



MONASH
University

Forecasting long-term peak half-hourly electricity demand for New South Wales

Dr Shu Fan

B.S., M.S., Ph.D.

Professor Rob J Hyndman

B.Sc. (Hons), Ph.D., A.Stat.

**Business & Economic
Forecasting Unit**

Telephone: (03) 9905 2358

Fax: (03) 9905 5474

ABN: 12 377 614 012

Report for

The Australian Energy Market Operator (AEMO)

3 June 2015

Contents

Summary 3

1 Modelling and forecasting electricity demand of summer 5

1.1 Historical data 5

1.2 Variable selection for the half-hourly model 10

1.3 Model predictive capacity 13

1.4 Half-hourly model residuals 16

1.5 Modelling and simulation of PV generation 17

1.6 Probability distribution reproduction 21

1.7 Demand forecasting 22

2 Modelling and forecasting electricity demand of winter 28

2.1 Historical data 28

2.2 Variable selection for the half-hourly model 32

2.3 Model predictive capacity 33

2.4 Half-hourly model residuals 36

2.5 Modelling and simulation of PV generation 37

2.6 Probability distribution reproduction 41

2.7 Demand forecasting 42

References 49

Summary

The Australian Energy Market Operator (AEMO) prepares seasonal peak electricity demand and energy consumption forecasts for the New South Wales (NSW) region of the National Electricity Market (NEM). Peak electricity demand and energy consumption in New South Wales are subject to a range of uncertainties. The electricity demand experienced each year will vary widely depending upon prevailing weather conditions (and the timing of those conditions) as well as the general randomness inherent in individual usage. This variability can be seen in the difference between the expected peak demand at the 1-in-10-year probability of exceedance (or 10% PoE) level and the peak demand level we expect to be exceeded 9 in 10 years (or 90% PoE). The electricity demand forecasts are subject to further uncertainty over time depending upon a range of factors including economic activity, population growth and changing customer behavior. Over the next decade we expect customer behaviour to change in response to changing climate and electricity prices, to changes in technology and measures to reduce carbon intensity.

Our model uses various drivers including recent temperatures at two locations (Sydney and Richmond), calendar effects and some demographic and economic variables. The report uses a semi-parametric additive model to estimate the relationship between demand and the driver variables. The forecast distributions are obtained using a mixture of seasonal bootstrapping with variable length and random number generation.

This report is part of Monash University's ongoing work with AEMO to develop better forecasting techniques. As such it should be read as part of a series of reports on modelling and forecasting half-hourly electricity demand. In particular, we focus on illustrating the historical data and forecasting results and discussing the implications of the results in this report. The underlying model and related methodology are explained in our technical report (Hyndman and Fan, 2015).

The focus of the forecasts presented in this report is operational demand, which is the demand met by local scheduled generating units, semi-scheduled generating units, and non-scheduled intermittent generating units of aggregate capacity larger than 30 MW, and by generation imports to the region. The operational demand excludes the demand met by non-scheduled non-intermittent generating units, non-scheduled intermittent generating units of aggregate capacity smaller than 30 MW, exempt generation (e.g. rooftop solar, gas tri-generation, very small wind farms, etc), and demand of local scheduled loads.

In this report, we estimate the demand models with all information up to February 2015 for the summer and winter season respectively, and we produce seasonal peak electricity demand forecasts for New South Wales for the next 20 years. In NSW, the annual peak tended to occur in winter before

2009, while the summer peak has tended to dominate since 2009. The forecasting capacity of the demand model for each season has been verified by reproducing the historical probability distribution of demand.

1 Modelling and forecasting electricity demand of summer

1.1 Historical data

1.1.1 Demand data

AEMO provided half-hourly New South Wales electricity demand values from 2002 to 2015. Each day is divided into 48 periods which correspond with NEM settlement periods. Period 1 is midnight to 0:30am Eastern Standard Time.

We define the period October–March as “summer” for the purposes of this report. Only data from October–March were retained for summer analysis and modelling. All data from April–September for each year were omitted. Thus each “year” consists of 182 days.

Time plots of the half-hourly demand data are plotted in Figures 1–3. These clearly show the intra-day pattern (of length 48) and the weekly seasonality (of length $48 \times 7 = 336$); the annual seasonality (of length $48 \times 182 = 8736$) is less obvious.

AEMO also provided half-hourly aggregated major industrial demand data for the summer season which is plotted in Figure 4. Although this load can vary considerably over time, it is not temperature sensitive, thus we do not include this load in the modelling.




Figure 1: Half-hourly demand data for New South Wales from 2002 to 2015. Only data from October–March are shown.

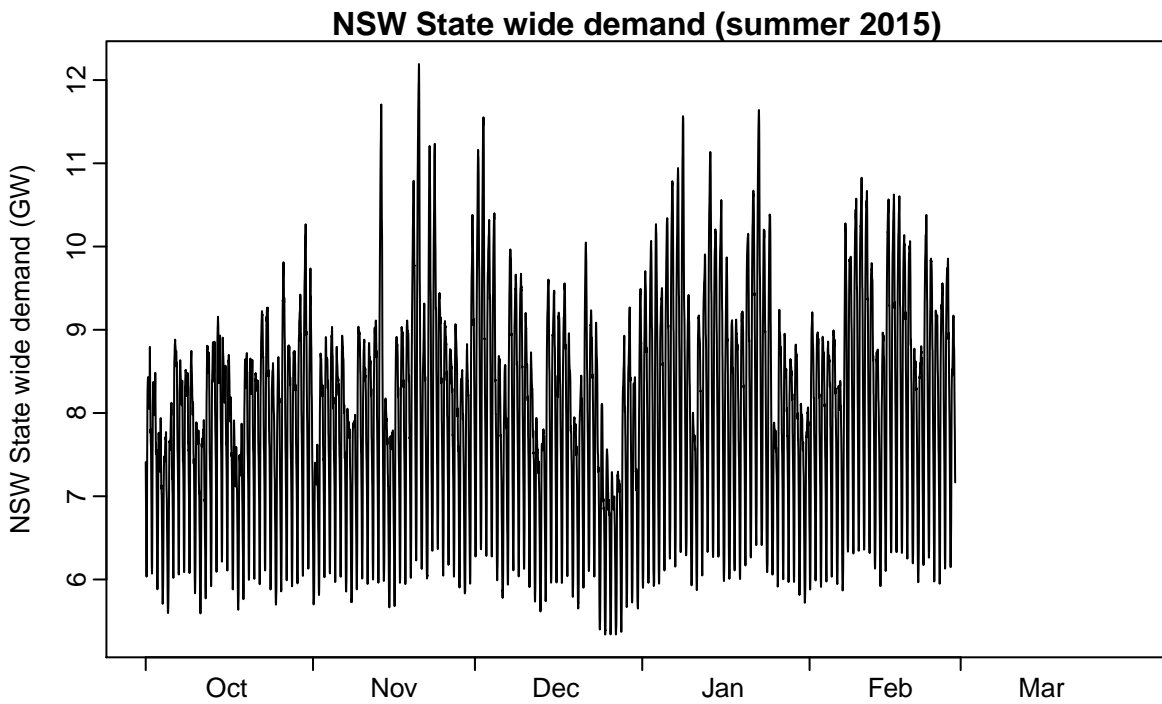


Figure 2: Half-hourly demand data for New South Wales for last summer.

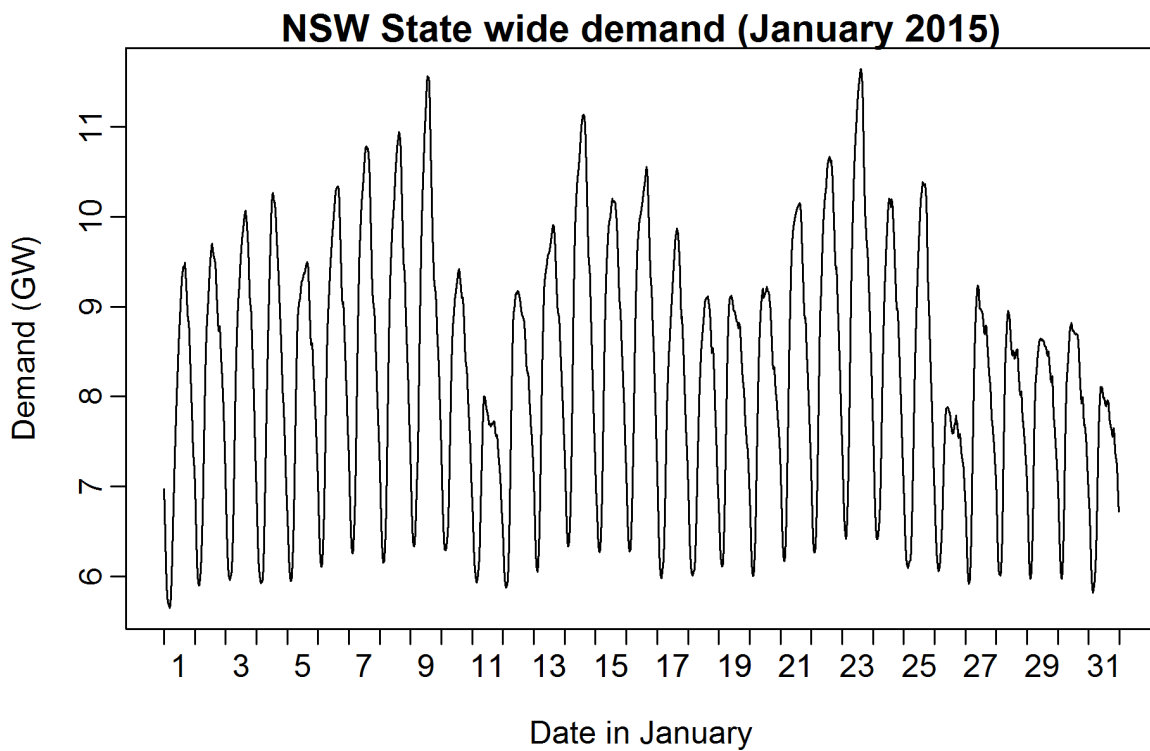


Figure 3: Half-hourly demand data for New South Wales, January 2015.

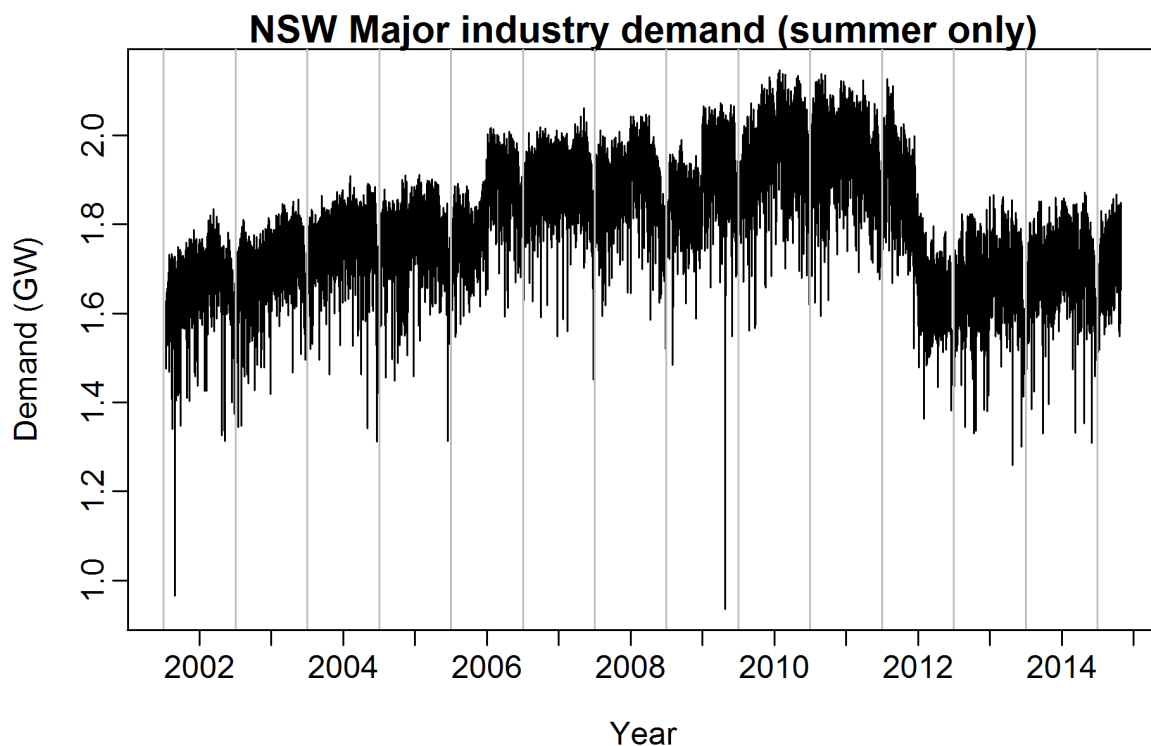


Figure 4: *Half-hourly demand data for major industries. 2002–2015.*

AEMO provided half-hourly rooftop generation data based on a 1MW solar system in New South Wales from 2003 to 2011. This data is shown in Figure 5, and the half-hourly values of total rooftop generation of New South Wales from 2003 to 2015 is shown in Figures 6, obtained by integrating the data from the 1MW solar system and the installed capacity of rooftop generation.




Figure 5: Half-hourly rooftop generation data based on a 1MW solar system. 2003–2011.

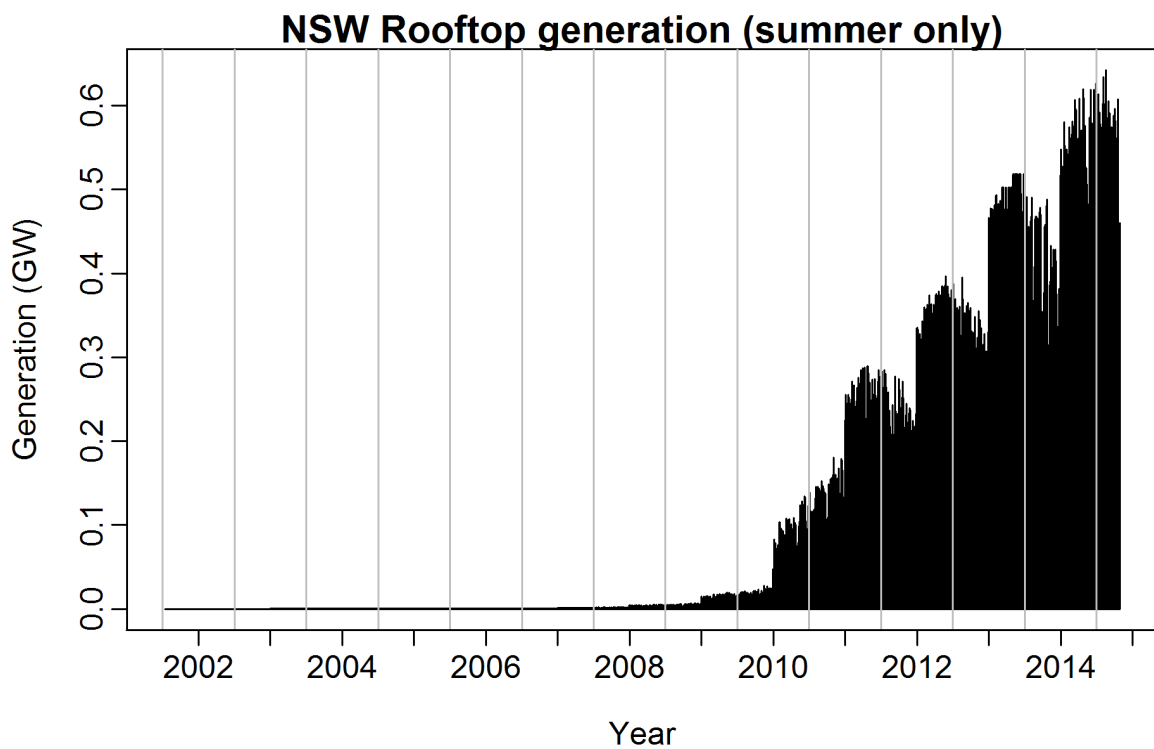


Figure 6: Half-hourly rooftop generation data. 2003–2015.

