DRAFT

Updates of

a regression model relating electricity spot price and hydro storage (PH Model)

and

a seasonal switching model for hydro storage (SH Model)

using South Island data.

undertaken for the

New Zealand Electricity Authority

by

Peter Thomson

Statistics Research Associates Ltd

Wellington, New Zealand

19 April 2018

Contents

Summary **1**

1 Scope of report and terms of reference **3**

2 Background **4**

2.1 2009 market reforms 4

2.2 Data 5

3 PH model update **9**

3.1 Electricity spot prices 9

3.2 Hydro storage 13

3.3 Relationship between hydro storage and electricity spot price 19

3.4 Summary 26

4 SH model update **27**

4.1 Exploratory analysis with non-seasonal switching model 29

4.2 Seasonal switching model 36

4.3 Price-storage relationship within seasonal state 43

4.4 Summary 48

Acknowledgements **49**

References **50**

Summary

This report updates the models and analyses given in two previous reports (Thomson 2013, 2014). The first report (Thomson 2013) explored the nature of any systematic, general relationship between South Island electricity spot prices and hydro storage. Among other findings, it proposed a framework regression model between suitably chosen transformations of price and storage (PH Model). The second report (Thomson 2014) developed a seasonal regime switching model for South Island hydro storage (SH Model) to better understand the seasonal structure and dynamics of storage and provide a suitable analytical framework for simulating and predicting storage time series. These two reports based their findings on data to 30 September 2012. The current report updates these analyses and models for data to 30 September 2017 (an additional 5 years) and accounts for the significant structural changes that took place in the regulatory frameworks governing the electricity marketplace and its institutions following a Ministerial Review of Electricity Market Performance undertaken in 2009 (MBIE 2009).

Here nominal South Island electricity spot prices were first inflation adjusted using the PPI (Producers Price Index). The real prices that resulted were further corrected for a decreasing trend equivalent to a productivity improvement of 1.7% per annum. Understanding the properties of the weekly time series of trend adjusted real spot prices and hydro storage remains the primary objective, with these key data sets underpinning the PH and SH models fitted. To account for the structural changes initiated by the 2009 Ministerial Review of Electricity Market Performance, a structural break was defined at the end of September 2009. The data before 30 September 2009 was assumed to be representative of the (stable) dynamics and statistical properties of the pre-reform electricity marketplace, whereas the data after 30 September 2009 was assumed to be representative of the post-reform marketplace.

Following Thomson (2013), both spot prices and storage levels were transformed to make them more Gaussian and amenable to regression modelling. The shifted logarithm transformation was applied to the adjusted real spot prices and the Johnson S_B transformation was applied to the storage levels. The latter transformation accounts for the changing shape of the long-run storage distributions by time-of-year and is essentially the logarithm of a storage ratio that measures the amount of storage available as a fraction of that already used. As before, these marginal transformations lead to a more Gaussian bivariate relationship which is exploited using linear regression analysis. Differences between the pre and post 30 September 2009 data are identified. In particular, both adjusted real prices and hydro storage show significant changes in their weekly seasonal patterns post 30 September 2009. However general correlations measuring strength of linear association remain much the same.

While the seasonal switching model used in Thomson (2014) has largely been revalidated, there is clear evidence that the 2009 Ministerial Review of Electricity Market Performance has led to changes in the dynamics of the SH model post 30 September 2009. In particular, the low storage season is less persistent in the post-2009 period with shorter sojourns and the probability of a transition from extreme to intermediate storage is higher in the post-2009 period. These and other results are consistent with greater risk aversion to

low storage in the post-2009 period following the 2009 Ministerial Review of Electricity Market Performance. However the basic structure of the SH model remains fit-for-purpose and provides a simple, yet flexible, stochastic framework within which to examine weekly hydro storage data and better understand its variability.

A preliminary exploration was undertaken of the relationship between spot price and storage within the four storage seasons identified by the SH model. In essence, a switching regression model was fitted between transformed weekly average spot prices and weekly average storage levels. While the fit of the switching regresssion model is reasonable, it is not quite as good as the fit of the PH model. The switching regression residuals also have much stronger residual seasonality than the PH model residuals. This is not unexpected since the switching regression model is based on dynamic storage seasons that are a function of hydro storage alone, whereas the PH model is based on static seasonal patterns that reflect seasonal demand for electricity in addition to seasonal storage and other possible covariates. Despite this limitation, the switching regression model based on storage seasons manages to provide a competitive and informative view of the relationship between price and storage.

Simple modifications to both models are suggested to account for the structural break caused by the 2009 Ministerial Review of Electricity Market Performance and other shortcomings. These and other issues remain topics for further research and development.

Further details are given in the report.

1 Scope of report and terms of reference

In 2013 Statistics Research Associates Ltd (SRA) undertook an exploratory analysis of the relationship between electricity spot price and hydro storage in New Zealand (Thomson 2013). Among other findings, this report proposed and developed a framework regression model between suitable transformations of price and storage (PH model). The PH model is designed to capture the static seasonal dependence relationship between price and storage. In 2014 SRA developed a seasonal regime switching model for South Island hydro storage (Thomson 2014). This model (SH model) is designed to reflect the evolving seasonality and dynamic structure of hydro storage which is episodic in nature due to seasonal rainfall inflows and managed seasonal outflows influenced by demand. Both models were based on weekly data to the end of September 2012 with the PH model based on 13 years of data and the SH model based on almost 16 years of data.

Given that five additional years of data are now available, the New Zealand Electricity Authority (Authority) wishes to update the PH and SH models and, in particular, better account for the structural changes that have taken place since the 2009 Ministerial Review of Electricity Market Performance.

SRA was commissioned to provide the following services.

(a) Update and re-fit the PH model using more recent data provided by the Authority.

(b) Using the PH Model, determine the weekly prices that would be expected given the current storage situation, and in particular, repeat the graphical diagnostic plots given in Thomson (2013).

(c) Explore and document any changes since the original PH model was estimated and, in particular, account for any changes caused by the 2009 Ministerial Review of Electricity Market Performance.

(d) Update and re-fit the SH model using more recent data provided by the Authority.

(e) Explore the relationship between price and storage within the four seasonal states identified by the SH model.

(f) Explore and document any changes since the original SH model was estimated and, if necessary, account for any changes caused by the 2009 Ministerial Review of Electricity Market Performance.

(g) Fully inform key Authority staff about the statistical models and techniques involved, and the statistical computing system R (R Development Core Team, 2004) used for the analysis.

These issues and others are addressed in the sections that follow. In addition, the development R code written for the report has been made available to the Authority.

2 Background

The impact of deregulation and more competitive electricity markets has led to the need for more appropriate models of electricity prices over daily, seasonal and inter-annual time scales. Electricity prices typically vary with time of day, week and year, since they are dependent on local demand, temperature and other variables. They are also highly volatile. This volatility reflects the inelasticity of demand due to the difficulty of substituting for electricity with other forms of energy, and the consequent shape of the marginal cost of supply function which rapidly increases as demand increases. The periodic cycles that are present in electricity demand are also present to some degree in electricity prices. These periodic cycles are a striking feature of New Zealand electricity demand which is dominated by domestic demand (see Bruce et al., 1994, for example). However lack of transportability (New Zealand is geographically isolated) makes price modelling more challenging leading to the need for purpose-built models tailored to the New Zealand market and environment. See Thomson (2013) for a fuller discussion of these issues and a selective literature review.

In New Zealand, electricity generation is dominated by hydro (around 60%), but the hydro catchments have limited total storage capacity of around 15% of annual demand. This is a point of difference between the New Zealand hydro generation system and those of Scandinavia or Canada, for example, where long-term storage is much greater. The lack of storage in New Zealand means that electricity prices are more sensitive to variations in hydro storage than electricity systems dominated by thermal generation or where there is greater long-term hydro storage. While not unexpected, the nature of this relationship and its dependence on time of year (seasonality) are less clear. The PH model (Thomson 2013) aims to capture any systematic, general seasonal relationship between New Zealand electricity spot prices and the levels of hydro storage.

While the PH model accounts for the static seasonal dependence between price and storage, it does not directly model the dynamics of price and storage. In particular, price forecasts or simulations based on the PH model need suitable forecasts or simulations of hydro storage. However inflows to New Zealand hydro reservoirs show stochastic seasons that may arrive early or late (see Harte and Thomson 2004, 2006, 2007), in some cases markedly so, while hydro outflows are managed by electricity generators to meet seasonal demand. The SH model (Thomson 2014) is a stochastic seasonal regime switching model that captures the episodic nature and seasonal evolution of hydro storage. In particular, it provides a dynamic and analytical framework suitable for simulating and forecasting weekly New Zealand hydro storage time series.

The PH model developed in Thomson (2013) is updated in Section 3; the SH model developed in Thomson (2014) is updated in Section 4.

2.1 2009 market reforms

The New Zealand electricity market was introduced in late 1996 and underwent further major structural changes during the period to 1999. Since 1999, the hydro lakes have experienced a number of periods of extremely low storage, some of which (2001, 2003 and 2008) resulted in national conservation campaigns. Following a general election in 2008, the new National government initiated a review of the New Zealand electricity market performance. In April 2009 an Electricity Technical Advisory Group was appointed to review the performance of the electricity market and governance arrangements and to make recommendations on improvements. On 12 August 2009 public feedback was sought on an initial discussion document setting out preliminary recommendations. Following this, the New Zealand government announced the outcome of the review on 9 December 2009. The new measures agreed included, among others, reconfiguration of electricity generators assets, more liquid hedging arrangements and measures to improve security of supply. The latter were designed to increase transparency and create suitable incentives for more conservative management of New Zealand's hydro generation resources. For further details see MBIE (2009)

As will be evident from the analysis that follows, the 2009 electricity market reforms have had a significant impact on both electricity spot prices and hydro storage with post 2009 experiencing less volatile spot prices and fewer periods of extremely low storage. In particular, the futures market for New Zealand electricity hedge contracts has grown rapidly since 2009 to become a liquid, transparent and mature market providing a clearer indication of future spot price movements as well as less volatile spot prices.

