts

It is intended that the forecast should not be unduly sensitive to changes in any one input. This section demonstrates that:

- if national GDP forecasts had been higher, the national peak forecast would not have been drastically higher
- if the regional GDP forecast for, say, Auckland, had been higher, the regional peak forecast would not have been drastically higher.

Two sensitivity runs have been carried out:

- “high national GDP” – NZIER projection is for NZ to quickly emerge from the recession and grow at 2.5% p.a. thereafter – moderately high population growth is also assumed
- “high Auckland GDP” – NZIER projection of Auckland GDP growth increases by 1 percentage point from 2009 onwards.²²

The results are shown below. All forecasts shown are prudent (P90 for the first five years, growing at the same rate as the expected forecast thereafter).

Figure 35: National peak forecasts – “high national GDP” vs base case




The forecast is higher, because the econometric HLFM is very sensitive to GDP projections.

However the difference between base case and sensitivity is moderate, because the econometric and MED²³ HLFMs are not affected, and we assume that the ad hoc HLFM would also not have changed. If only the econometric model was used, the gap would be much wider.

²² Other regions’ GDP forecasts are derated by the same amount, so as to leave the national GDP forecast unchanged.

²³ The MED-derived model does include GDP assumptions, but these are not based on the NZIER forecast – so the MED HLFM is not affected by changes to NZIER’s projections.

Figure 36: Auckland regional forecasts – “high Auckland GDP” vs base case



The forecast is higher, because the AM takes projected GDP growth into account when projecting regional peaks.

However, the difference between base case and sensitivity is moderate. If the AM was purely driven by GDP and population projections, the gap would be much wider.

Our conclusion is that the sensitivity of the forecast to GDP forecasts is reasonable.

C.3 A one-year backcast

It is useful to consider what results this methodology might have produced if it had been employed in previous years. It is hoped that there would not have been extensive “flip-flopping” from year to year.

It is always difficult, however, to determine what might have been done given the state of knowledge at the time.

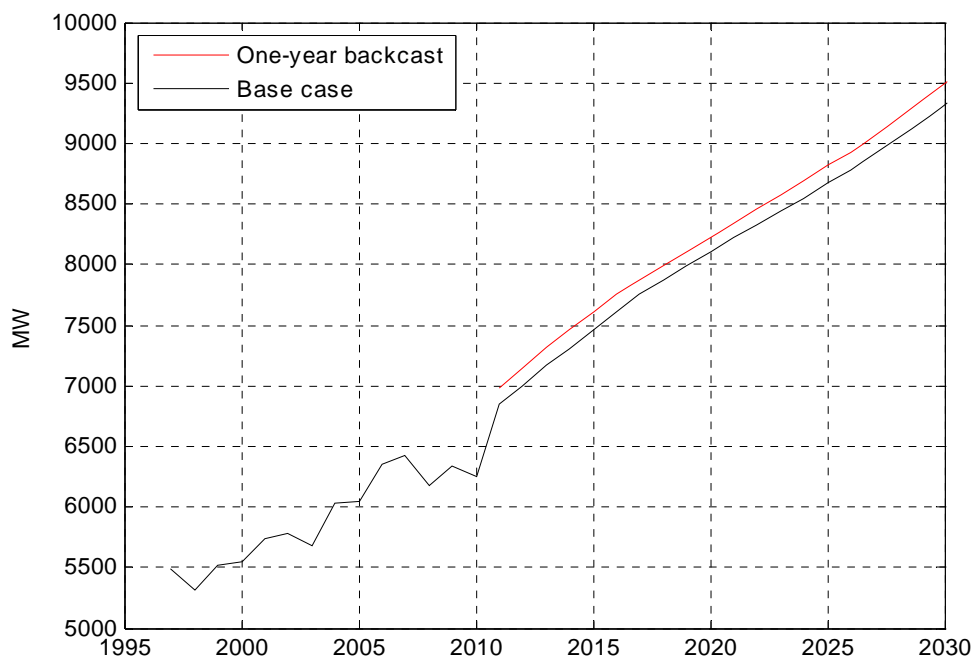
A one-year backcast has been produced, seeking to determine what results would have been obtained if this methodology had been applied a year ago. Relative to the current forecast, the changes are:

- the econometric HLFM is run using out-of-date GDP forecasts (though the model has not been refitted);
- the endogenous HLFM is refitted without 2010 data;
- the ad-hoc HLFM now uses underlying growth paths that might have seemed realistic at the time
- the AM does not use 2010 actuals in the allocation process.

No attempt has been made to reproduce a set of step changes that might have been put forward a year ago – the current set of step changes was used without modification.

The results are shown below. All forecasts shown are prudent (P90 for the first five years, growing at the same rate as the expected forecast thereafter).

Figure 37: National peak forecasts – one-year backcast vs base case



The new national forecast is lower than the backcast, because three out of four HLFMs have gained additional information that suggests that demand will be lower than previously thought. However, the difference is very moderate.

At the regional level, the differences are also typically moderate. In some cases, the new forecast is lower (Figure 38); in some cases, higher (Figure 39).

Figure 38: North Isthmus peak forecasts – one-year backcast vs base case





Figure 39: Wellington peak forecasts – one-year backcast vs base case



Generally these results are reassuring – the forecast can be expected to have a reasonable level of stability from year to year.

C.4 Heat pumps

In 2009, We commissioned a study by BRANZ Ltd on the potential effects of domestic heat pumps on electricity demand.²⁴

The study suggests that heat pump uptake may lead to substantial changes in peak demand in some areas, including:

- increases in summer peak (due to cooling load)
- increases in winter peak (due to heating load)
- in some areas at some times, a net decrease in winter peak (due to replacement of inefficient resistance heaters).