1.1.2 Temperature data and degree days

AEMO provided half-hourly temperature data for Sydney and Richmond, from 2002 to 2015. The relationship between demand (excluding major industrial loads) and the average temperature is shown in Figure 7.




Figure 7: Half-hourly NSW electricity demand (excluding major industrial demand) plotted against average temperature (degrees Celsius).

The half-hourly temperatures are used to calculate cooling degree days in summer and heating degree days in winter, which will be considered in the seasonal demand model.

For each day, the cooling degrees is defined as the difference between the mean temperature (calculated by taking the average of the daily maximum and minimum values of Sydney and Richmond average temperature) and 18°C. If this difference is negative, the cooling degrees is set to zero. These values are added up for each summer to give the cooling degree-days for the summer, that is,

$$CD = \sum_{\text{summer}} \max(0, t_{\text{mean}} - 18^{\circ}).$$

Accordingly, the heating degrees is defined as the difference between 19°C and the mean temperature for each day. If this difference is negative, the heating degrees is set to zero. These values are added

up for each winter to give the heating degree-days for the winter,

$$HD = \sum_{\text{winter}} \max(0, 19^\circ - t_{\text{mean}}).$$

The historical cooling and heating degree days for summer and winter are calculated in Figure 8.

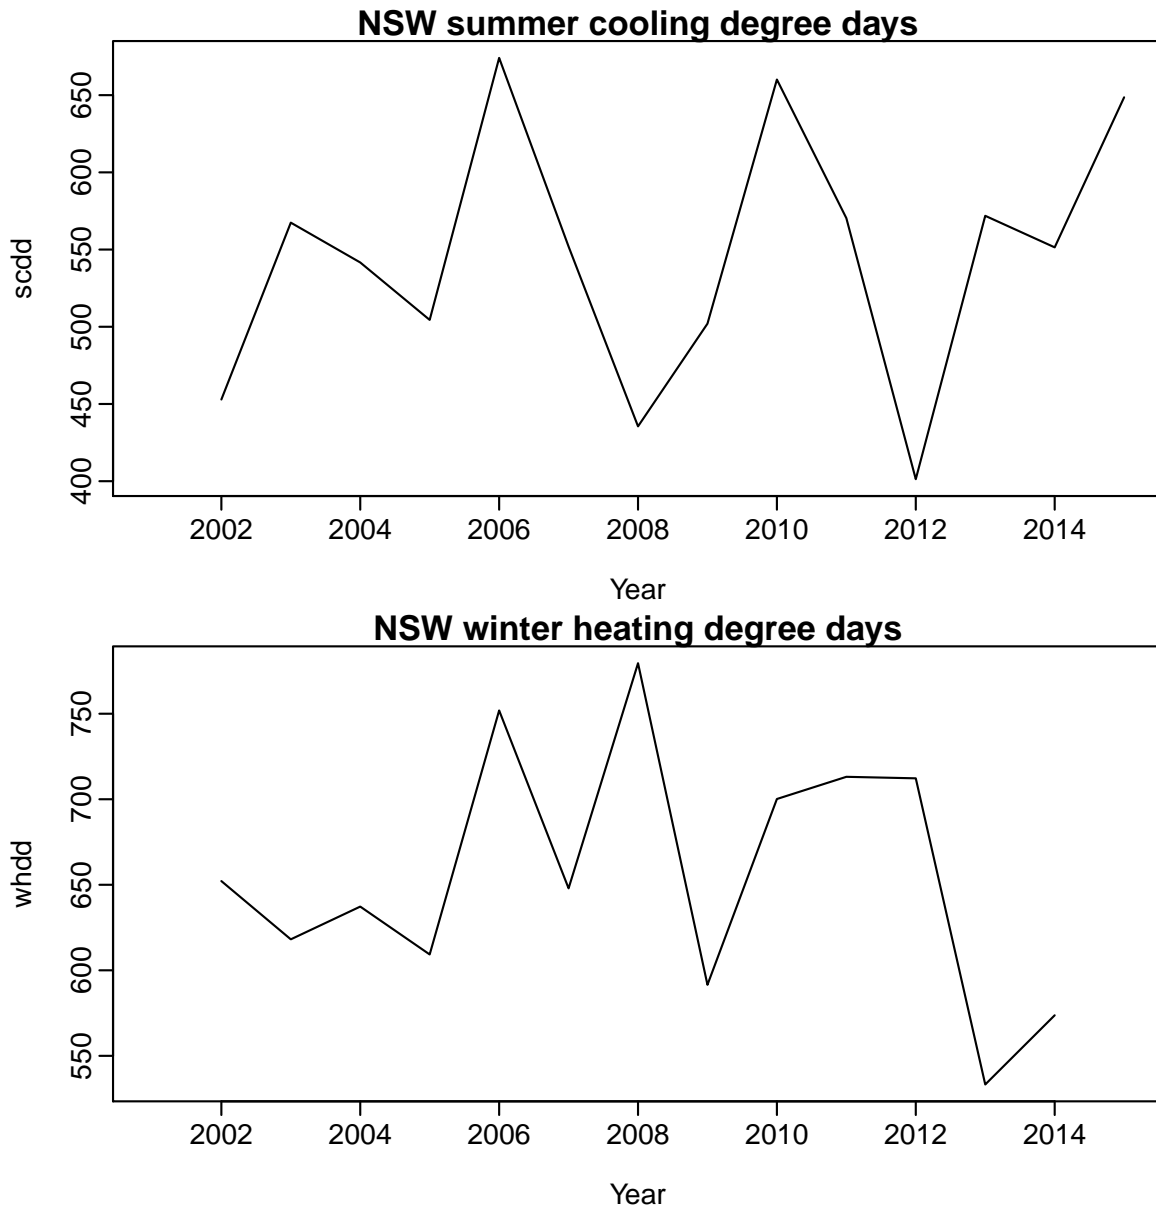


Figure 8: Cooling (above) and heating (bottom) degree days for summer and winter respectively.

1.2 Variable selection for the half-hourly model

To fit the model to the demand excluding major industrial loads, we normalize the half-hourly demand against seasonal average demand of each year. The top panel of Figure 9 shows the original

demand data with the average seasonal demand values shown in red, and the bottom panel shows the half-hourly adjusted demand data.

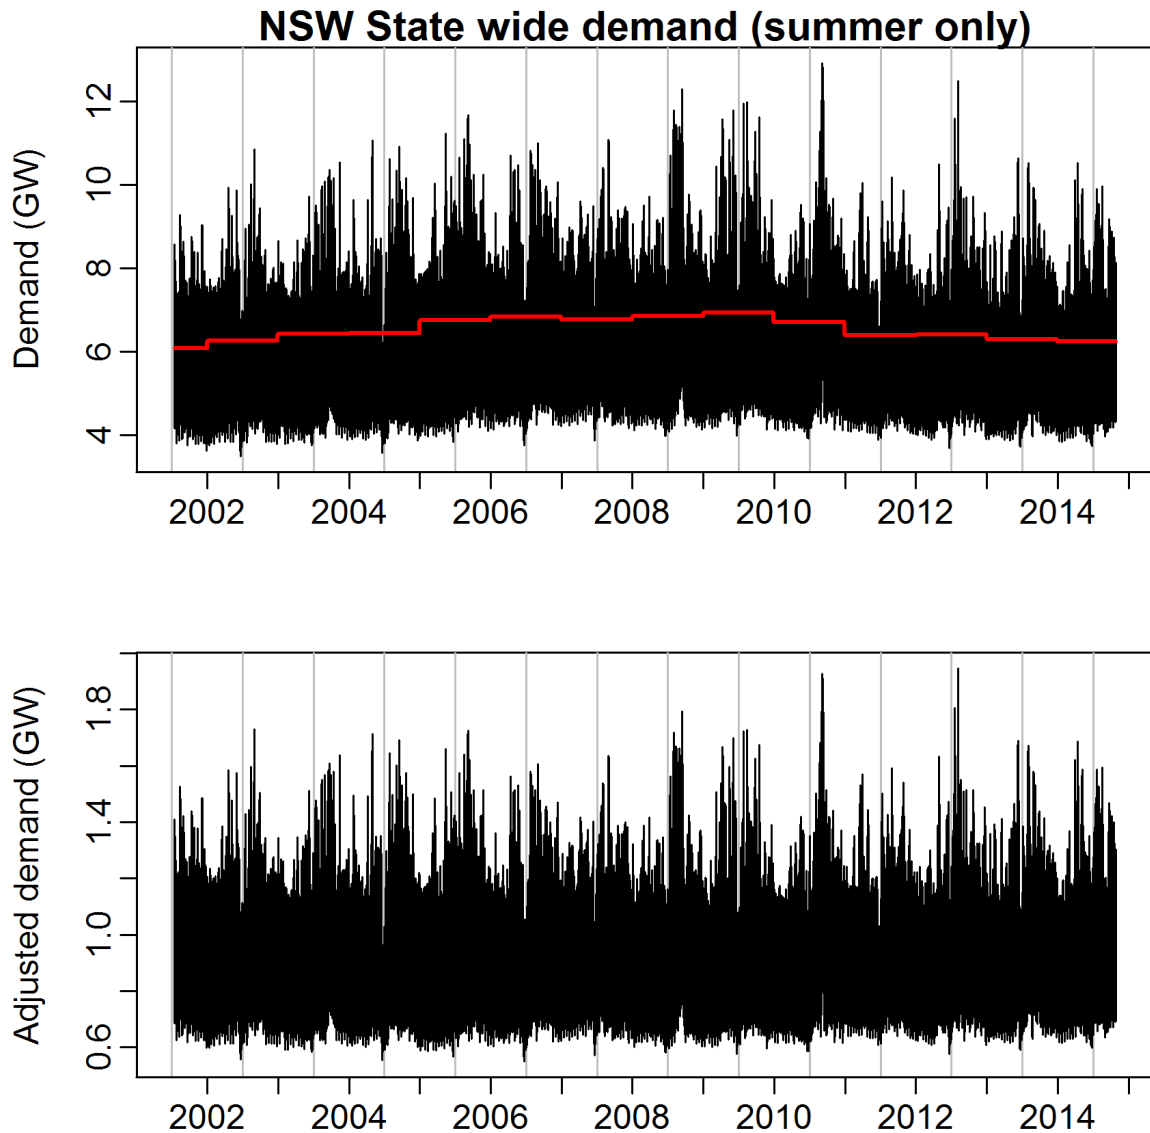


Figure 9: *Top: Half-hourly demand data for New South Wales from 2002 to 2015. Bottom: Adjusted half-hourly demand where each year of demand is normalized against seasonal average demand. Only data from October–March are shown.*

We fit a separate model to the adjusted demand data from each half-hourly period. In addition, we allow temperature and day-of-week interactions by modelling the demand for work days and non-work days (including weekends and holidays) separately. Specific features of the models we consider are summarized below.

- ▶ We model the logarithm of adjusted demand rather than raw demand values. Better forecasts were obtained this way and the model is easy to interpret as the temperature and calendar variable have a multiplicative effect on demand.
- ▶ Temperature effects are modelled using additive regression splines.

The following temperatures & calendar variables are considered in the model,

- ▶ Temperatures from the last three hours and the same period from the last six days;
- ▶ The maximum temperature in the last 24 hours, the minimum temperature in the last 24 hours, and the average temperature in the last seven days;
- ▶ Calendar effects include the day of the week, public holidays, days just before or just after public holidays, and the time of the year.

Four boosting steps are adopted to improve the model fitting performance (Hyndman and Fan, 2015), the linear model in boosting stage 1 contained the following variables:

- ▶ the current temperature and temperatures from the last 3 hours;
- ▶ temperatures from the same time period for the last 4 days;
- ▶ the current temperature differential and temperature differential from the last 3 hours;
- ▶ temperature differentials from the same period of previous 6 days;
- ▶ the average temperature in the last seven days;
- ▶ the maximum temperature in the last 24 hours;
- ▶ the minimum temperature in the last 24 hours;
- ▶ the day of the week;
- ▶ the holiday effect;
- ▶ the day of season effect.

The variables used in the seasonal demand model include per-capita state income, total electricity price and cooling and heating degree days, which is consistent with the energy model by AEMO. As explained in our technical report (Hyndman and Fan, 2015), the price elasticity varies with demand. For NSW, the price elasticity coefficient estimated at the 50% demand quantile was -0.17 , while the price elasticity coefficient estimated at the 95% demand quantile was -0.2 , indicating price elasticity of peak demand is larger than that applying to the average demand. In this report, therefore, we adjust the price coefficient in the seasonal demand model when forecasting the peak electricity demand. Specifically, we adjust the price coefficient as

$$c_p^* = c_p \times 0.2/0.17,$$

where c_p price coefficient is the original coefficient. Then we re-estimate the remaining coefficients using a least squares approach.

1.3 Model predictive capacity

We investigate the predictive capacity of the model by looking at the fitted values. Figure 10 shows the actual historical demand (top) and the fitted (or predicted) demands. Figure 11 illustrates the model prediction for January 2015. It can be seen that the fitted values follow the actual demand remarkably well, indicating that the vast majority of the variation in the data has been accounted for through the driver variables. Both fitted and actual values shown here are after the major industrial load has been subtracted from the data.