2.2 Data

The data provided by the Authority for the analysis comprised South Island (Benmore) daily average electricity spot prices from 1 October 1996 to 30 September 2017, and the total daily storage capacity of the South Island hydro reservoirs at Lakes Tekapo, Pukaki and Ohau from 1 January 1996 to 30 September 2017. Prices are quoted in New Zealand dollars per megawatt-hour ($NZ/MWh) and storage levels are measured in terms of generation potential in terawatt-hours (TWh). The conversion from hydro storage to generation potential and a discussion of the operational constraints on the hydro reservoirs is discussed in Paine and McConchie (2010).

As in Thomson (2013, 2014) the analysis is restricted to South Island prices and the combined storage of Lakes Tekapo, Pukaki and Ohau which will be referred to collectively as the Waitaki storage. The focus on South Island data is partly for expediency (Waitaki storage makes up the bulk of New Zealand's total available hydro storage) and partly because, as noted in Tipping et al. (2004), it minimises any distortions due to the HVDC (High-Voltage Direct Current) link between the South and North Islands. As before, the analysis will be based on weekly average price and storage data rather than the original high frequency daily data. As noted in Thomson (2013), there are two reasons for this. First, since the processed data always has exactly 52 weeks each year, more conventional seasonal analysis techniques are possible. Second, the weekly averages enhance the systematic components of each time series enabling a better understanding of any structural relationships between them.

Weekly averages were constructed from the original daily time series by forming successive weekly averages from the start of each year with 29 February included in the week of 28

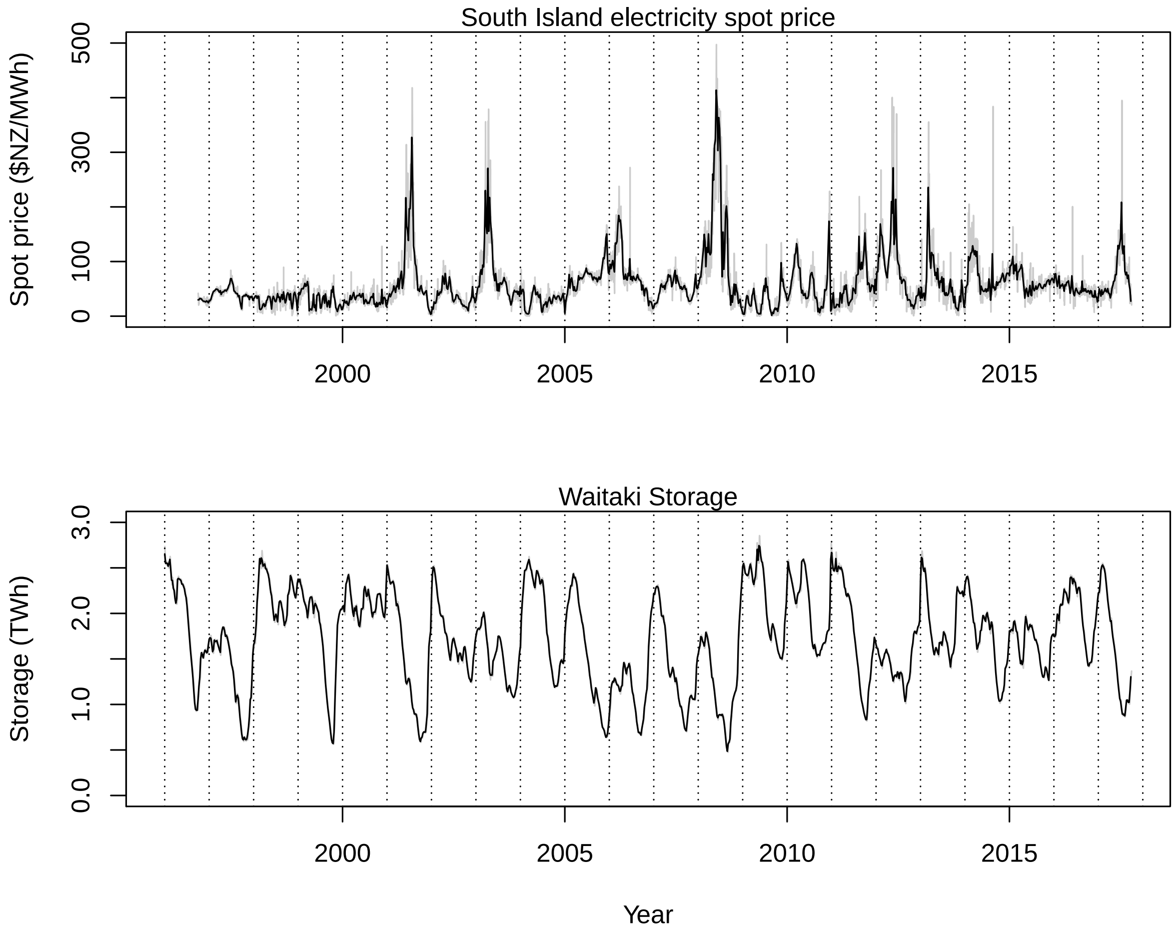

Figure 1: The **upper panel** plots the South Island daily average electricity spot prices (grey) over the period 1 October 1996 to 30 September 2017 with weekly average prices (black) superimposed. The **lower panel** plots Waitaki daily storage levels (grey) over the period 1 January 1996 to 30 September 2017 with weekly average storage levels (black) superimposed.

February (week 9) in leap years, and 31 December included in the week of 30 December (week 52). This yielded 52 weeks for each year with the first week being the average of the first 7 days in January and the last week being the average of the last 8 days of December.

Figure 1 shows the plots of the South Island daily average electricity spot prices and Waitaki daily storage over the period 1 November 1996 to 30 September 2017 with weekly averages superimposed. The weekly average prices closely approximate their daily counterparts with even closer agreement between Waitaki daily and weekly storage. Little information is lost by considering weekly averages, especially if the goal is to better understand the variation of the systematic components of both electricity spot prices and storage.

The nature and scale of the variation in the price series before 2000 is markedly different to that after 2000. This is largely due to the major structural reforms of the New Zealand electricity market that took place in 1996 and over the period to 1999. As a consequence, Thomson (2013) based its analysis on electricity spot prices from 1 October 1999 to 30 September 2012 (13 complete years). The same starting point is adopted for the updated analysis that follows which uses South Island weekly average electricity spot prices over the period 1 October 1999 to 30 September 2017 (18 complete years).

The 2009 Ministerial Review of Electricity Market Performance has also had an impact on prices, although more subtle. Post 2009 spot prices appear to have fewer periods of very high volatility and more periods of quite low volatility. These effects will be further examined and delineated in the sections that follow.

The storage data is much more homogeneous and trend free, but does appear to have fewer very low storage levels post 2009. Although the joint analysis of spot price and storage will use the common time window of 1 October 1999 to 30 September 2017, the analysis of the properties of storage alone will use the data from 1 October 1996 to 30 September 2017. Thomson (2013) analysed data from 1 November 1996 to 30 September 2012 (almost 16 years) whereas the updated analysis that follows analyses 21 complete years of data.

The South Island spot prices shown in Figure 1 appear to show an increasing trend reflecting, in part, the impact of inflation over the 21 years concerned. To account for this effect, weekly average electricity spot prices were inflation adjusted to common (30 September 2017) dollars using the electricity component of the New Zealand Producers Price Index (PPI). This quarterly index is prepared by, and available from, Statistics New Zealand (www.stats.govt.nz) with a weekly version formed using linear interpolation of the logarithms of the index. Thomson (2013) used the New Zealand Consumers Price Index (CPI) rather than the PPI. Although the differences between these two inflation adjusted series are modest, the PPI is the more appropriate measure and has the beneficial effect of down-weighting very large price spikes arising from extremely low storage levels.

Figure 2 plots the nominal and inflation adjusted South Island weekly average electricity spot prices. Robust linear time trends were fitted to these time series using robust regression (M-estimation; see Venables and Ripley, 2002, Chapter 6, Section 6.5) to account for non-Gaussian price variation. The fitted lines show a significant positive slope for the nominal spot prices and a significant negative slope for the inflation adjusted spot prices. To remove the latter, a linear time trend was fitted to the logarithms of the inflation adjusted prices using linear regression. Real South Island weekly average electricity spot prices are now obtained by correcting the inflation adjusted prices for this trend. It is noted that this trend correction corresponds to a productivity improvement of 1.7%. The real spot prices are shown in Figure 2 with a robust linear time trend superimposed. As expected the latter is not significant and the real prices are trend free.

Subsequent analysis will now focus on the 18 complete years of real South Island weekly electricity spot prices from week 40, 1999 to week 39, 2017 (1 October 1999 to 30 September 2017) and 21 complete years of Waitaki weekly storage from week 40, 1996 to week 39, 2017 (1 October 1996 to 30 September 2017).