We have created two scenarios based on the BRANZ report.

- In the “high heat pump” scenario, BRANZ’s base-case cooling and heating loads are multiplied by 4/3.
- In the “low heat pump” scenario, BRANZ’s base-case cooling and heating loads are multiplied by 2/3.

In each scenario, a “summer cooling” effect is added to each region’s summer peak, and a “winter heating” effect is added to each region’s winter peak. The effects are calculated by taking the base-case summer cooling and winter heating loads from BRANZ’s Tables 1 and 2, and subtracting a “business as usual” load calculated as BRANZ’ figure for 2009 plus 1.5% growth per annum. (For the purpose of these scenarios, the gap between the tabulated figure and this “business as usual” calculation is considered to be an additional, unanticipated load.)

For some regions, the effect is substantial. For instance, in the “high heat pump” scenario, by 2015:

- Auckland and North Isthmus combined have 120 MW additional summer load, and 115 MW additional winter load
- the Upper South Island has 55 MW additional summer load, and 65 MW additional winter load.

These scenarios could be used as sensitivities for grid planning purposes.

C.5 Electric vehicles

Electric vehicle charging load is widely seen as a potentially significant influence on future demand.

Some key uncertainties include the amount of charging load and when it occurs (from a grid planning point of view, it would generally be preferable for vehicles to charge off-peak).

We have created four scenarios based on MED publications.

- In the “high uptake, flat load” scenario, we add a new load, spread evenly over seasons of the year and times of the day. The amount of load increases from near-nil in 2015, to 100 GWh in 2020, 450 GWh in 2025, 1 TWh in 2030, and 2 TWh in 2040.
- In the “high uptake, trough-filling” scenario, we add the same amount of load, but assume it is largely off-peak.

²⁴ <http://www.transpower.co.nz/n3554.html>

- In the “low uptake, flat load” scenario, the amount of load increases from near-nil in 2015, to 75 GWh in 2020, 200 GWh in 2025, 350 GWh in 2030, and 700 GWh in 2040.
- In the “low uptake, trough-filling” scenario, we add the same amount of load, but assume it is largely off-peak.

Electric vehicle charging load is assumed to be distributed across regions in proportion to population.

These scenarios have relatively little impact on national and regional peak demand between now and 2025, and are less likely to be used for grid planning purposes.

C.6 Tiwai decommissioning

The NZAS aluminium smelter at Tiwai is by far New Zealand’s biggest single load, and clearly if the smelter reduced its output there would be a significant impact on transmission flows. We have no reason to expect this to happen, but we have included an “early Tiwai decommissioning” scenario for completeness.

Figure 40: Annual regional peaks in three scenarios, c.f. the base case

Scenario	2015 prudent peak (MW)		2030 prudent peak (MW)	
	Auckland	Otago/Southland	Auckland	Otago/Southland
Base case	1,640	1,140	2,230	1,250
High heat pumps	1,710	1,180	2,280	1,250
Electric vehicles – high uptake, flat load	1,640	1,140	2,260	1,260
Tiwai decommissioning	1,640	1,140	2,230	610

Appendix D Validation

This section compares 2011 winter peak demand with the predictions made in the May draft forecast.

Validation is always worthwhile, but is particularly important at this point because the polar blast of August 2011 set unusually high peaks. If these peaks were substantially higher than the prudent forecasts published in May, it would suggest that the forecasting methodology was insufficiently prudent.

Region	Winter peak demand (MW)				Comment
	2010 actual	2011 actual (to end of August)	Draft forecast for 2011 - expected	Draft forecast for 2011 - prudent	
New Zealand	6,254	6,557	6,642	6,848	Below expected
North Island	4,189	4,470	4,560	4,740	"
South Island	2,079	2,132	2,147	2,206	"
North Isthmus	788	880	856	897	Between expected & prudent
Auckland	1,255	1,368	1,415	1,501	Below expected
Waikato	463	470	472	496	"
Bay of Plenty	495	480	535	559	"
Central Districts	267	289	312	330	"
Hawkes Bay	302	295	302	315	"
Taranaki	188	183	204	212	"
Wellington	679	729	701	750	Between expected & prudent
West Coast	60	56	57	63	Below expected
Nelson/ Marlborough	210	217	224	238	"
Canterbury	731	775	709	747	Above prudent
South Canterbury	158	144	160	179	Below expected
Otago/ Southland	1,024	1,061	1,078	1,102	"

It turns out that, nationally and for most regions, the actual peak for winter 2011 was within the draft forecast bounds. This suggests the methodology is appropriately prudent (in this regard at least).²⁵

The exception was the Canterbury regional peak, which was higher than the draft prudent forecast. It appears that, despite the recent earthquakes, Christchurch can still experience high demand peaks in very adverse weather conditions. In this final report, we have responded by halving the assumed reduction in demand due to earthquake effects.

If time permits, we will institute an annual validation report that compares predictions from previous years with actual demand. Validation is an important part of the forecasting cycle and can be used to drive improvements in the methodology.

It should be added that validation is most effective when the methodology remains relatively stable from year to year.

²⁵ Some might argue that the methodology was excessively prudent, and that an extreme weather event should lead to demand peaks close to (or even above) the prudent forecast. We disagree – our view is that national peak demand in winter 2011 could easily have been substantially higher than it actually was, if, for example, the economy had recovered faster than expected.

Appendix E Technical review by NZIER

Appendix F Technical review by Professor Hyndman