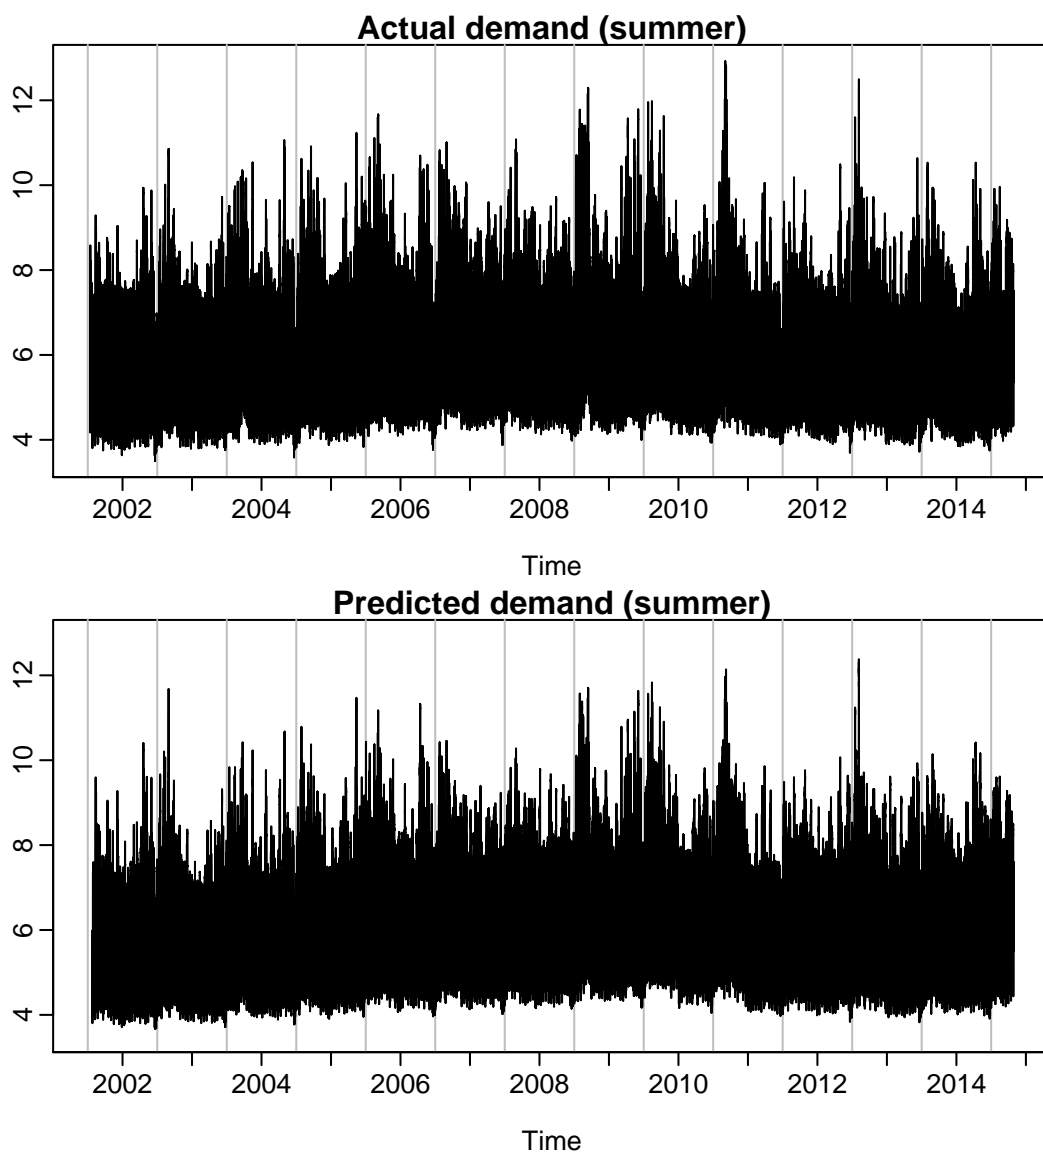


Figure 10: Time plots of actual and predicted demand.




Figure 11: Actual and predicted demand for January 2015.

Note that these “predicted” values are not true forecasts as the demand values from these periods were used in constructing the statistical model. Consequently, they tend to be more accurate than what is possible using true forecasts.

1.4 Half-hourly model residuals

The time plot of the half-hourly residuals from the demand model is shown in Figure 12.




Figure 12: *Half-hourly residuals (actual – predicted) from the demand model.*

Next we plot the half-hourly residuals against the predicted adjusted log demand in Figure 13. That is, we plot e_t against $\log(y_{t,p}^*)$ where these variables are defined in the demand model (Hyndman and Fan, 2015).

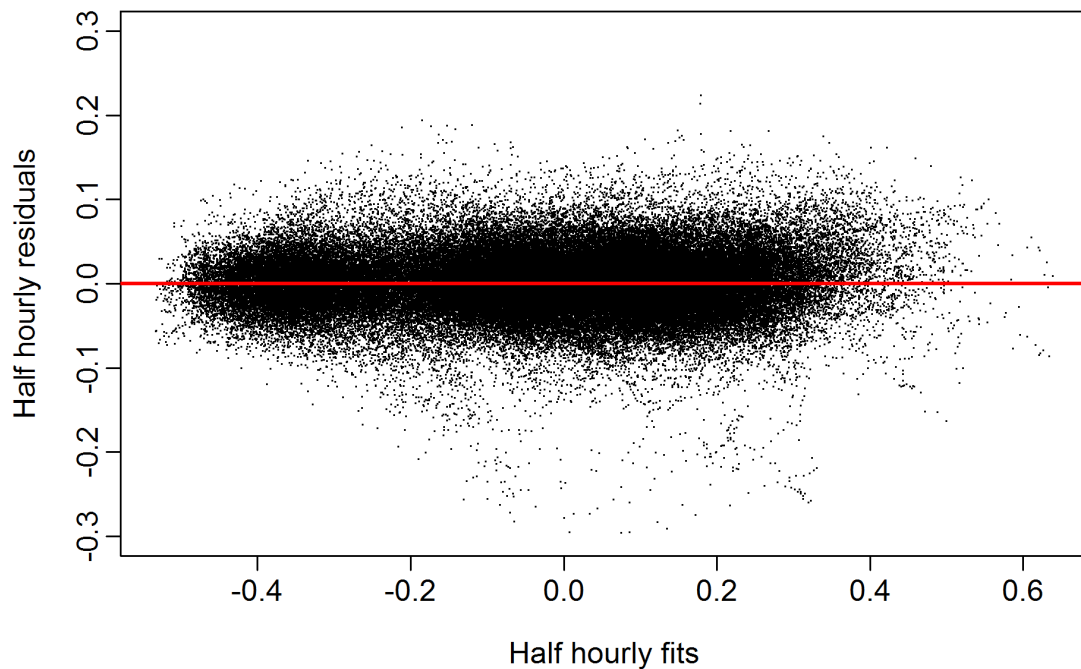


Figure 13: *Residuals vs predicted log adjusted demand from the demand model.*

1.5 Modelling and simulation of PV generation

Figure 14 shows the relationship between the daily PV generation and the daily solar exposure in NSW from 2003 to 2011, and the strong correlation between the two variables is evident. The daily PV generation against daily maximum temperature is plotted for the same period in Figure 15, and we can observe the positive correlation between the two variables. Accordingly, the daily PV exposure and maximum temperature are considered in the PV generation model (Hyndman and Fan, 2015).

We simulate future half-hourly PV generation in a manner that is consistent with both the available historical data and the future temperature simulations. To illustrate the simulated half-hourly PV generation, we plot the boxplot of simulated PV output based on a 1 MW solar system in Figure 17, while the boxplot of the historical ROAM data based on a 1 MW solar system is shown in Figure 16. Comparing the two figures, we can see that the main features of the two data sets are generally consistent. Some more extreme values are seen in the simulated data set — these arise because there are many more observations in the simulated data set, so the probability of extremes occurring somewhere in the data is much higher. However, the quantiles are very similar in both plots showing that the underlying distributions are similar.

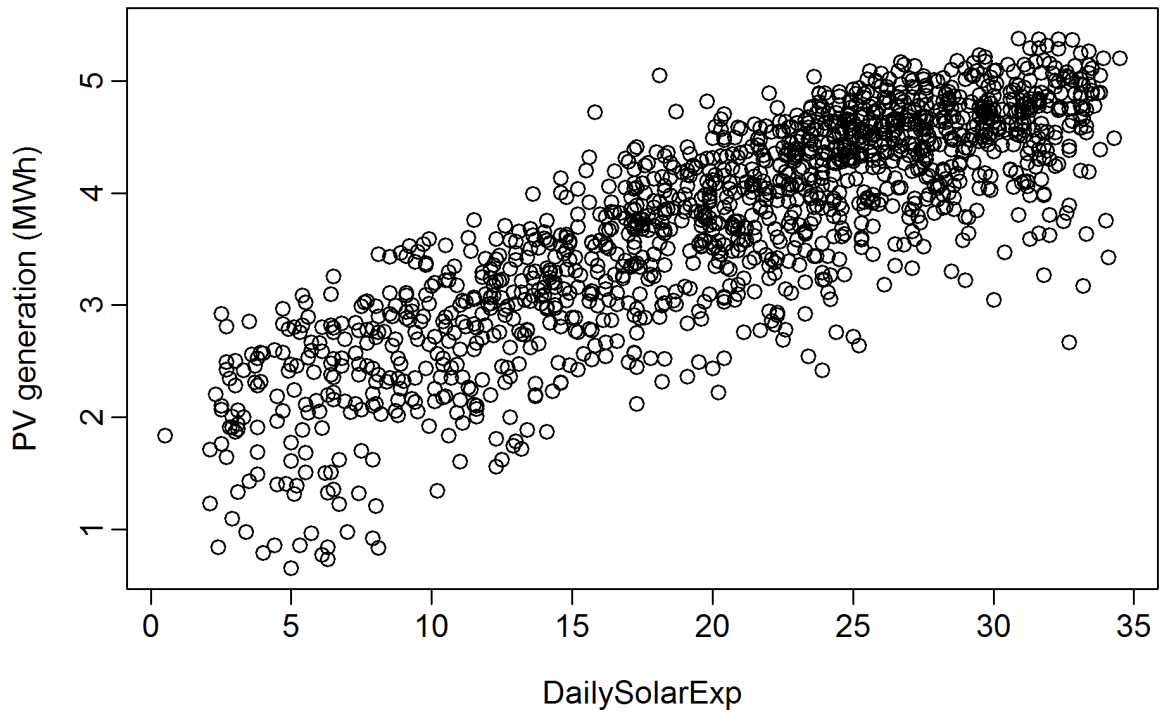


Figure 14: Daily solar PV generation plotted against daily solar exposure data in NSW from 2003 to 2011. Only data from October–March are shown.

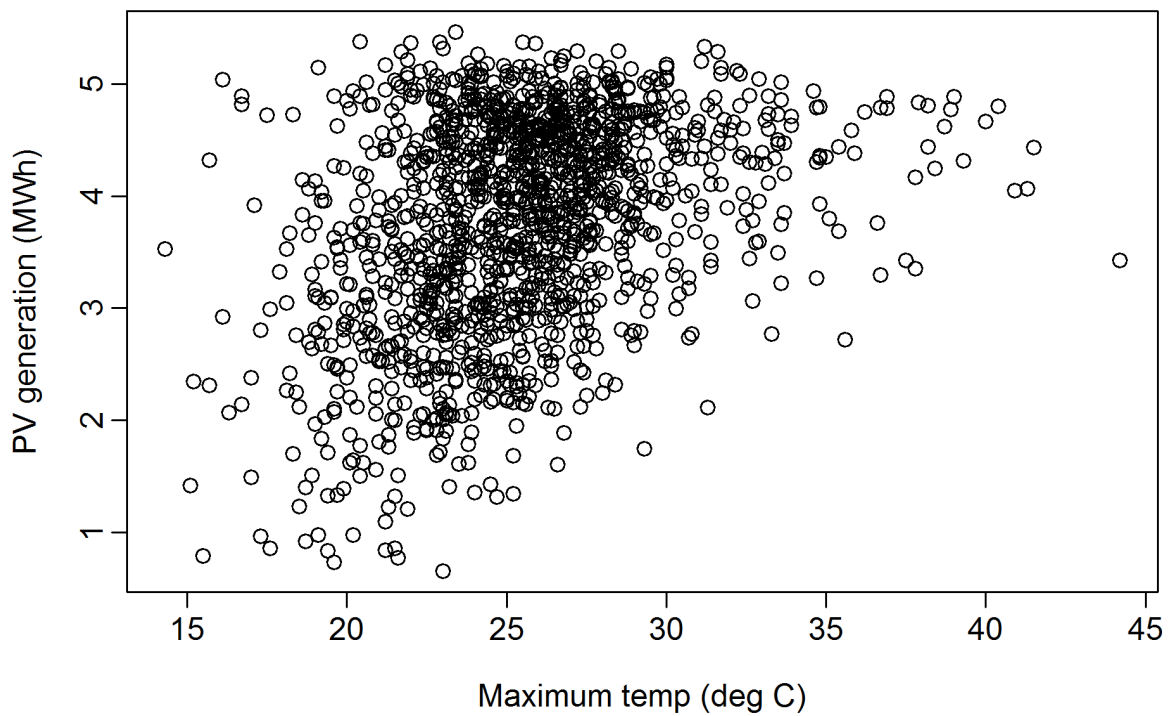


Figure 15: Daily solar PV generation plotted against daily maximum temperature in NSW from 2003 to 2011. Only data from October–March are shown.

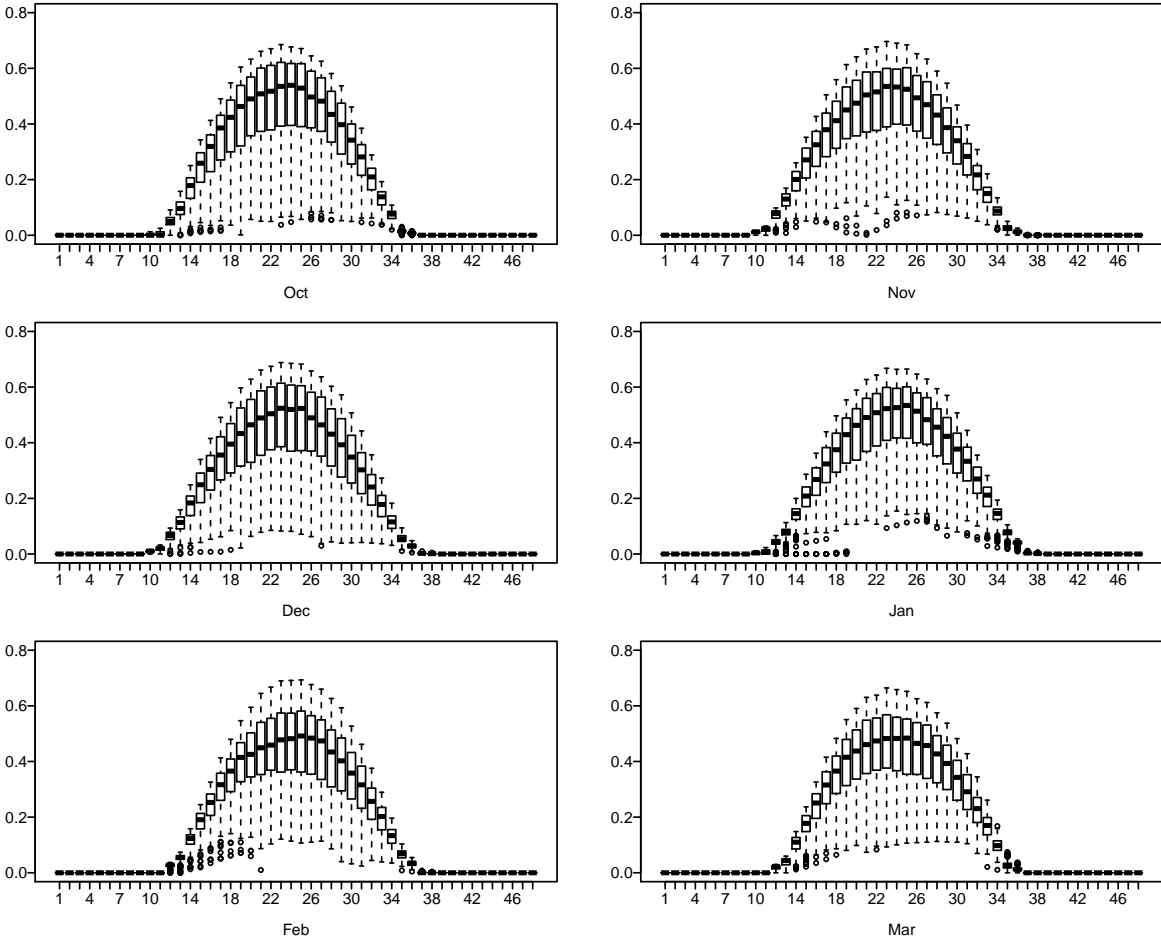


Figure 16: Boxplot of historical PV output based on a 1 MW solar system. Only data from October–March are shown.