To examine the impact of the 2009 Ministerial Review of Electricity Market Performance, the data were further partitioned into two periods. The **pre-2009** data comprise weekly prices and storage to 30 September 2009 (up to and including week 39, 2009) and the

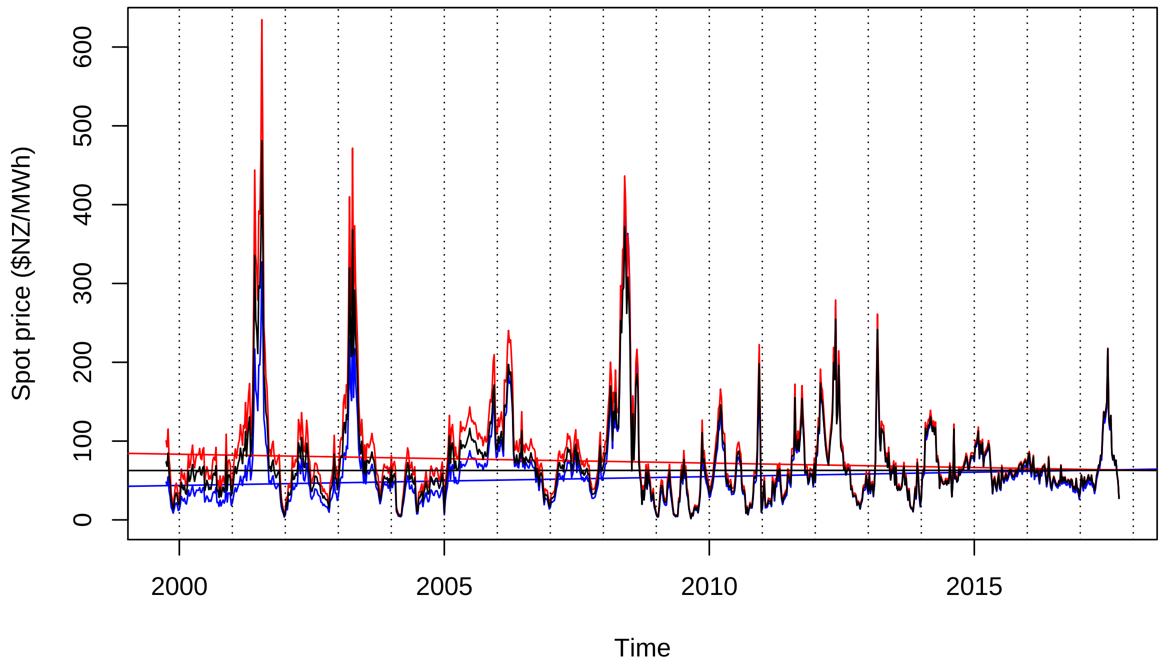

Figure 2: Nominal (blue), inflation-adjusted (red) and real (black) South Island weekly average electricity spot prices over the period 1 October 1999 to 30 September 2017 with robust linear trends superimposed. The inflation-adjusted prices were obtained from the nominal prices using the PPI and the real prices are the inflation-adjusted prices after trend correction.

post-2009 data comprise weekly prices and storage after 30 September 2009 (from week 40, 2009). The split chosen is somewhat arbitrary since some of the outcomes of the review did not come into effect until the following year or even later. A later date such as 30 September 2010 would be closer to the date of the establishment of the Electricity Authority as part of the New Zealand Electricity Industry Act 2010 which enacted the reforms. However, the review outcomes were largely signalled to industry participants before 30 September 2009 and, judging from the plots of the South Island electricity spot prices and Waitaki storage (Figures 1 and 2), had already begun to be factored into the industry's operations. Moreover, the choice of the earlier date of 30 September 2009 provides 8 complete years of seasonal weekly data which, although modest, is sufficient for informative analysis.

The following sections explore these data sets with Section 3 updating the long-run general seasonal statistical relationships (PH model) found in Thomson (2013) and Section 4 updating the seasonal regime switching model (SH model) proposed in Thomson (2014).

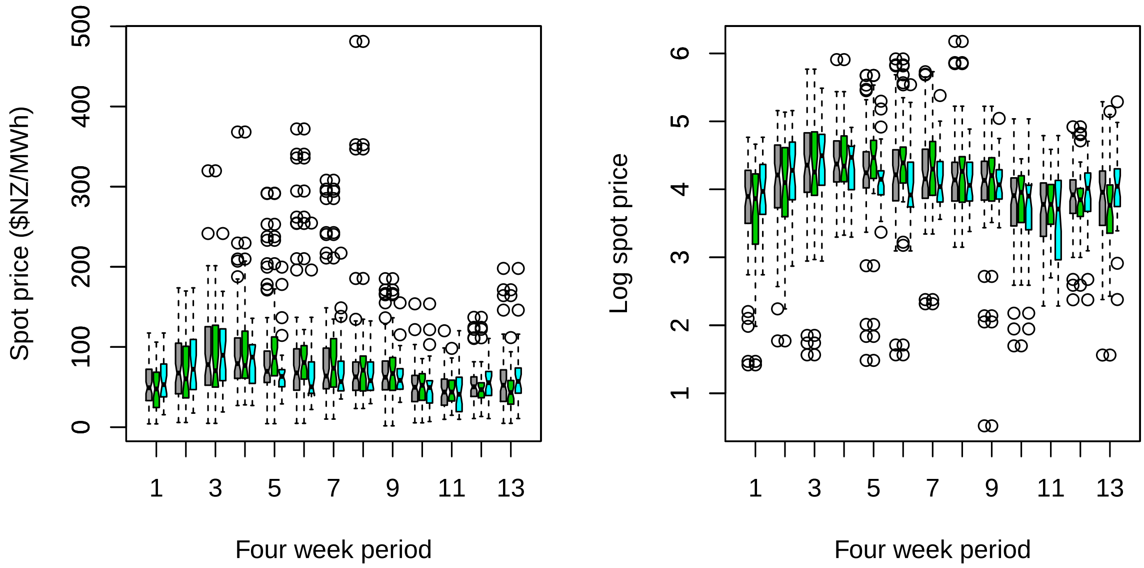

Figure 3: Notched boxplots of real South Island weekly average electricity spot prices (**left panel**) and their logarithms (**right panel**) by four-week period of the year. For each four-week period the component boxplots are for all (grey), pre-2009 (green) and post-2009 (cyan) data.

3 PH model update

To better examine the systematic seasonal components in the data we now block the 52 weeks of the calendar year into 13 consecutive four-week seasonal periods. As noted in Thomson (2013), these periods could have been defined slightly differently (for example, the last week of December could have been included in the first period to better reflect the New Zealand Christmas holiday season) or over finer intervals such as weeks. However the results are unlikely to differ greatly and so we have chosen to maintain the same definitions as before.

The following sections consider the marginal distributions by season of South Island weekly average electricity spot prices (Section 3.1) and also the Waitaki weekly average storage levels (Section 3.2). The joint relationship between weekly spot price and weekly storage is explored in Section 3.3 and conclusions drawn in Section 3.4. Three time periods are considered: **all** data to 30 September 2017 (18 complete years), **pre-2009** data to 30 September 2009 (10 complete years) and **post-2009** data after 30 September 2009 (8 complete years).

3.1 Electricity spot prices

The asymmetric variation of the weekly average, electricity spot prices about their mean levels shown in Figure 2 suggest that a transformation such as the logarithm may well be appropriate. This and any seasonal dependence is further considered in Figure 3 which plots the boxplots of the real prices and their logarithms by four-week period of the year.

Figure 3 shows that the logarithm transformation reduces the positive skewness in the price distributions for each period, making them more symmetric and Gaussian. However they now show some evidence of negative skewness. Although systematic seasonal patterns are difficult to discern in the spot price time series shown in Figures 1 and 2, they are clearly evident in Figure 3 where typical prices, as exemplified by the median, tend to be lower in spring and summer, and higher in autumn and winter. The variation of the price logarithms about their medians (the interquartile range) is also more constant with much less seasonal character. These observations support the use of the logarithm transformation, or similar.

Following Thomson (2013) we consider the shifted logarithm transformation

$$Y_t = \log(P_t - \theta_t) \qquad (P_t > \theta_t) \tag{1}$$

where P_t denotes the weekly average electricity spot price and the threshold parameters θ_t satisfy $\theta_t = \theta_{t+52}$. Note that θ_t represents the lowest possible price for the week of the year corresponding to week t. This transformation is monotonic (preserves order relationships), one-to-one (uniquely maps spot prices to transformed spot prices and vice-versa) and includes the familiar logarithm transformation as a special case ($\theta_t = 0$). In effect, such transformations stretch or contract the shape of an original distribution to make it more or less Gaussian and symmetric.

The threshold parameters θ_t are estimated using the same approach as Thomson (2013) and Harte and Thomson (2006) with each θ_t estimated from all weekly spot prices that fall within a moving time-of-year window of 12 weeks, or 3 consecutive four-week periods. Here each window is centred on the middle four-week period and wraps around the 52 weeks of the year in a circular fashion with week 1 following week 52 etc. The transformed prices Y_t are assumed to follow approximate independent Gaussian distributions with θ_t assumed to be constant over the window, and the means and standard deviations of the transformed spot prices assumed to be constant within each four-week period, but different across four-week periods. Given these assumptions, the θ_t can now be estimated by local maximum likelihood in the manner described in Appendix A.2 of Harte and Thomson (2006) yielding smooth moving estimates of the thresholds θ_t on a four-week time scale. Weekly estimates can be constructed using linear interpolation.