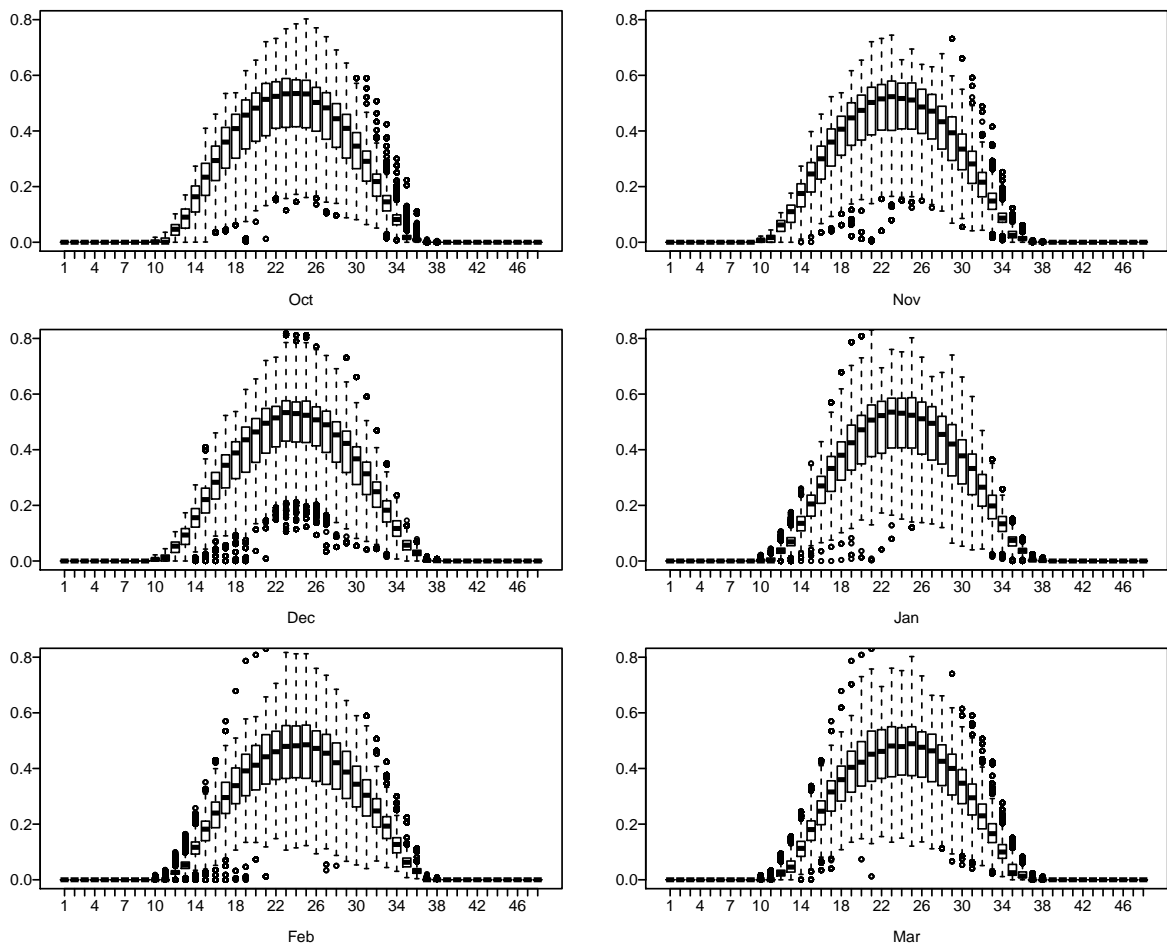


Figure 17: Boxplot of simulated PV output based on a 1 MW solar system. Only data from October–March are shown.

1.6 Probability distribution reproduction

To validate the effectiveness of the proposed temperature simulation, we reproduce the probability distribution of historical demand using known economic values but simulating temperatures, and then compare with the real demand data.

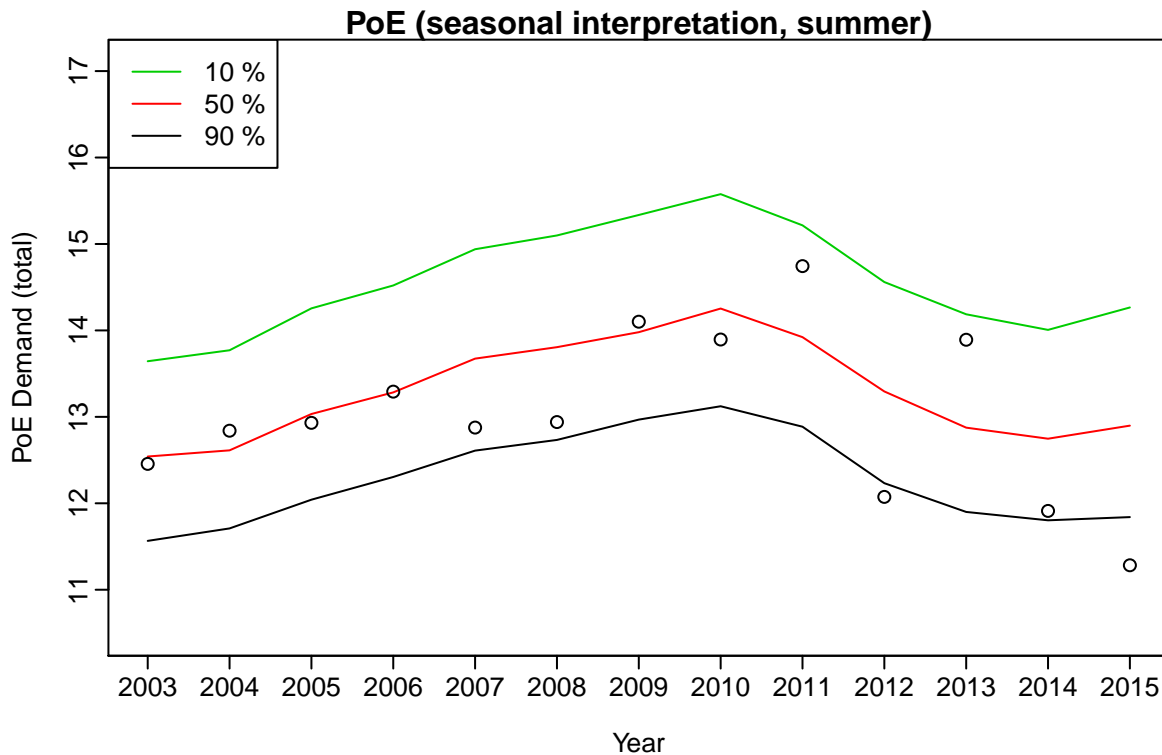


Figure 18: *PoE levels of demand excluding offset loads for all years of historical data using known economic values but simulating temperatures.*

Figure 18 shows PoE levels for all years of historical data using known economic values but simulating temperatures. These are based on the distribution of the seasonal maximum in each year. Figure 19 shows similar levels for the weekly maximum demand. Of the 13 historical seasonal maximum demand values, there are 5 seasonal maximums above the 50% PoE level, no seasonal maximums above the 10% PoE levels and 2 seasonal maximums below the 90% PoE level. It can be inferred that the results in Figure 18 and Figure 19 are within the ranges that would occur by chance with high probability if the forecast distributions are accurate. This provides good evidence that the forecast distributions will be reliable for future years.

On the other hand, Figure 18 shows that the most recent seasonal maximum demand value is very low compared to the forecasted probability distribution, and the previous year's value was also relatively low. If another low value was seen in the next year, this may suggest that our model needs adjusting to take account of recent changes in demand response to the predictor variables.

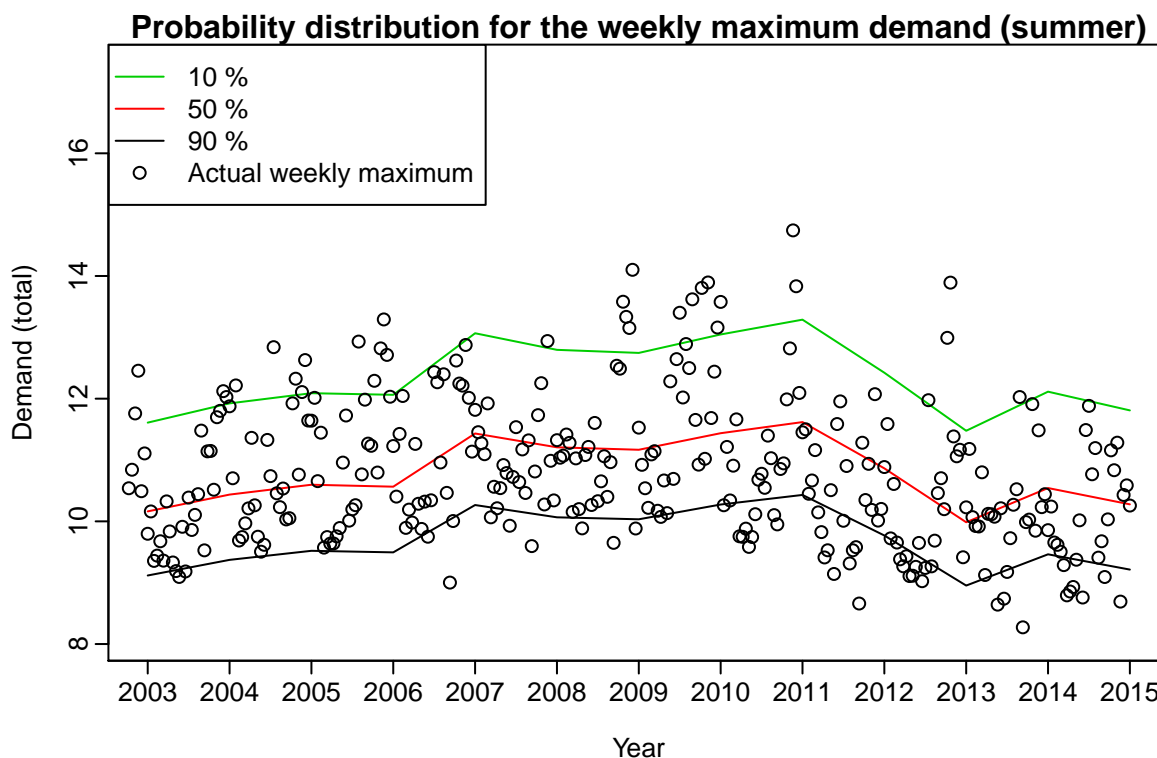


Figure 19: Probability distribution of the weekly maximum demand for all years of historical data using known economic values but simulating temperatures.

1.7 Demand forecasting

Forecasts of the distribution of demand are computed by simulation from the fitted model as described in Hyndman and Fan (2015).

The temperatures at Sydney and Richmond are simulated from historical values observed in 2002–2015 using a double seasonal bootstrap with variable length (Hyndman and Fan, 2015). The temperature bootstrap is designed to capture the serial correlation that is present in the data due to weather systems moving across New South Wales. We also add a small amount of noise to the simulated temperature data to allow temperatures to be as high as any observed since 1900. A total of more than 1000 years of temperature profiles were generated in this way for each year to be forecast.

In this report, some simple climate change adjustments are made to allow for the possible effects of global warming. Estimates were taken from CSIRO modelling (Climate Change in Australia). The shifts in temperature till 2030 for mainland Australia are predicted to be 0.6°C, 0.9°C and 1.3°C for the 10th percentile, 50th percentile and 90th percentile respectively, and the shifts in temperature till 2030 for Tasmania are predicted to be 0.5°C, 0.8°C and 1.0°C for the 10th percentile, 50th percentile and 90th percentile respectively (CSIRO, 2015). The temperature projections are given relative to the period 1986–2005 (referred to as the 1995 baseline for convenience). These shifts are implemented

linearly from 2015 to 2035. CSIRO predicts that the change in the standard deviation of temperature will be minimal in Australia. This matches work by (Räisänen, 2002).

Three different scenarios for population, GSP, total electricity price, major industrial loads and installed capacity of rooftop photo-voltaic (PV) cells are considered. The historical and assumed future values for the population, GSP and total electricity price are shown in Figure 20, and the 50% POE major industrial loads and PV installed capacity are shown in Figure 21 and Figure 22 respectively.




Figure 20: Three scenarios for New South Wales population, GSP and total electricity price.

The major industrial demand is simulated simultaneously with temperature using the same double seasonal bootstrap with variable length (Hyndman and Fan, 2015). We do not model the major industrial demand as there are no reliable drivers available (this demand is not correlated to temperature). The simulated industrial demand is then scaled so that the 50% POE and 10% POE are equal to the projections provided by AEMO. Specifically, the simulated major industrial demand o_t is transformed using the following linear transformation $o_t^* = (o_t - a)/b$ where a and b are chosen to ensure the 50% POE and 10% POE of o_t^* are equal to the projected values. If \bar{o}_t denotes the projected 50% POE, o_t^* denotes the projected 10% POE, and o_t^+ is the 10% POE of the simulated o_t , o_t^- is the 50% POE of the simulated o_t , then $b = (o_t^+ - o_t^-)/(o_t^* - \bar{o}_t)$ and $a = o_t^- - b\bar{o}_t$.

The simulated temperatures, known calendar effects, assumed values of GSP, electricity price and PV installed capacity and simulated residuals are all combined using the fitted statistical model to give simulated electricity demand for every half-hourly period in the years to be forecast. Thus,




Figure 21: Three scenarios for major industrial loads.




Figure 22: Three scenarios for installed capacity of rooftop photo-voltaic (PV) cells.

we are predicting what could happen in the years 2015/16–2034/35 under these fixed economic

and demographic conditions, but allowing for random variation in temperature events and other conditions.

1.7.1 Probability distributions

Figure 23 shows the simulated seasonal maximum demand densities for 2015/16–2034/35, and Figure 24 shows quantiles of prediction of seasonal maximum demand.