Although these assumptions are at best approximate, they provide a reasonable basis for estimating the θ_t. This procedure could have been based on weeks, rather than four-week periods, to achieve a higher time-of-year resolution. However the lack of data available for any given week (8 years in the case of the post-2009 data) will lead to less robust estimates of the Gaussian means and standard deviations. Working on the coarser time-of-year scale (four-weeks) eliminates this issue at the expense of a lower resolution of the variation of θ_t over the year.

Figure 4 plots estimates of θ_t for the South Island weekly average electricity spot prices (all, pre-2009 and post-2009 data) by week of the year together with week-of-year minimum, median and maximum prices for comparison. The seasonal estimates of θ_t vary about the constant ($\theta_t = \theta$) estimates of -\$18.91 (all data), -\$19.99 (pre-2009 data) and -\$7.01 (post-2009 data). These figures are all negative and not dissimilar to the figure (-\$15.54) given in the earlier study Thomson (2013) which analysed CPI (rather than PPI) adjusted

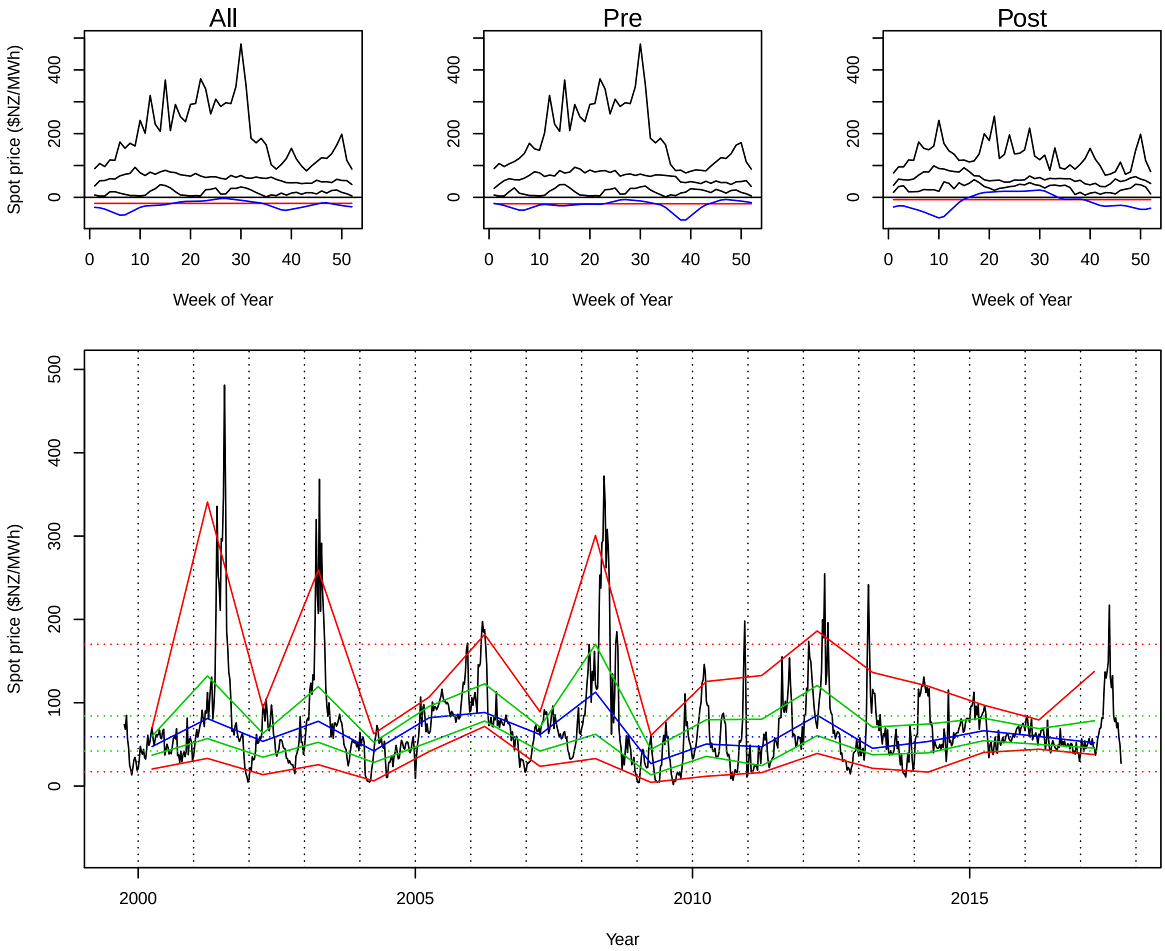

Figure 4: The **upper panels** show zero (black), constant (red) and seasonal (blue) estimates of the threshold parameter θ_t for real South Island weekly average electricity spot prices (all, pre-2009 and post-2009 data) with week-of-year minimum, median and maximum prices superimposed. The **lower panel** shows real South Island weekly average electricity spot prices with 5% and 95% (red), 25% and 75% (green), 50% (blue) annual quantiles and overall quantiles (dotted) superimposed.

weekly spot prices to 30 September 2012. However, while the constant threshold estimate for the pre-2009 data (10 years) is close to that for all data (18 years), it is much closer to zero for the post-2009 data although still negative. This would appear to be due to the dramatic reduction in volatility of the post-2009 weekly average spot prices, particularly extreme prices, about a week-of-year median that is much the same over all, pre-2009 and post-2009 data. These observations are further reinforced by the plot of the weekly average spot prices with annual quantiles (5%, 25%, 50%, 75% and 95%) and overall quantiles superimposed. While the annual medians and quartiles vary about their overall counterparts, the annual extreme quantiles (5th and 95th percentiles) are markedly closer to the median than the overall extreme quantiles. There is evidently a marked reduction in volatility post the 2009 Ministerial Review of Electricity Market performance.

Here the thresholds θ_t have been estimated on a purely statistical basis and need not have

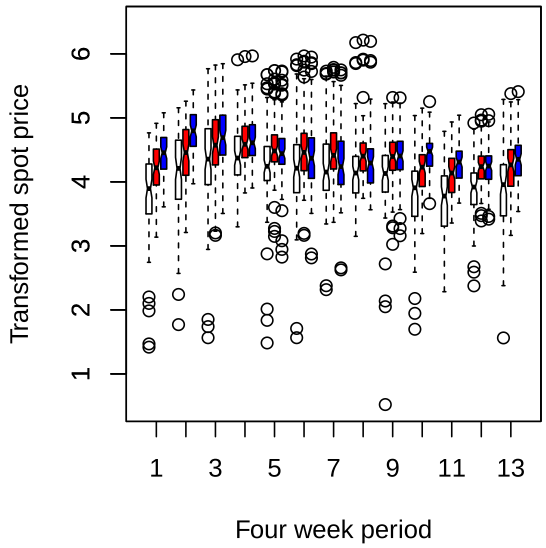

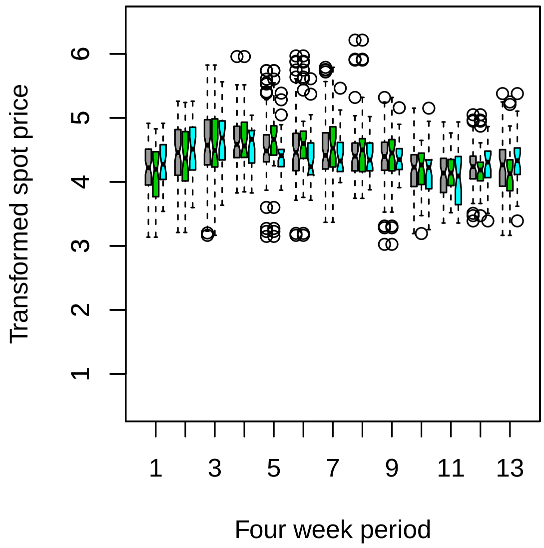

Figure 5: The **left panel** shows notched boxplots of the transformed real South Island weekly average electricity spot prices for all the data by four-week period of the year. For each four-week period the component boxplots are for the shifted logarithm transformation with zero (white), constant (red) and seasonal (blue) thresholds. The **right panel** shows notched boxplots of the transformed real South Island weekly average electricity spot prices with constant threshold by four-week period of the year. For each four-week period the component boxplots are for all (grey), pre-2009 (green) and post-2009 (cyan) data.

any special interpretation. However the estimates of θ_t imply a model in which electricity prices can be negative. As noted in Thomson (2013), this may yet prove to be a feature of the model, rather than a deficiency, since negative commodity prices can sometimes occur (see Fenton et al., 2011, for example).

To aid comparison between the various data sets (all, pre-2009 and post-2009) we now seek one overall transformation (1) and choice of threshold parameters θ_t that makes all data sets as Gaussian and symmetric as possible and therefore more amenable to techniques such as linear regression. A comparison of the impact of the shifted logarithm transformation for all the data by four-week period of the year is given in Figure 5 for zero threshold (equivalent to the logarithm transformation), constant threshold and seasonal threshold. All boxplots show much more symmetry that the original untransformed data (see Figure 3) with those for the constant and seasonal thresholds being more consistently symmetric than those for the logarithm transform (zero threshold). Here the estimated thresholds θ_t are negative so the boxplots for the logarithm transforms will always lie below those for the shifted logarithm transform with constant or seasonal thresholds.

Notched boxplots of the transformed prices Y_t with constant threshold are also shown in Figure 5. By comparison to the logarithm transformation (see Figure 3), the Y_t are less negatively skew and more symmetric, but preserve seasonal variation in location and possibly scale across the year. Judging from the boxplot medians, the differences between the seasonal patterns pre-2009 and post-2009 are generally small with the exception of four-week periods 5 and 6 (late Autumn and early Winter) where the post-2009 medians of the transformed prices are significantly lower than the pre-2009 medians. Note that boxplot notches are designed so that non-overlapping notches indicate that the difference between the respective medians is statistically significant at the 5% level.

As in Thomson (2013) we now suppose that the transformed prices can be modelled as

$$Y_t = \mu_t^Y + \sigma_t^Y V_t \tag{2}$$

where $\mu_t^Y = \mu_{t+52}^Y$, $\sigma_t^Y = \sigma_{t+52}^Y$ are the long-run seasonal mean and standard deviation of Y_t (both periodic with period 52), and the standardised process V_t has mean zero and unit standard deviation. The distribution of the V_t should throw some light on the nature of the distribution of the transformed prices Y_t and, as a consequence, the real weekly spot prices P_t. Simple estimates of μ_t^Y and σ_t^Y are given by the sample means and variances of the transformed prices Y_t by week of the year. These are then smoothed by using a triangular moving average of length 13 which wraps around the 52 weeks of the year in a circular fashion. This simple (seasonal) standardisation process can also be used with other transformations including the logarithm transform and the special case of no transformation.

Figure 6 plots the histograms and normal Q-Q plots of the standardised real, South Island weekly average electricity spot prices, the standardised price logarithms ($\theta_t = 0$), the standardised shifted price logarithms with constant shift ($\theta_t = \theta$), and the standardised shifted price logarithms with seasonal shift θ_t. The normal Q-Q plots graph the sample quantiles of the standardised data against the standard Gaussian quantiles and, for each histogram, a best fitting Gaussian density has been superimposed.

The shifted logarithm transformation with constant or seasonal shift would appear to provide a suitable transformation to Gaussianity for real, weekly average electricity spot prices, in preference to either the real prices or their logarithms. However, even in these cases there is still evidence of lack of fit with the standardised transformed data showing a slightly lighter tail than the Gaussian distribution. It is likely that this is due to the Christmas - New Year holiday period when there is a sharp drop in electricity usage. A more adaptive estimate of θ_t should help to minimise this deficiency. Alternatively, this period could be treated separately from the rest of the year.

Figure 5, and especially Figure 6, show that there is little to pick between the weekly average spot price data transformed by the shifted logarithm transformation with constant or seasonal threshold. For simplicity we adopt the constant (non-seasonal) threshold ($\theta_t = \theta$) in the following sections where θ has been estimated as -\$18.91. This choice of non-seasonal transformation also has the advantage of forcing (2) to model systematic weekly seasonality only through the means μ_t^Y and standard deviations σ_t^Y leading to simpler interpretations and understandings.

3.2 Hydro storage

Here we explore the long-run distributional properties and seasonality of Waitaki weekly average storage levels using data from 30 September 1996 to 30 September 2017 (21 com-

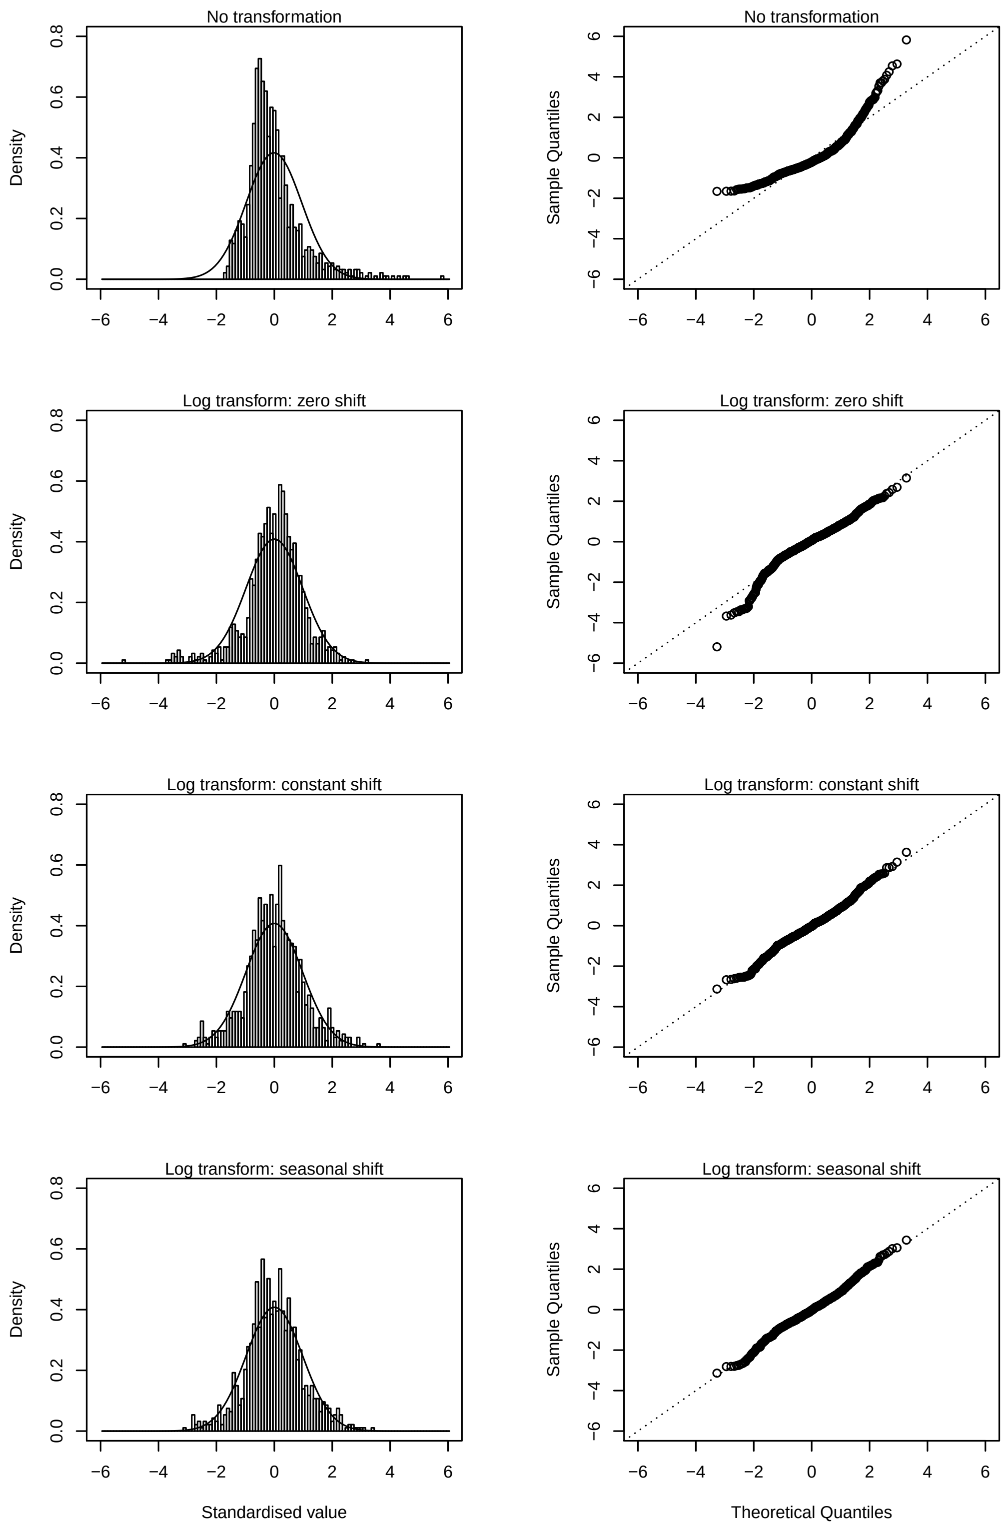

Figure 6: Histograms (left) and normal Q-Q plots (right) of the standardised real South Island weekly average electricity spot prices, the standardised price logarithms, the standardised shifted price logarithms with constant shift, and the standardised shifted price logarithms with seasonal shift. Each histogram has a best fitting Gaussian distribution superimposed.