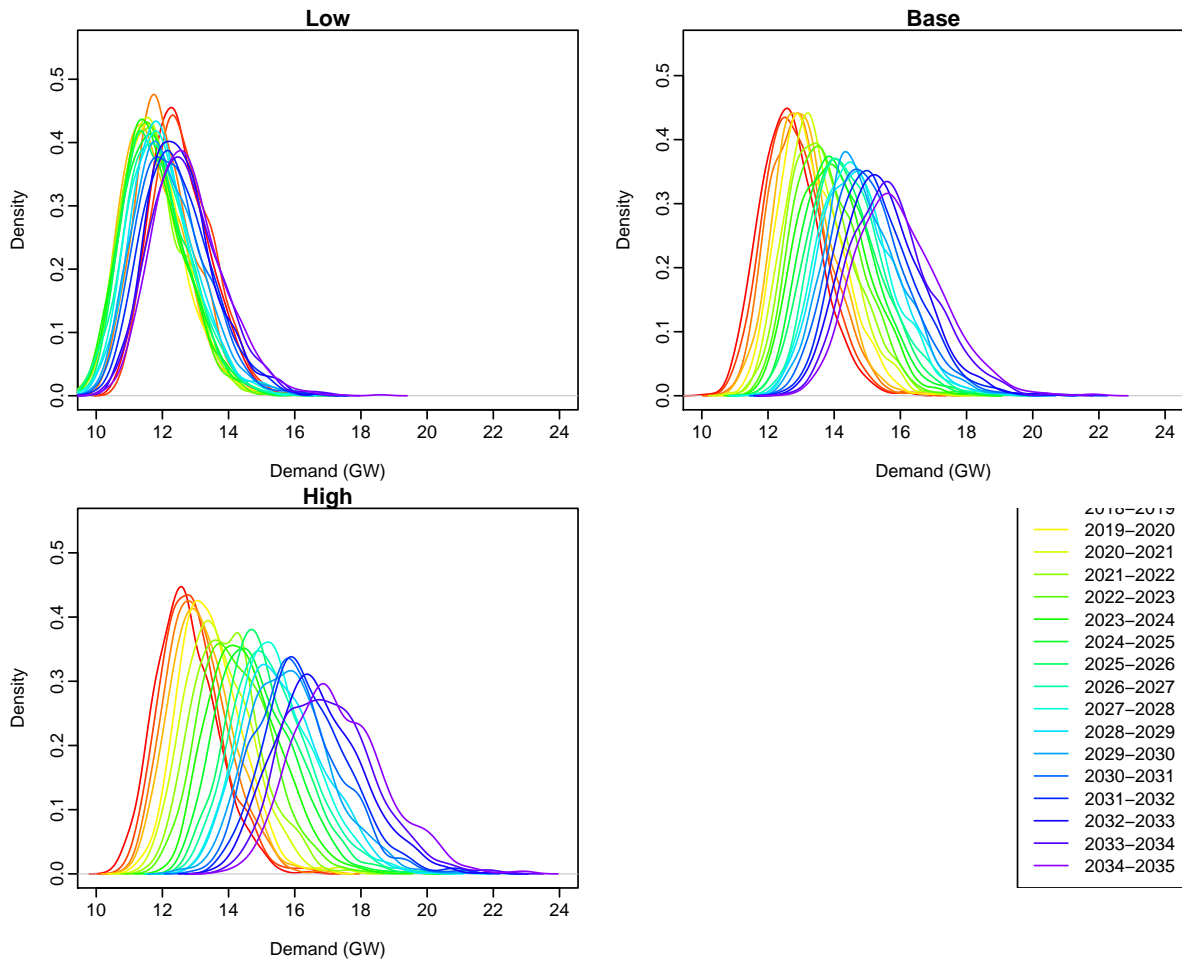


Figure 23: Distribution of simulated seasonal maximum demand for 2015/16–2034/35.

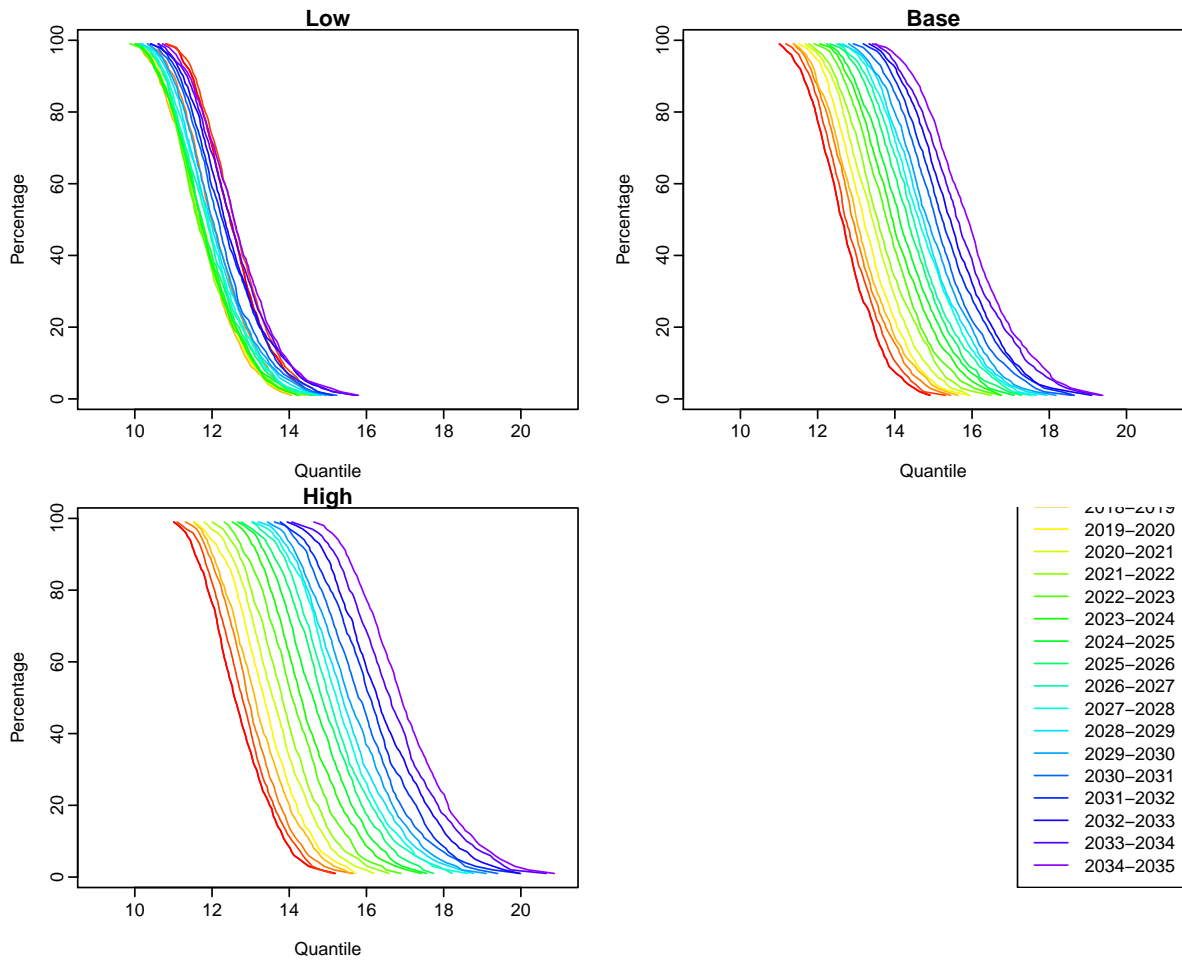


Figure 24: *Quantiles of prediction of seasonal maximum demand for 2015/16–2034/35.*

1.7.2 Probability of exceedance

AEMO also requires calculation of “probability of exceedance” levels. These can now be obtained from the simulated data. Suppose we are interested in the level x such that the probability of the maximum demand exceeding x over a 1-year period is p .

If Y_1, \dots, Y_m denote m simulated seasonal maxima from a given year, then we can estimate x as the q th quantile of the distribution of $\{Y_1, \dots, Y_m\}$ where $q = 1 - p$ (Hyndman and Fan, 1996).

PoE values based on the seasonal maxima are plotted in Figure 25, along with the historical PoE values and the observed seasonal maximum demand values. Figure 26 shows similar levels for the weekly maximum demand.

It should be noted that all of the PoE values given in the tables in this report are conditional on the economic and demographic driver variables observed in those years. They are also conditional on the fitted model. However, they are not conditional on the observed temperatures. They take into account the possibility of different temperature patterns from those observed.




Figure 25: Probability of exceedance values for past and future years. Forecasts based on the base scenario are shown as solid lines; forecasts based on the low scenario are shown as dashed lines; Forecasts based on the high scenario are shown as dotted lines.




Figure 26: Probability distribution levels for past and future years. Forecasts based on the base scenario are shown as solid lines; forecasts based on the low scenario are shown as dashed lines; Forecasts based on the high scenario are shown as dotted lines.

2 Modelling and forecasting electricity demand of winter

2.1 Historical data

We define the period April–September as “winter” for the purposes of this report. Only data from April–September were retained for analysis and modelling winter demand. Thus, each “year” consists of 183 days.

Time plots of the half-hourly demand data are plotted in Figures 27–29. These clearly show the intra-day pattern (of length 48) and the weekly seasonality (of length $48 \times 7 = 336$); the seasonal seasonality (of length $48 \times 183 = 8784$) is less obvious.

The half-hourly aggregated major industrial demand is given in Figure 30.

AEMO provided half-hourly rooftop generation data based on a 1MW solar system in New South Wales from 2003 to 2011 in Figures 31, and the half-hourly values of total rooftop generation of New South Wales from 2003 to 2015 in Figures 32, by integrating the data from the 1MW solar system and the installed capacity of rooftop generation.

The relationship between demand (excluding major industrial loads) and the average temperature is shown in Figure 33.

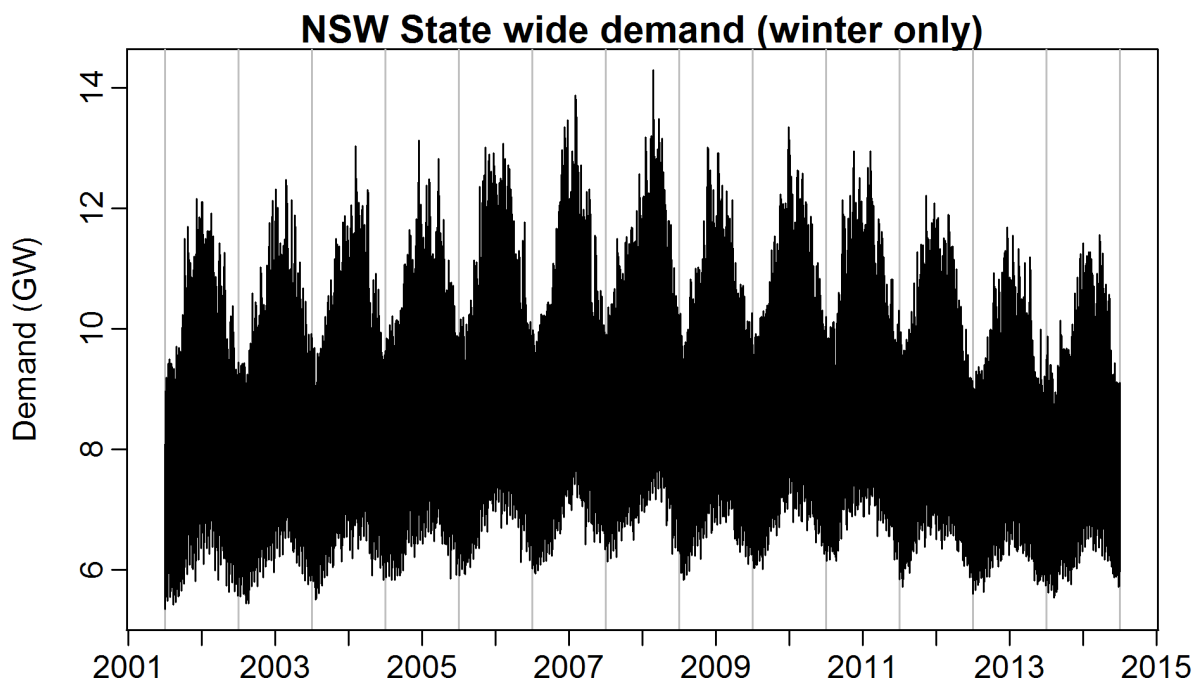


Figure 27: Half-hourly demand data for New South Wales from 2002 to 2015. Only data from April–September are shown.

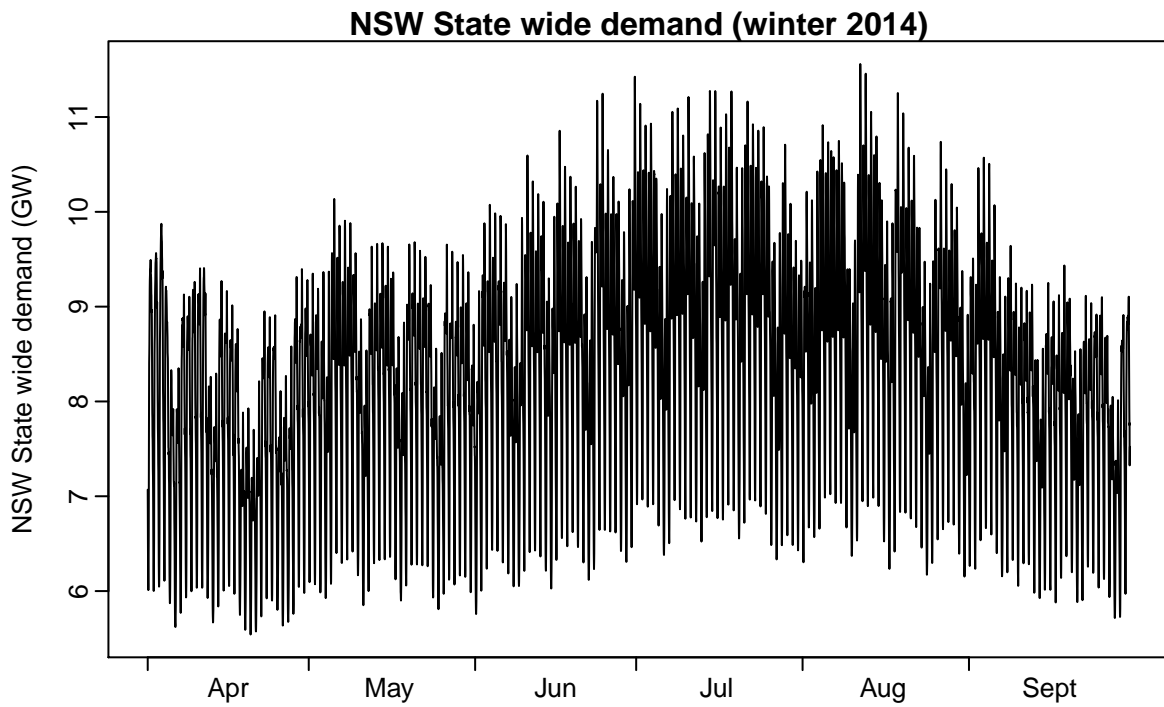


Figure 28: Half-hourly demand data for New South Wales for last winter.

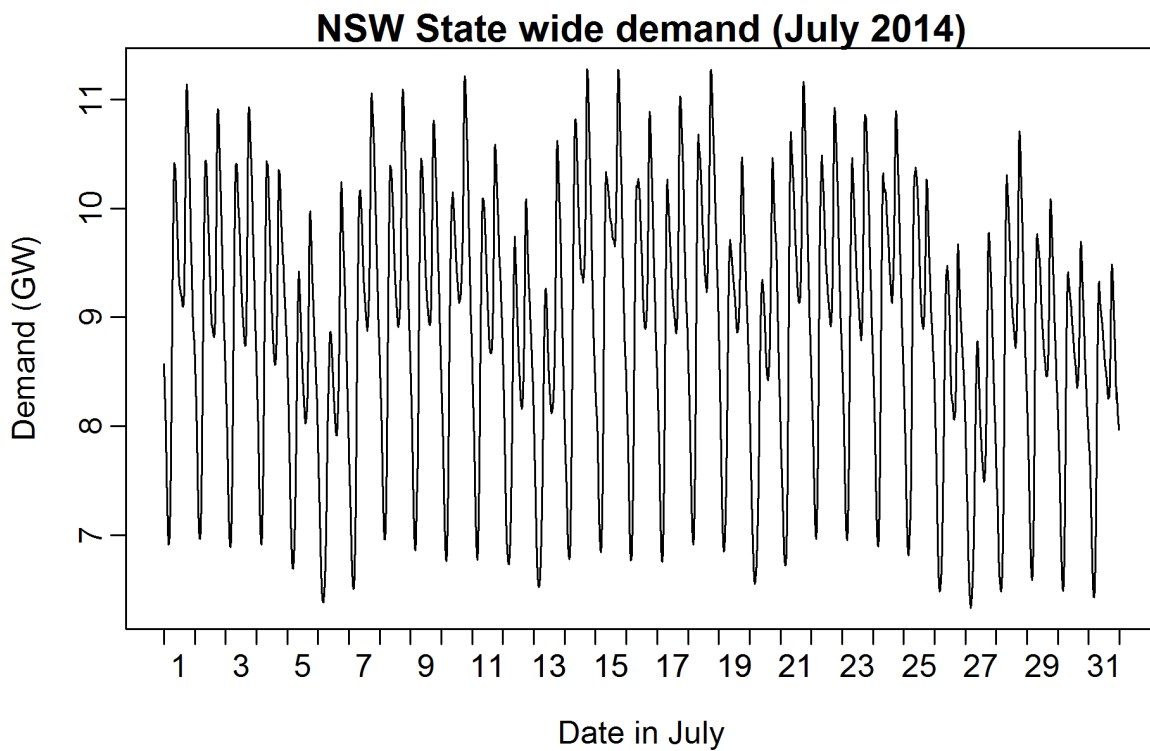


Figure 29: Half-hourly demand data for New South Wales, July 2014.




Figure 30: Half-hourly demand data for aggregated major industries from 2002 to 2015.

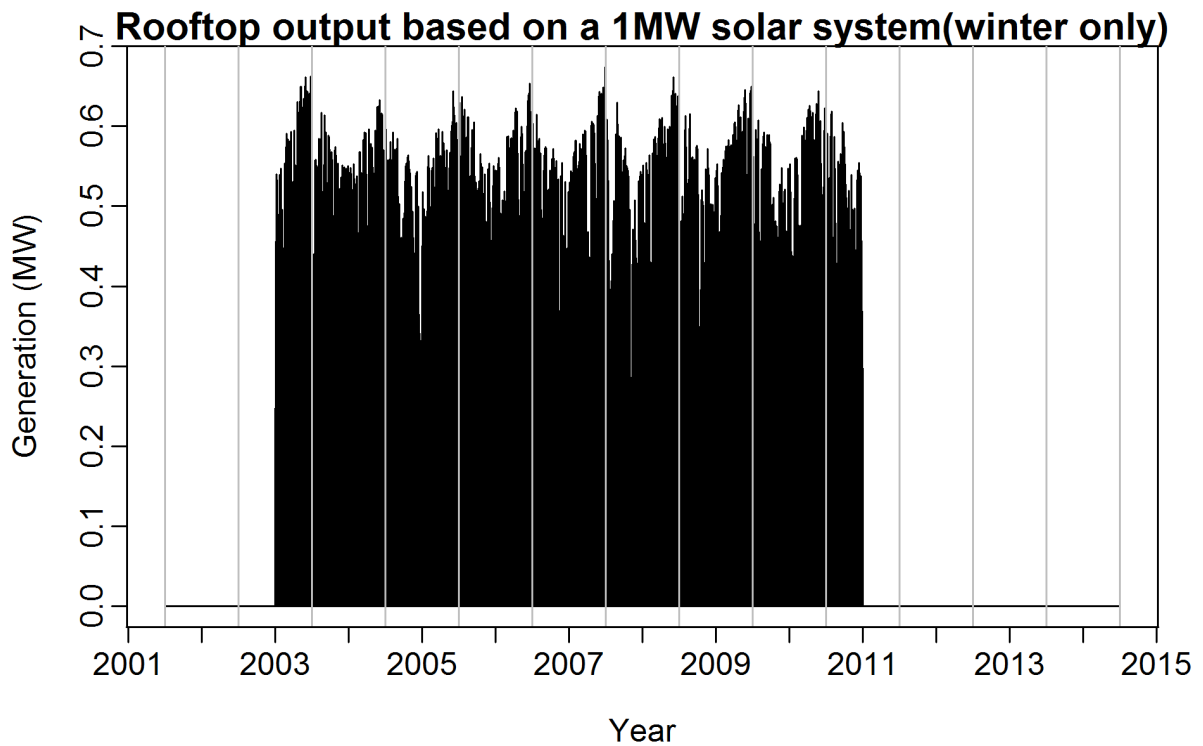


Figure 31: Half-hourly rooftop generation data based on a 1MW solar system. 2003–2011.




Figure 32: Half-hourly rooftop generation data. 2003–2015.




Figure 33: Half-hourly NSW electricity demand (excluding major industrial demand) plotted against average temperature (degrees Celsius).

2.2 Variable selection for the half-hourly model

To fit the model to the demand excluding major industrial loads, we normalize the half-hourly demand against seasonal average demand of each year. The top panel of Figure 34 shows the original demand data with the average seasonal demand values shown in red, and the bottom panel shows the half-hourly adjusted demand data.

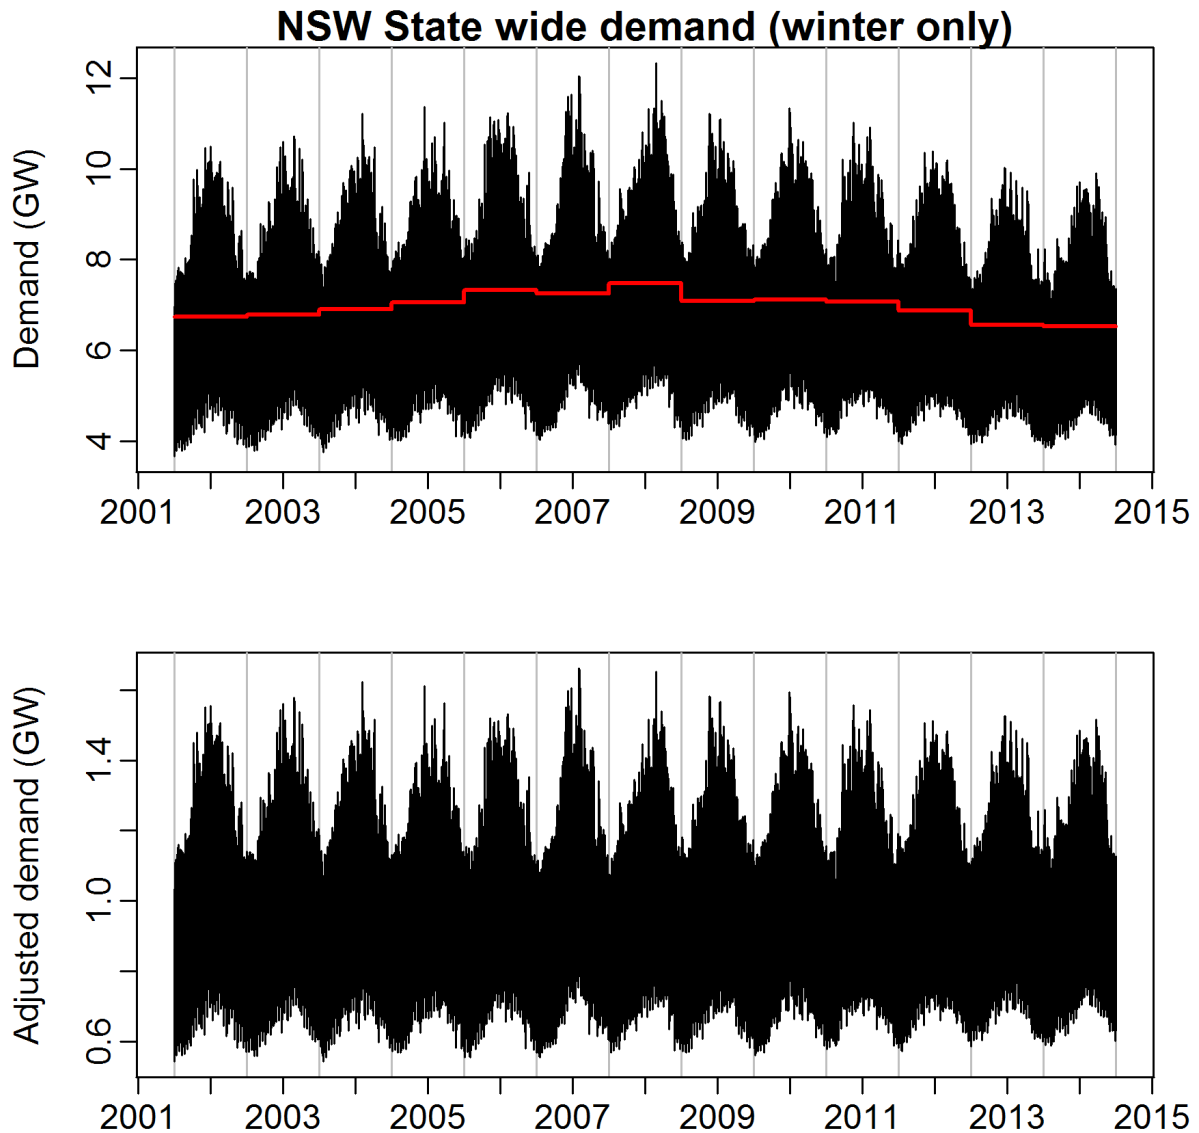


Figure 34: Top: Half-hourly demand data for New South Wales from 2002 to 2015. Bottom: Adjusted half-hourly demand where each year of demand is normalized against seasonal average demand. Only data from April–September are shown.

We fit a separate model to the adjusted demand data from each half-hourly period. Four boosting steps are adopted to improve the model fitting performance (Hyndman and Fan, 2015), the linear model in boosting stage 1 contained the following variables:

- ▶ the current temperature and temperatures from the last 2 hours;
- ▶ temperatures from the same time period for the last 6 days;
- ▶ the current temperature differential and temperature differentials from the last 2.5 hours;
- ▶ temperature differentials from the same period of previous day;
- ▶ the maximum temperature in the last 24 hours;
- ▶ the minimum temperature in the last 24 hours;
- ▶ the average temperature in the last 24 hours;
- ▶ the day of the week;
- ▶ the holiday effect;
- ▶ the day of season effect.

2.3 Model predictive capacity

We investigate the predictive capacity of the model by looking at the fitted values. Figure 35 shows the actual historical demand (top) and the fitted (or predicted) demands for the entire winter season. Figure 36 illustrates the model prediction for July 2014. It can be seen that the actual and fitted values are almost indistinguishable, indicating that the vast majority of the variation in the data has been accounted for through the driver variables. Both fitted and actual values shown here are after the major industrial load has been subtracted from the data.

Note that these “predicted” values are not true forecasts as the demand values from these periods were used in constructing the statistical model. Consequently, they tend to be more accurate than what is possible using true forecasts.

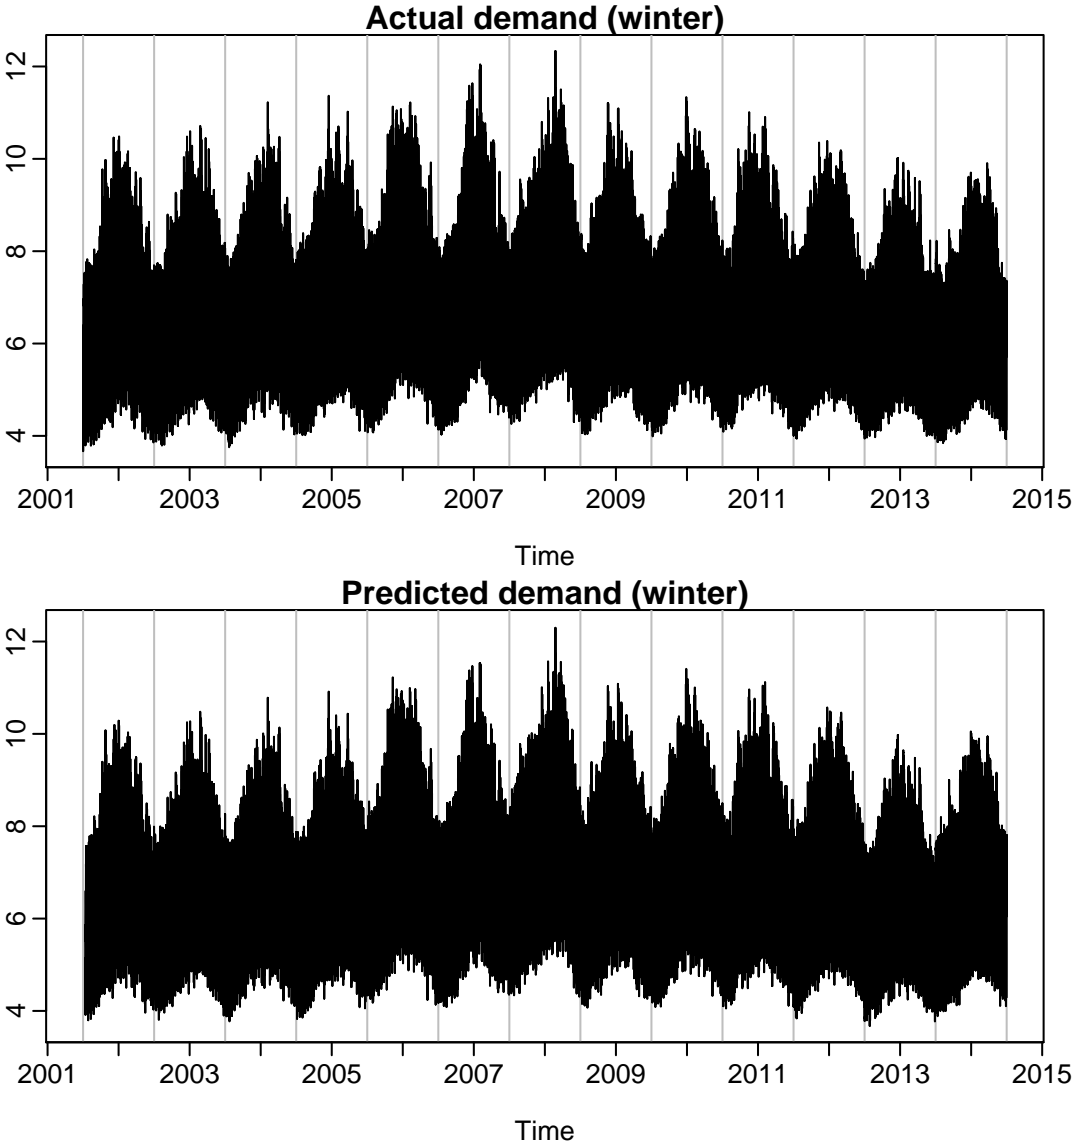


Figure 35: Time plots of actual and predicted demand.