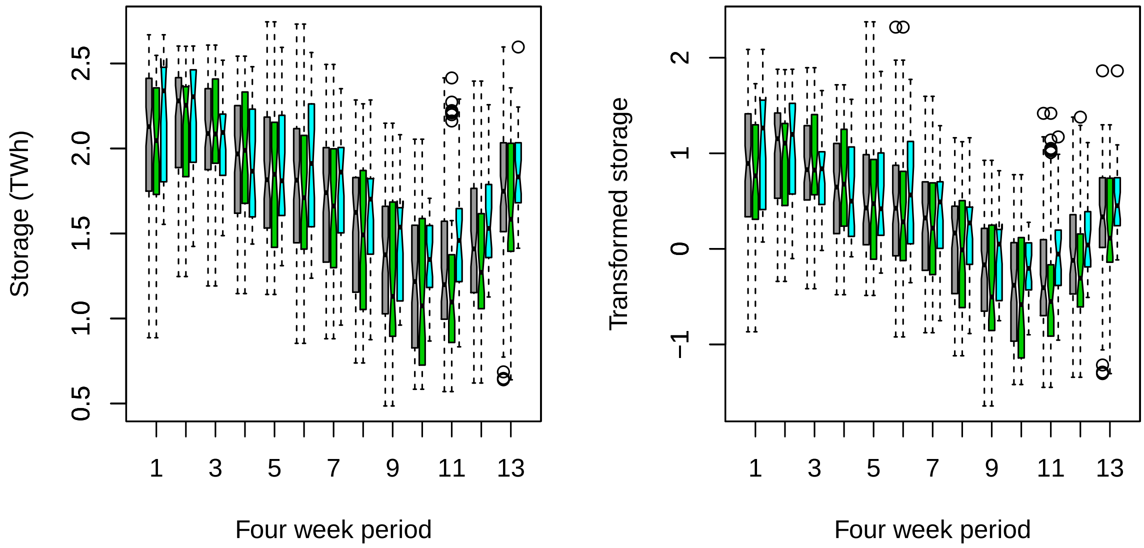

Figure 7: Notched boxplots of Waitaki weekly average storage levels (**left panel**) and their transformed values (**right panel**) by four-week period of the year. The fixed threshold (0,3) Johnson S_B transformation is used. For each four-week period the component boxplots are for all (grey), pre-2009 (green) and post-2009 (cyan) data.

plete years) with the pre-2009 data now covering 13 complete years and the post-2009 data covering 8 complete years as before. Figure 6 shows boxplots of Waitaki weekly average storage levels by four-week period of the year. The seasonal character of this data is clearly evident with weekly storage levels typically lowest in Spring and highest in February. The disposition of the extremes and quartiles relative to the median also suggests that the shape of these distributions depends on time of year with the distributions typically being positively skew when the storage levels are low, and negatively skew when storage levels are high.

As before we seek a transformation that will make the data more symmetric and Gaussian and therefore more amenable to techniques such as linear regression. However New Zealand's hydro reservoirs are managed by electricity generators subject to controls (operating consents) that come into play when storage levels lie outside specified trigger limits. These limits can vary by time of year and, if exceeded, are subject to further operational restrictions including control of long-run recurrence rates. The interaction of these factors and the natural seasonal inflows over time can be complex (see Paine and McConchie, 2010, for details). Nevertheless it is likely that any (statistical) transformation to approximate Gaussianity will reflect both these limits.

Following Thomson (2013), the weekly average storage data are transformed using the Johnson S_B transformation. If H_t denotes the weekly average hydro storage levels then the Johnson S_B transformation is given by

$$X_t = \log(\frac{H_t - \alpha_t}{\beta_t - H_t}) = \operatorname{logit}(\frac{H_t - \alpha_t}{\beta_t - \alpha_t}) \qquad (\alpha_t < H_t < \beta_t) \tag{3}$$

where the lower threshold α_t and upper threshold β_t satisfy $\alpha_t = \alpha_{t+52}$, $\beta_t = \beta_{t+52}$. Here $\text{logit}(x) = \log(x/(1-x))$ is the *logit transformation* defined for $0 < x < 1$. Note that $X_t = \log R_t$ where

$$R_t = \frac{H_t - \alpha_t}{\beta_t - H_t} \tag{4}$$

can be interpreted as the ratio of storage available $(H_t - \alpha_t)$ to that already used $(\beta_t - H_t)$. As a consequence we shall refer to R_t as the *storage ratio*. Furthermore, X_t is a monotonically increasing function (the logit transformation or log odds ratio) of $(H_t - \alpha_t)/(\beta_t - \alpha_t)$ which is the proportion of storage available. These simple relationships show that the Johnson S_B transformation is a direct and interpretable measure of the amount of storage available.

An example of the Johnson S_B transformation is shown in Figure 7 where limits of 0 TWh and 3 TWh ($\alpha_t = 0$, $\beta_t = 3$) have been chosen. These two limits are conservative since storage is always non-negative and, to date, no Waitaki weekly average storage level has ever exceeded 3 TWh (the maximum daily storage level recorded since 1 January 1996 is 2.745 TWh on 18 May 2009). In this particular case there is little difference and the transformation has had only a marginal impact on the shape of the week-of-year distributions.

Using the same estimation strategy as that proposed in Thomson (2013), the thresholds α_t, β_t are estimated from all weekly average storage levels falling within a moving time-of-year window of 3 consecutive four-week periods with each window centred on the middle four-week period. The transformed storage levels X_t are assumed to follow approximate independent Gaussian distributions with α_t, β_t assumed to be constant over the window, and the means and standard deviations of the storage levels assumed to be constant within each four-week period, but different across four-week periods. Given these assumptions, the α_t, β_t can be estimated by local maximum likelihood yielding smooth moving estimates of the thresholds α_t, β_t on a four-week time scale, with weekly estimates constructed using linear interpolation. As before, this estimation strategy should provide reasonable estimates of α_t, β_t although the same caveats apply.

Figure 8 plots the estimates of α_t, β_t for Waitaki weekly average storage levels together with week-of-year minimum, median and maximum weekly average storage levels. In general the seasonal estimates of α_t and β_t vary about the constant thresholds ($\alpha_t = \alpha, \beta_t = \beta$) which are estimated as (0.4, 2.8) TWh (all data), (0.5, 2.8) TWh (pre-2009) and (0.7, 2.7) TWh (post-2009). The size of the post-2009 data (8 complete years) was too small to reliably estimate seasonal thresholds. As expected, the results for all and pre-2009 data are very similar to those given in the previous study (Thomson 2013). Although the post-2009 estimate of the constant upper threshold is much the same as its pre-2009 estimate, the post-2009 estimate of the constant lower threshold differs markedly from its pre-2009 estimate. This is due to the relative lack of low weekly average storage levels in the post-2009 data compared to the pre-2009 data. This is further illustrated in the lower panel of Figure 8 which shows the Waitaki weekly average storage levels with annual quantiles (5%, 25%, 50%, 75% and 95%) and overall quantiles superimposed. While the annual medians and higher quantiles vary about their overall counterparts, the annual lower quantiles (5th and 25th percentiles) are closer to the median that their overall

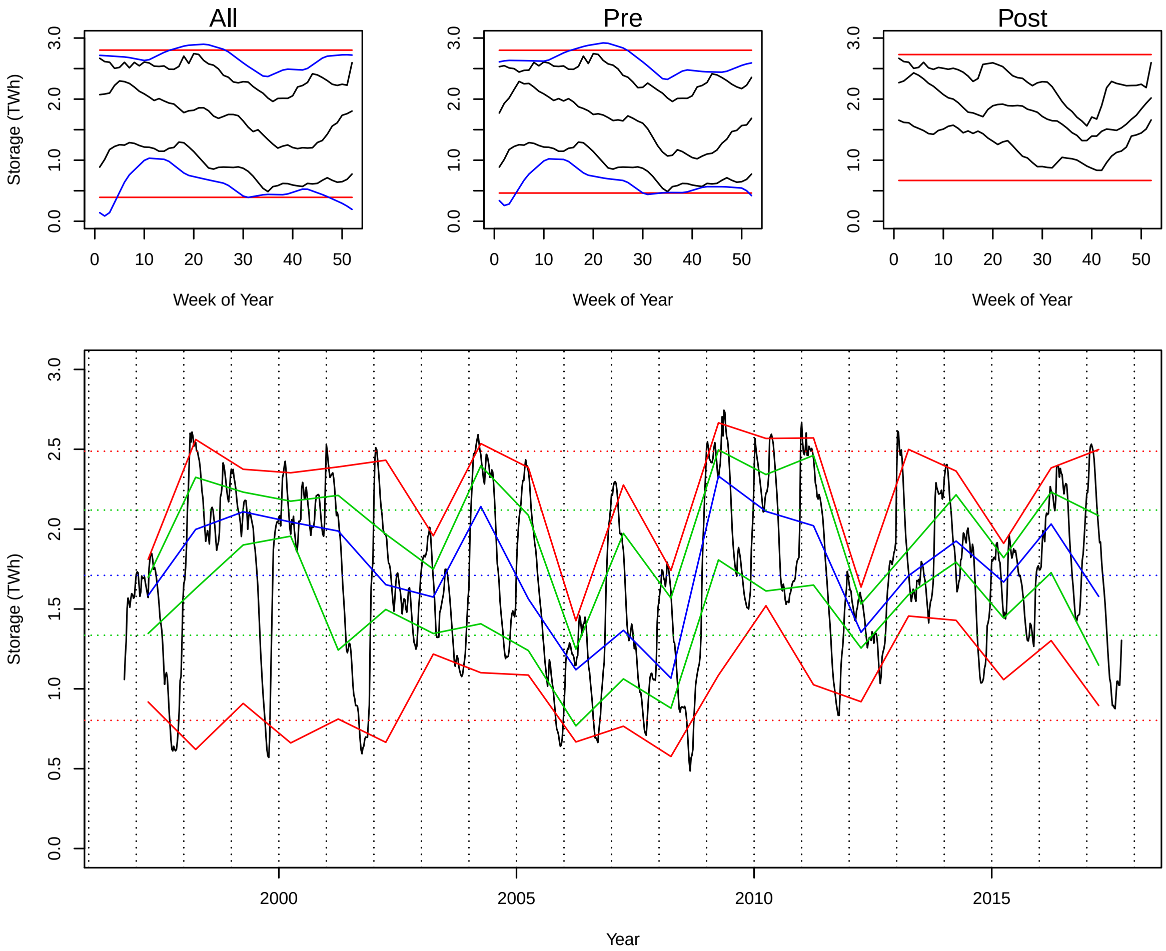

Figure 8: The **upper panels** show constant (red) and seasonal (blue) estimates of the threshold parameters α_t (lower) and β_t (upper) for Waitaki weekly average storage levels (all, pre-2009 and post-2009 data) with week-of-year minimum, median and maximum storage levels superimposed. The post-2009 seasonal threshold estimates are omitted due to insufficient data. The **lower panel** shows Waitaki weekly average storage levels with 5% and 95% (red), 25% and 75% (green), 50% (blue) annual quantiles and overall quantiles (dotted) superimposed.

counterparts. This is particularly so for the lower extreme quantile (5th percentile). The lower tails of the post-2009 seasonal distributions of weekly average storage levels have evidently contracted by comparison to those for the pre-2009 data. This contraction suggests risk aversion and would appear to be mainly the result of changes in hydro storage management post the 2009 Ministerial Review of Electricity Market Performance.