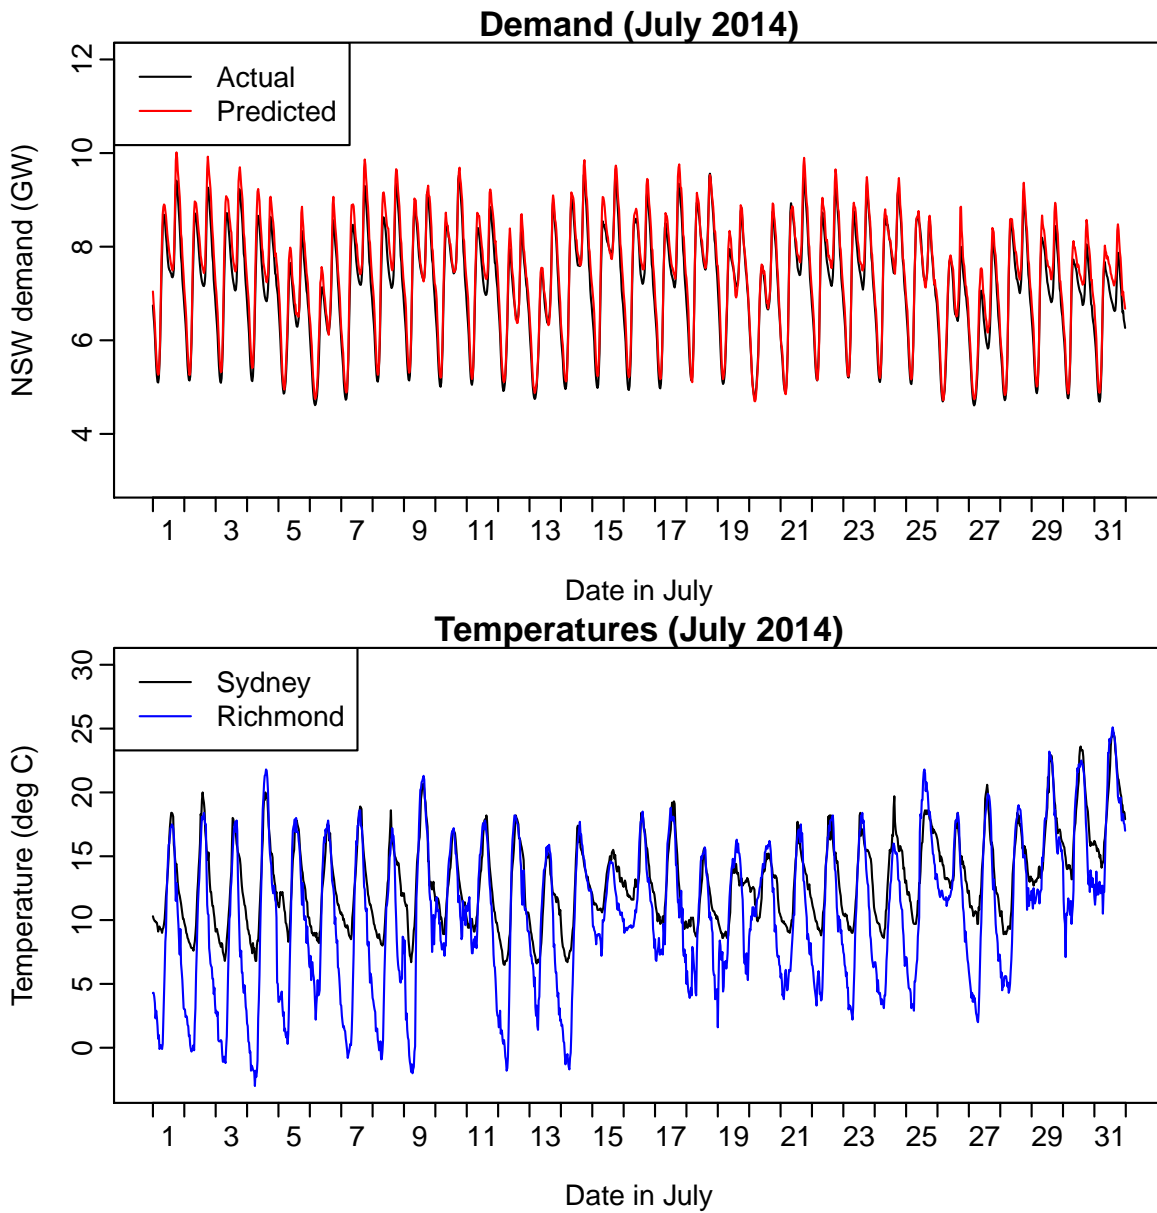


Figure 36: Actual and predicted demand for July 2014.

2.4 Half-hourly model residuals

The time plot of the half-hourly residuals from the demand model is shown shown in Figure 37.




Figure 37: *Half-hourly residuals (actual – predicted) from the demand model.*

Next we plot the half-hourly residuals against the predicted adjusted log demand in Figure 38. That is, we plot e_t against $\log(y_{t,p}^*)$ where these variables are defined in the demand model (Hyndman and Fan, 2015).




Figure 38: *Residuals vs predicted log adjusted demand from the demand model.*

2.5 Modelling and simulation of PV generation

Figure 39 shows the relationship between the daily PV generation and the daily solar exposure in NSW from 2003 to 2011, and the strong correlation between the two variables is evident. The daily PV generation against daily maximum temperature is plotted for the same period in Figure 40, and we can observe the positive correlation between the two variables. Accordingly, the daily PV exposure and maximum temperature are considered in the PV generation model (Hyndman and Fan, 2015).

We simulate future half-hourly PV generation in a manner that is consistent with both the available historical data and the future temperature simulations. To illustrate the simulated half-hourly PV generation, we plot the boxplot of simulated PV output based on a 1 MW solar system in Figure 42, while the boxplot of the historical ROAM data based on a 1 MW solar system is shown in Figure 41. Comparing the two figures, we can see that the main features of the two data sets are generally consistent. Some more extreme values are seen in the simulated data set — these arise because there are many more observations in the simulated data set, so the probability of extremes occurring somewhere in the data is much higher. However, the quantiles are very similar in both plots showing that the underlying distributions are similar.




Figure 39: Daily solar PV generation plotted against daily solar exposure data in NSW from 2003 to 2011. Only data from April–September are shown.




Figure 40: Daily solar PV generation plotted against daily maximum temperature in NSW from 2003 to 2011. Only data from April–September are shown.

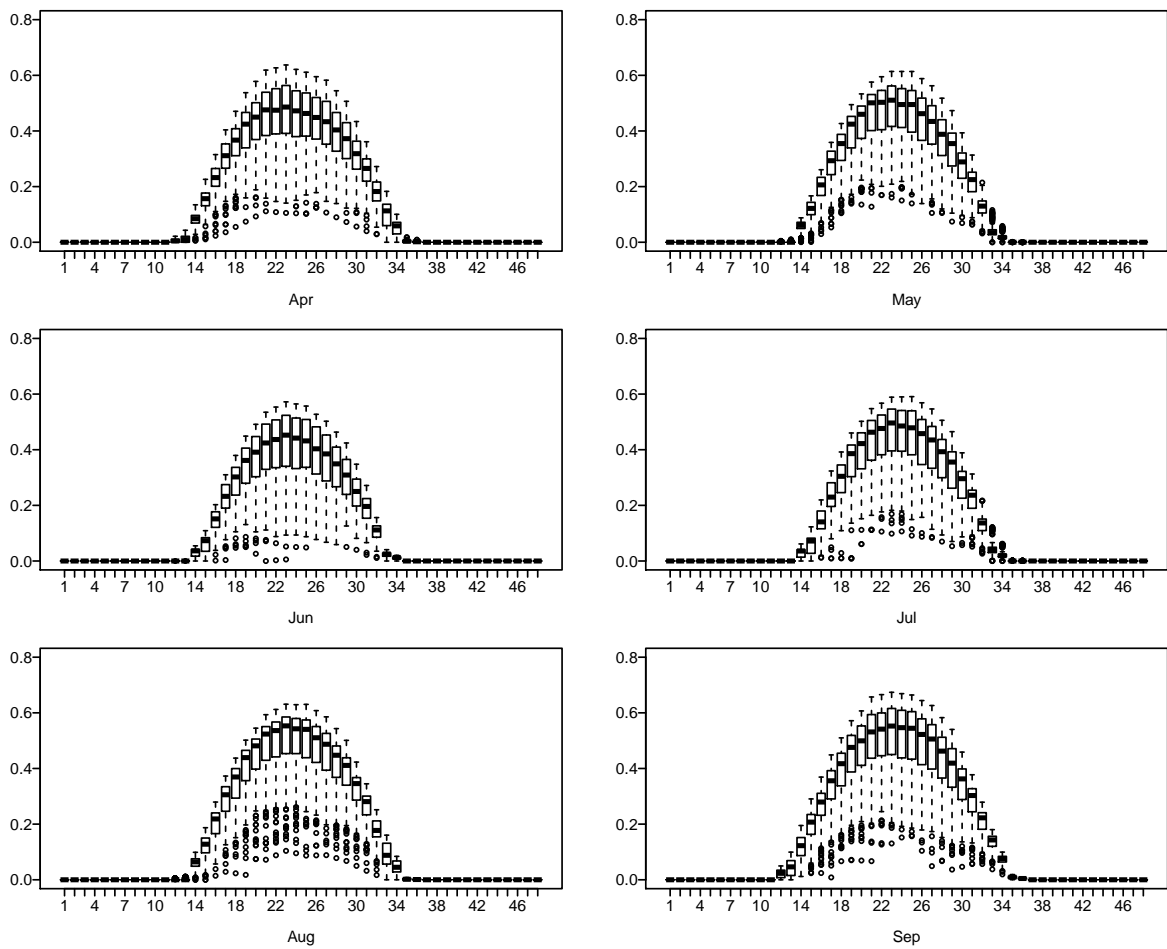


Figure 41: Boxplot of historical PV output based on a 1 MW solar system. Only data from April–September are shown.

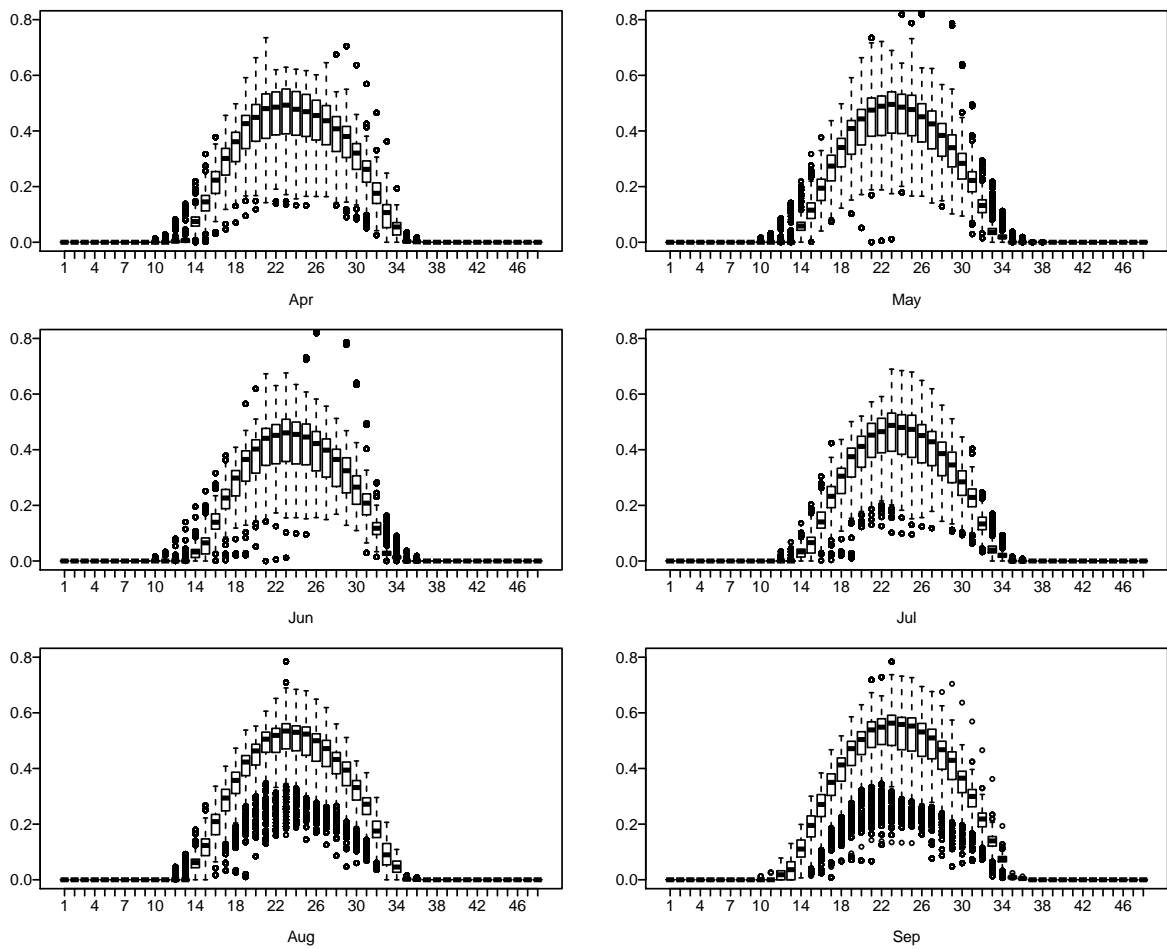


Figure 42: *Boxplot of simulated PV output based on a 1 MW solar system. Only data from April–September are shown.*

2.6 Probability distribution reproduction

To validate the effectiveness of the proposed temperature simulation, we reproduce the probability distribution of historical demand using known economic values but simulating temperatures, and then compare with the real demand data.




Figure 43: PoE levels of demand excluding offset loads for all years of historical data using known economic values but simulating temperatures.

Figure 43 shows PoE levels for all years of historical data using known economic values but simulating temperatures. These are based on the distribution of the seasonal maximum in each year. Figure 44 shows similar levels for the weekly maximum demand. Of the 12 historical seasonal maximum demand values, there are 6 seasonal maximums above the 50% PoE level, 1 seasonal maximum above the 10% PoE levels and 1 seasonal maximum below the 90% PoE level. It can be inferred that the results in Figure 43 and Figure 44 are within the ranges that would occur by chance with high probability if the forecast distributions are accurate. This provides good evidence that the forecast distributions will also be reliable for future years.

On the other hand, Figure 43 shows that the most recent seasonal maximum demand value is very low compared to the forecasted probability distribution, and the previous year's value was also relatively low. If another low value was seen in the next year, this may suggest that our model needs adjusting to take account of recent changes in demand response to the predictor variables.