As before, we now seek an overall transformation (3) and choice of threshold parameters α_t*,* β_t *that makes all data sets as Gaussian and symmetric as possible.* This makes for more secure comparisons between the various data sets (all, pre-2009 and post-2009) with the transformed data being more amenable to Gaussian techniques such as linear regression. A comparison of the impact of the Johnson S_B transformation for all the data by four-week period of the year is given in Figure 9 for fixed (0,3) thresholds, constant (non-seasonal) thresholds (α, β) and seasonal thresholds (α_t, β_t). Although the differences

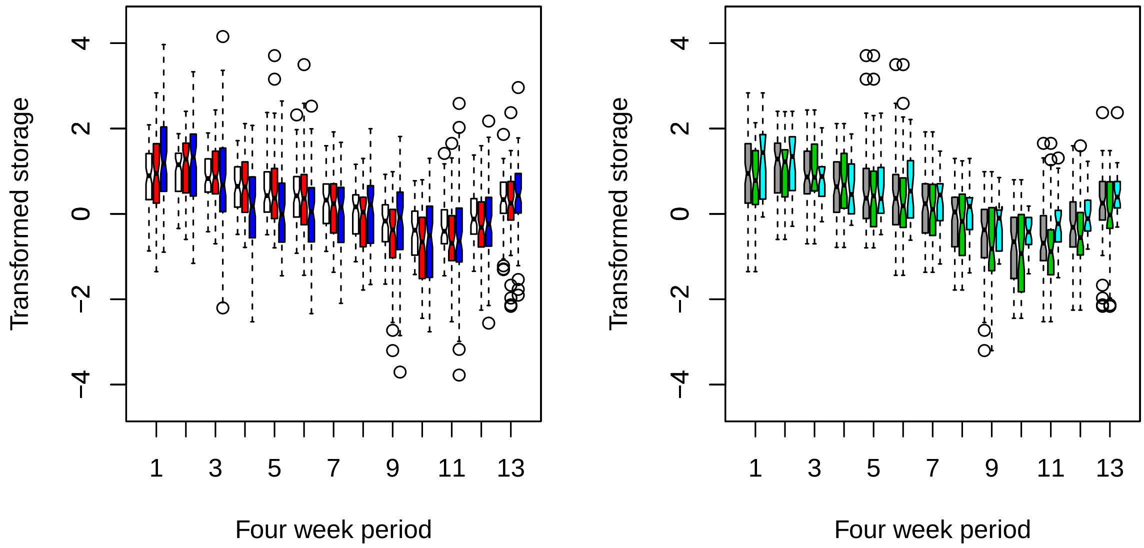

Figure 9: The **left panel** shows notched boxplots of the transformed Waitaki weekly average storage levels for all the data by four-week period of the year. For each four-week period the component boxplots are for the Johnson S_B transformation with fixed (0,3) thresholds (white), constant thresholds (red) and seasonal thresholds (blue). The **right panel** shows notched boxplots of the transformed Waitaki weekly average storage levels with constant thresholds for all the data by four-week period of the year. For each four-week period the component boxplots are for all (grey), pre-2009 (green) and post-2009 (cyan) data.

between the three S_B transformations is modest, they are all consistently more symmetric than the untransformed data shown in Figure 7 and the S_B transformations with constant and seasonal thresholds appear to perform slightly better than the transformation with fixed (0,3) thresholds.

Figure 9 also shows notched boxplots of the transformed prices X_t using the Johnson S_B transformation with constant (non-seasonal) thresholds. By comparison to the transformation with fixed (0,3) thresholds (see Figure 7), the X_t are generally more symmetric while still retaining the dominant seasonal variation across the year. From the boxplot medians and notches, the differences between the seasonal patterns pre-2009 and post-2009 are significantly different for four-week periods 9, 11, 12 and 13 and close to significant for four-week periods 1 and 10. In all these cases the post-2009 medians exceed their pre-2009 counterparts. The seasonal pattern for weekly average storage has evidently changed since 2009. While the variation of post-2009 weekly average storage is much the same as pre-2009 weekly average storage in late summer and autumn (four-week periods 2-5), the post-2009 seasonal medians all exceed those for the pre-2009 data over the rest of the year. This is particularly evident in late winter and spring (four-week periods 9-12). These observations provide further evidence of seasonal change and risk aversion post the 2009 Ministerial Review of Electricity Market Performance.

Now model the transformed weekly average storage levels X_t as

$$X_t = \mu_t^X + \sigma_t^X U_t \tag{5}$$

where $\mu_t^X = \mu_{t+52}^X$, $\sigma_t^X = \sigma_{t+52}^X$ are the long-run seasonal mean and standard deviation of X_t and the standardised process U_t has mean zero and unit standard deviation. Estimates of μ_t^X, σ_t^X are calculated by smoothing the sample means and variances of X_t for each week of the year in the same way as described following (2). These estimates, in turn, yield estimates of U_t whose distribution can be checked for Gaussianity.

Figure 10 plots the histograms and normal Q-Q plots of the standardised Waitaki weekly average storage levels (no transformation) and the standardised transformed Waitaki weekly average storage levels using the Johnson S_B transformation with fixed (0,3) thresholds, constant (non-seasonal) thresholds and seasonal thresholds. The normal Q-Q plots graph the sample quantiles of the standardised data against the standard Gaussian quantiles and, for each histogram, a best fitting Gaussian density has been superimposed. The distribution of the standardised storage levels (no transformation) is platykurtic and possibly bimodal, reflecting the changing shapes of the individual time-of-year distributions. The Johnson S_B transformation with fixed (0,3) thresholds improves the upper tail, but not the lower tail. The Johnson S_B transformations with constant (non-seasonal) and seasonal thresholds are better with the seasonal transformation best. However the differences between the non-seasonal and seasonal transformations is modest.

For simplicity the Johnson S_B transformation with constant or non-seasonal thresholds (α, β) is adopted in the following sections where (α, β) are estimated as (0.4, 2.8) TWh. As in Section 3.1, the non-seasonal transformation has the advantage of forcing (5) to model systematic weekly seasonality only through the means μ_t^X and standard deviations σ_t^X leading to simpler interpretations and understandings.

3.3 Relationship between hydro storage and electricity spot price

The transformations developed in Sections 3.1 and 3.2 have made the marginal distributions of the transformed weekly average spot prices and the transformed weekly average storage levels more Gaussian. While not guaranteed, these transformations should lead to joint distributions that are approximately Gaussian and, as a consequence, more secure correlation and regression analyses. To this end we now consider various scatter plots of transformed price against transformed storage by time of year and overall for all data, pre-2009 data and post-2009 data. Least squares linear regression lines and local regression functions are fitted and assessed.

Figure 11 plots the standardised real South Island weekly average electricity spot prices against the standardised Waitaki weekly average storage levels for no transformation and after transformation. The shifted logarithm transformation was used to transform prices and the Johnson S_B transformation was used to transform storage, both with constant thresholds. As in the previous sections, the standardisation adjusts each variable by a periodic weekly mean and standard deviation where these are calculated by smoothing the sample means and variances of the (transformed) data for each week of the year. In