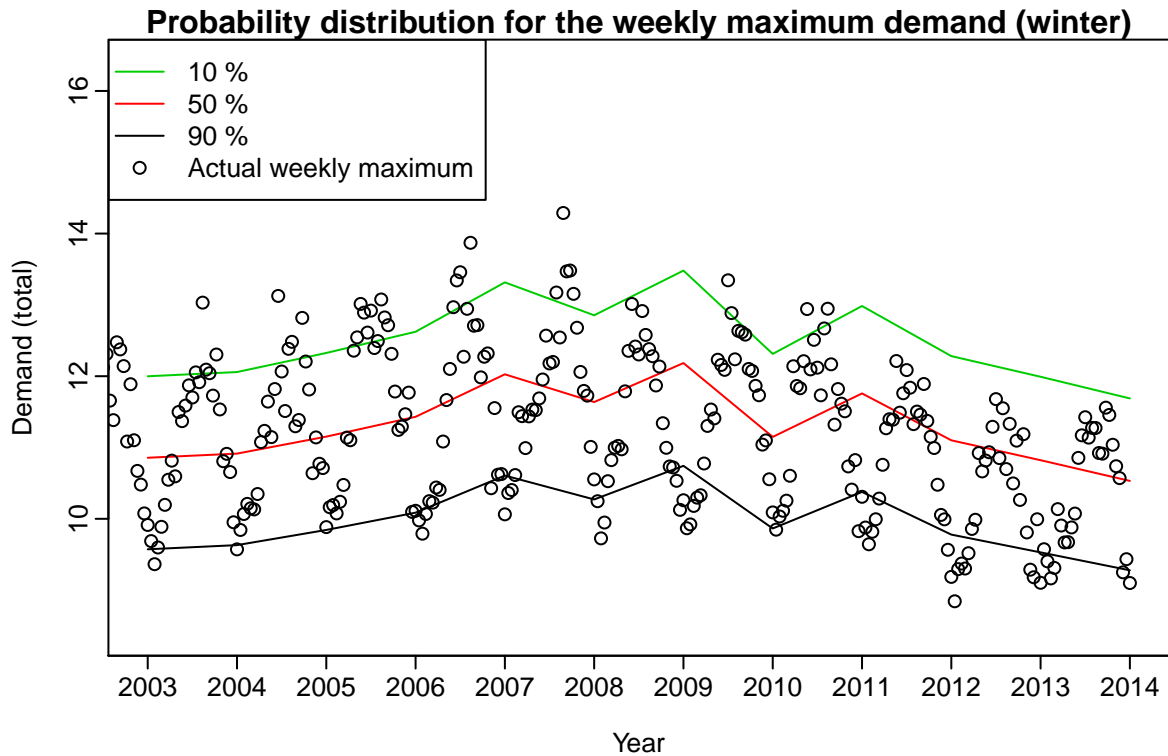


Figure 44: Probability distribution of the weekly maximum demand for all years of historical data using known economic values but simulating temperatures.

2.7 Demand forecasting

2.7.1 Probability distributions

The historical and assumed future values for the population, GSP and total electricity price are shown in Figure 45, and 50% POE of the major industrial loads and PV installed capacity are shown in Figure 46 and Figure 47 respectively.

Figure 48 shows the simulated seasonal maximum demand densities for 2015 – 2034, and Figure 49 shows quantiles of prediction of seasonal maximum demand.




Figure 45: Three future scenarios for New South Wales population, GSP and total electricity price.




Figure 46: Three scenarios for major industrial loads.




Figure 47: Three scenarios for installed capacity of rooftop photo-voltaic (PV) cells.

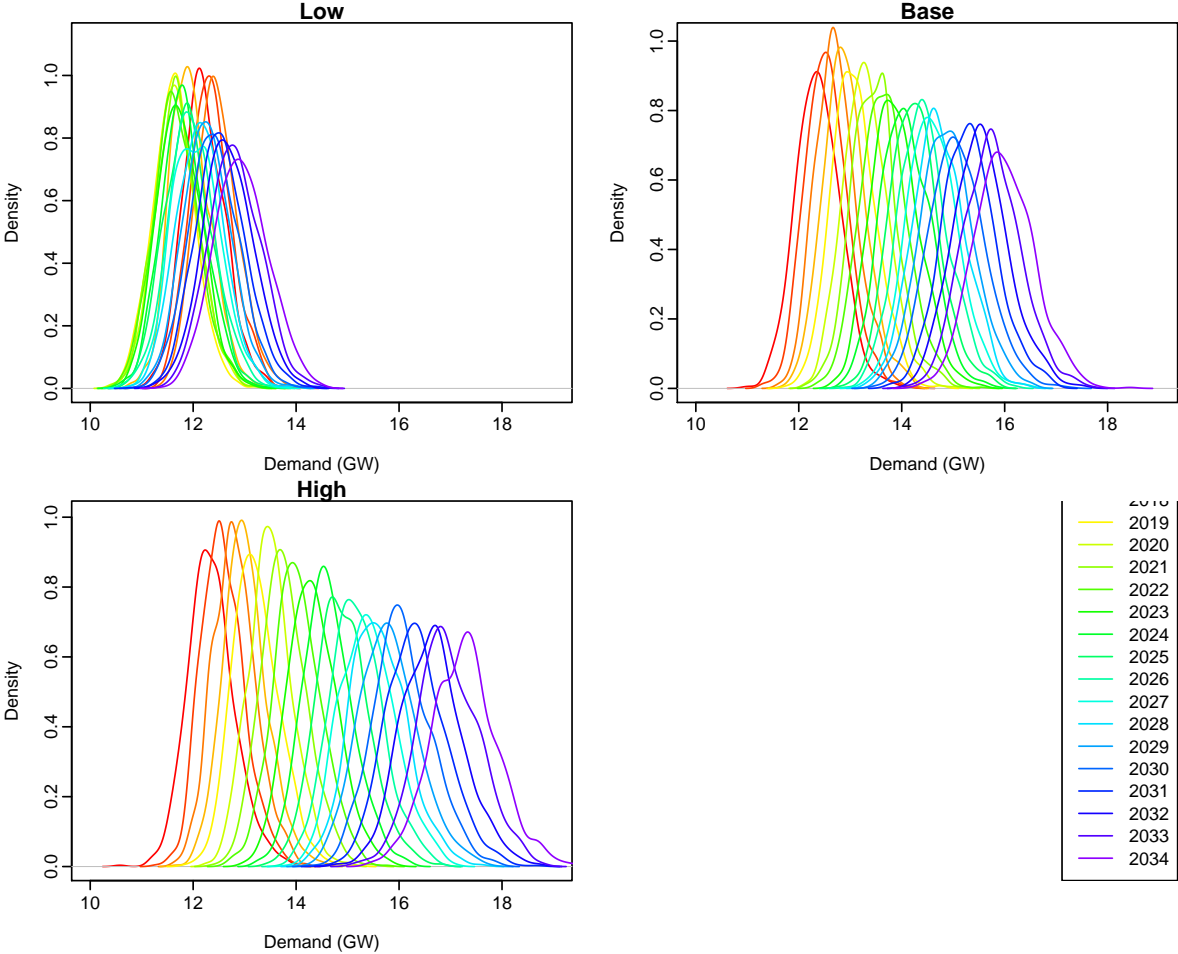


Figure 48: Distribution of simulated seasonal maximum demand for 2015 – 2034.

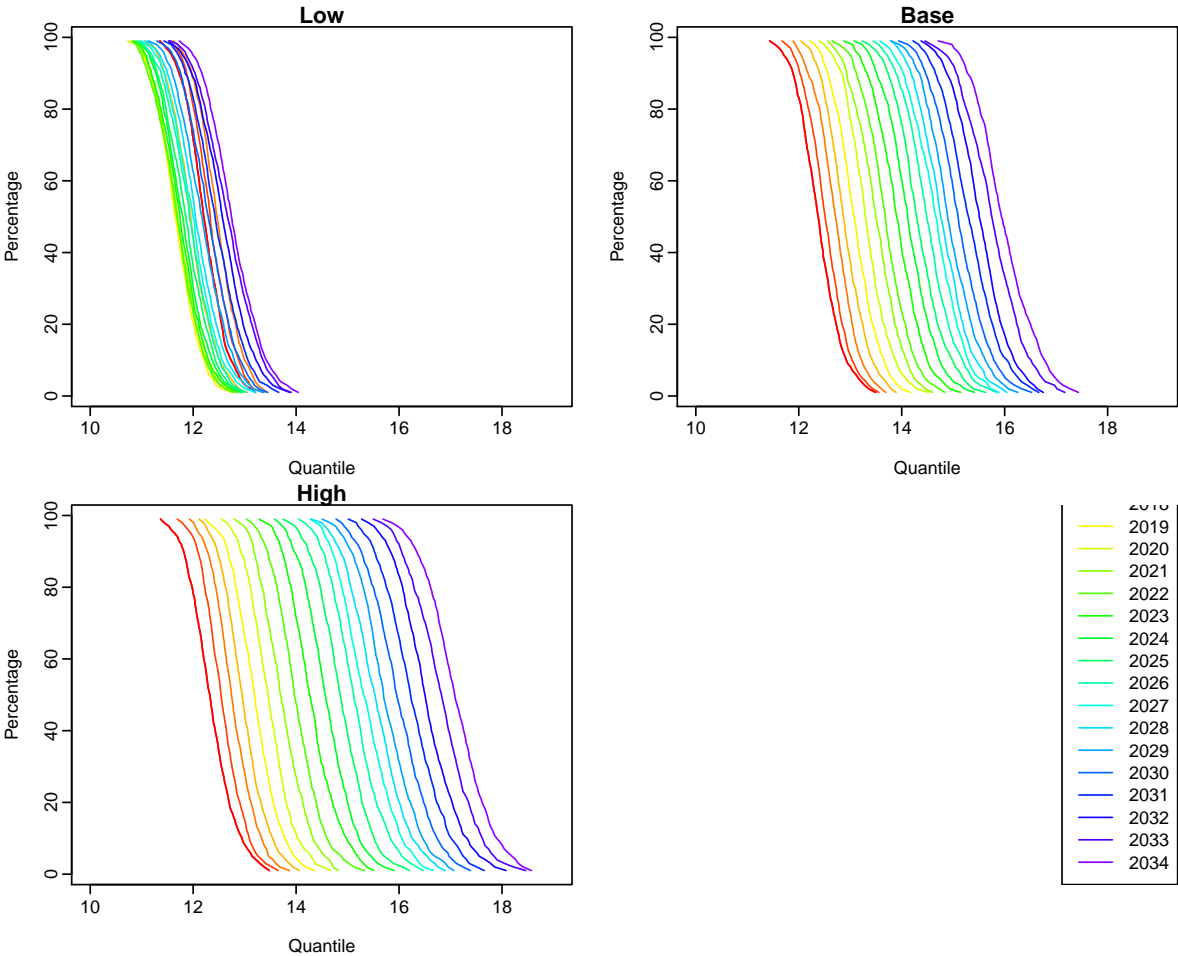


Figure 49: Quantiles of prediction of seasonal maximum demand for 2015 – 2034.

2.7.2 Probability of exceedance

PoE values based on the seasonal maxima are plotted in Figure 50, along with the historical PoE values and the observed seasonal maximum demand values. Figure 51 shows similar levels for the weekly maximum demand.

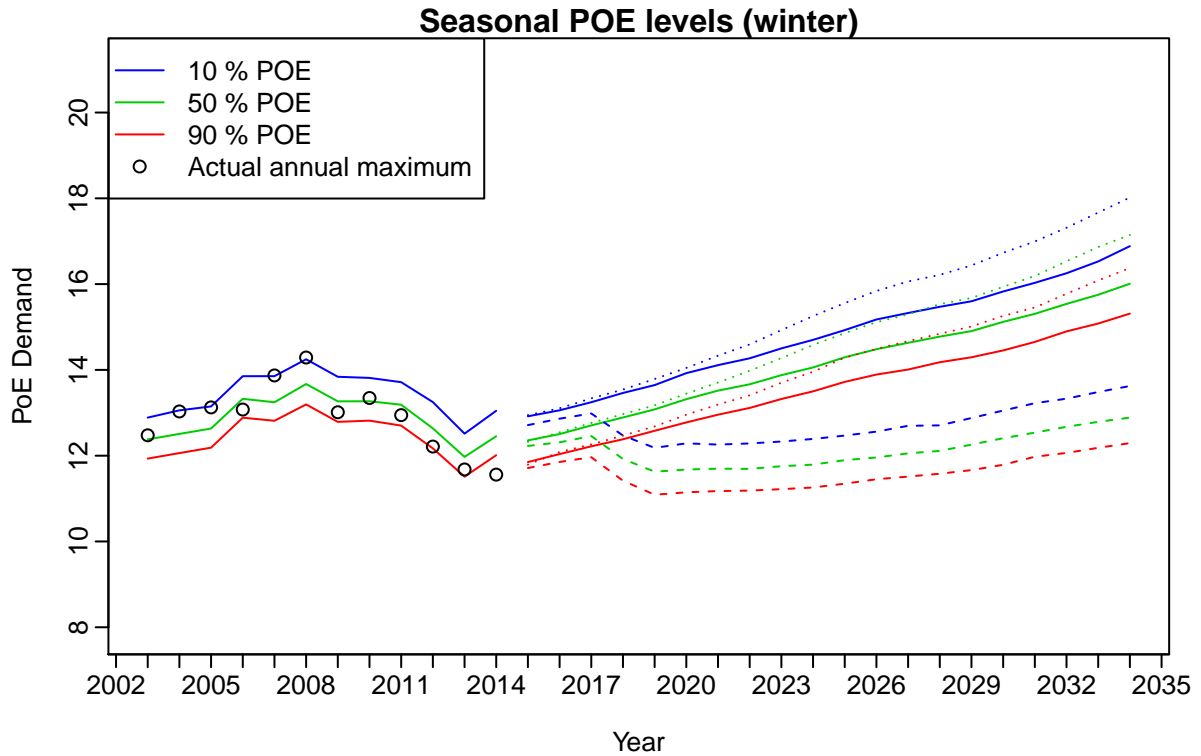


Figure 50: Probability of exceedance values for past and future years. Forecasts based on the base scenario are shown as solid lines; forecasts based on the low scenario are shown as dashed lines; Forecasts based on the high scenario are shown as dotted lines.




Figure 51: Probability of exceedance values for past and future years. Forecasts based on the base scenario are shown as solid lines; forecasts based on the low scenario are shown as dashed lines; Forecasts based on the high scenario are shown as dotted lines.

References

- CSIRO (2015). *Technical Report - Climate Change in Australia - Projections for Australia's NRM Regions*. Accessed: 2015-3-05. <http://www.climatechangeinaustralia.gov.au/en/publications-library/technical-report/>.
- Hyndman, RJ and S Fan (2015). *Monash Electricity Forecasting Model*. Report for Australian Energy Market Operator (AEMO). <http://robjhyndman.com/papers/MEFMR1.pdf>.
- Hyndman, RJ and Y Fan (1996). Sample quantiles in statistical packages. *The American Statistician* **50**(4), 361–365.
- Räisänen, J (2002). CO₂-Induced Changes in Interannual Temperature and Precipitation Variability in 19 CMIP2 Experiments. *Journal of Climate* **15**(17), 2395–2411.