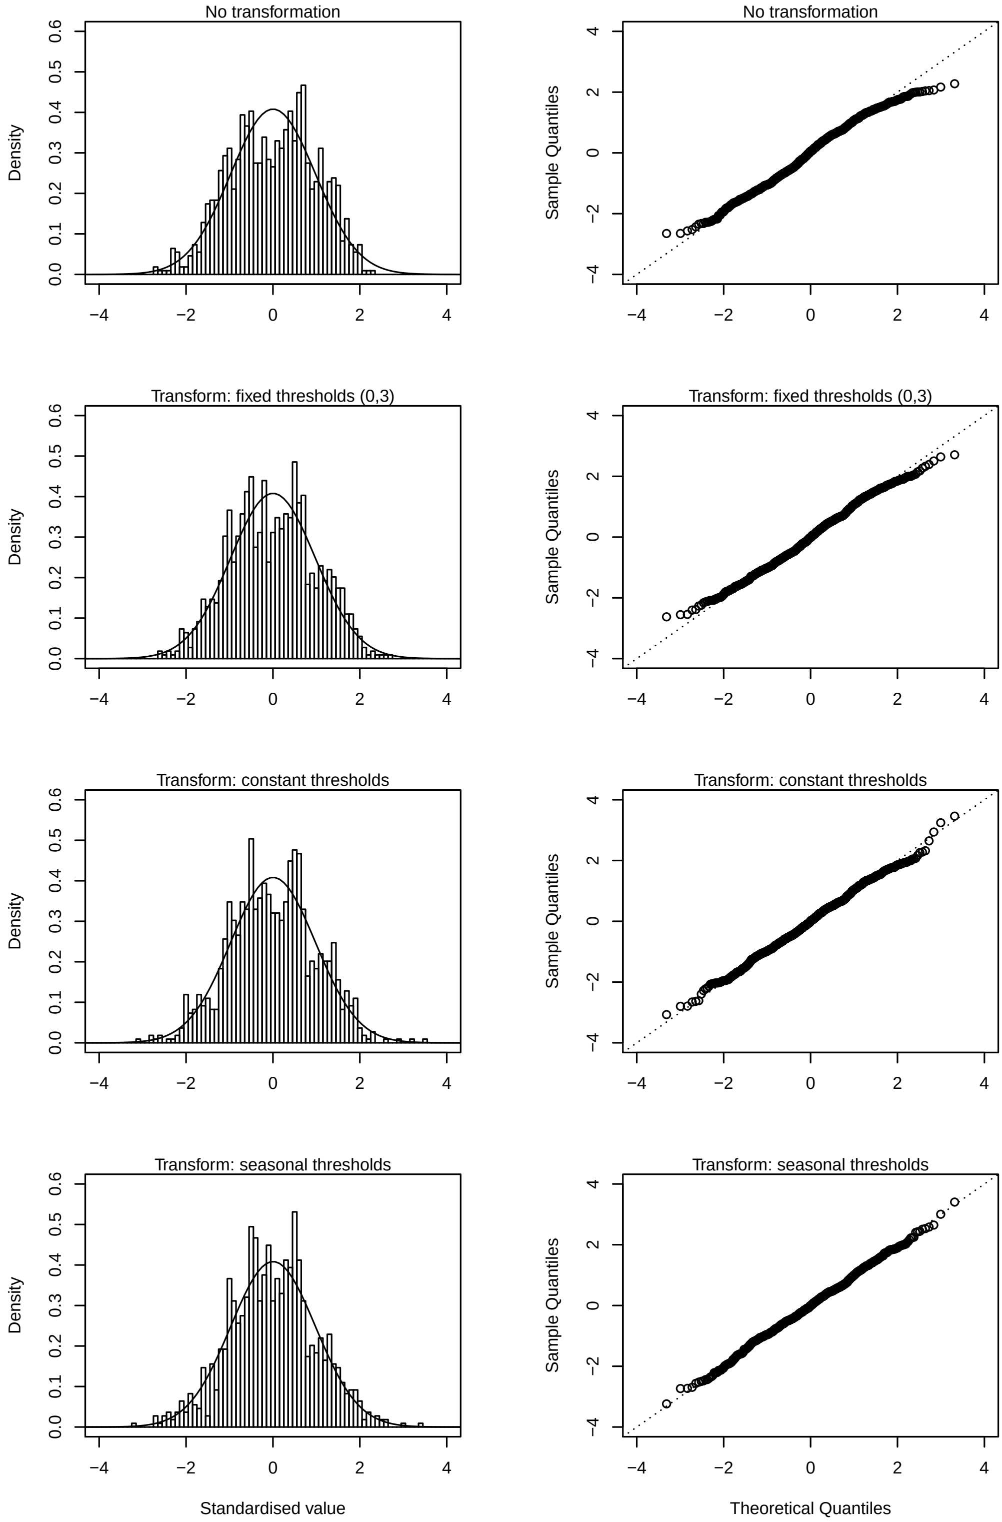

Figure 10: Histograms (left plots) and Gaussian QQ plots (right plots) of standardised Waitaki weekly average storage levels and the standardised transformed Waitaki weekly average storage levels using the Johnson S_B transformation with fixed (0,3), constant (non-seasonal) and seasonal thresholds. Each histogram has a best fitting Gaussian distribution superimposed.

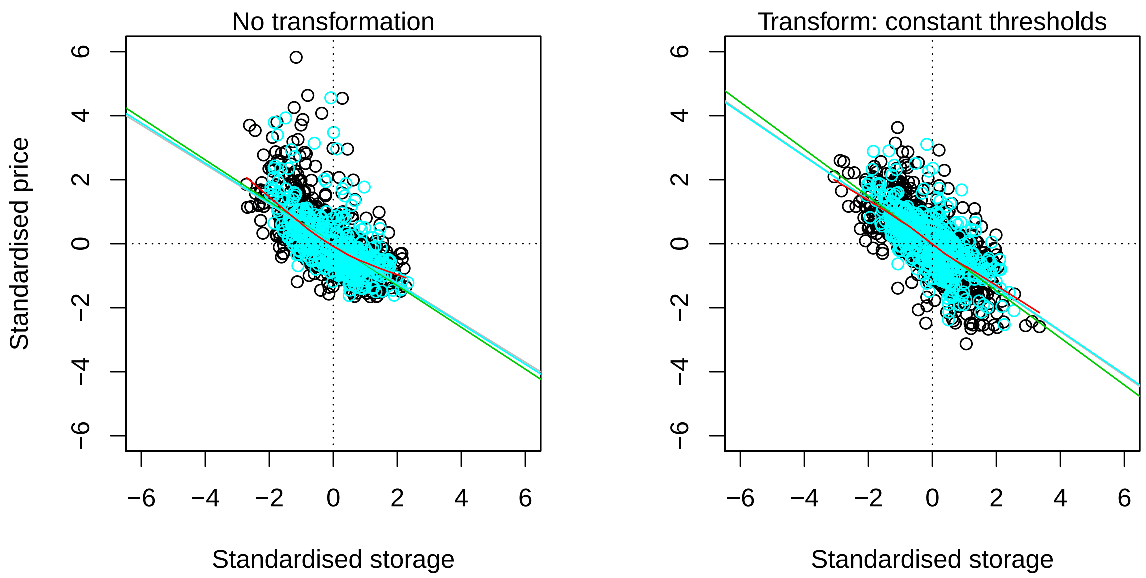

Figure 11: Scatterplots of standardised real South Island weekly average electricity spot prices versus standardised Waitaki weekly average storage levels for no transformation and after transformation (shifted logarithm transformation for prices and Johnson S_B transformation for storage, both with constant thresholds). All data points are shown with post-2009 data points highlighted (cyan). Least squares regression lines for all (grey), pre-2009 (green) and post-2009 (cyan) data are superimposed as is a *loess* local regression functions (red) for all the data.

particular, for the transformed data (X_t, Y_t), Figure 11 plots estimates of the standardised values (U_t, V_t) and their regression lines.

As expected, the scatter plot for the standardised transformed variables looks more Gaussian (elliptical clustering) than that for the untransformed data. Adaptive local regression functions are fitted to both scatter plots (all data) using *loess* (*lo*cal regr*ess*ion; see Cleveland et al., 1992) where these functions are nonparametric estimates of the true, possibly non-linear, regression relationship between (transformed) price and storage. Since a linear regression relationship is further evidence of Gaussianity, the standardised transformed variables again look much more Gaussian compared to the case of no transformation, where the regression function shows marked non-linearity. In the case of the standardised transformed variables, the linear relationships shown (all, pre-2009 and post2009 data) are in reasonable agreement with the *loess* curve. Note that the slopes of the best fitting regression lines for standardised variables are direct estimates of the correlations between the two variables for the three data sets (all, pre-2009 and post-2009).

Table 1 gives the estimated slopes (correlations) and R-squared values for the best fitting regression lines of standardised price versus standardised storage before and after transformation. In all cases the linear relationship is strongest after transformation and the differences between slopes is not statistically significant. As might be expected, the values for the pre-2009 data are in close agreement with those found in Thomson (2013) for data to 30 September 2012.

No transform			Transform		
All	Pre-2009	Post-2009	All	Pre-2009	Post-2009
-0.62 (0.39)	-0.65 (0.46)	-0.63 (0.40)	-0.69 (0.47)	-0.74 (0.54)	-0.68 (0.46)

Table 1: Slopes (correlations) and R-squared values (in brackets) for the best fitting regression lines of the standardised real South Island weekly average electricity spot prices versus standardised Waitaki weekly average storage levels for no transformation and after transformation (shifted logarithm for prices and Johnson S_B transformation for storage, both with constant thresholds).

The regression relationship between the standardised transformed spot prices and the standardised transformed storage levels may be different at different times of the year. To check whether this is the case, Figure 12 shows the same plots as Figure 11 but over the standard seasons of the year where Spring corresponds to September, October, November (weeks 36–48), Summer corresponds to December, January, February (weeks 49–52 and 1–9), Autumn corresponds to March, April, May (weeks 10–22) and Winter corresponds to June, July, August (weeks 23–35). In general the scatter plots for the standardised transformed variables are more Gaussian (elliptical clusters) than those for the untransformed data, and the corresponding linear regression lines are closer to the actual regression relationships estimated by *loess*. A summary of the regression results is given in Table 2.

	No transform			Transform		
	All	Pre-2009	Post-2009	All	Pre-2009	Post-2009
Spring	-0.52 (0.32)	-0.48 (0.31)	-0.68 (0.50)	-0.57 (0.34)	-0.56 (0.35)	-0.74 (0.52)
Summer	-0.68 (0.46)	-0.79 (0.65)	-0.65 (0.38)	-0.69 (0.45)	-0.80 (0.60)	-0.73 (0.48)
Autumn	-0.64 (0.39)	-0.70 (0.49)	-0.52 (0.26)	-0.75 (0.56)	-0.82 (0.68)	-0.57 (0.34)
Winter	-0.63 (0.41)	-0.63 (0.40)	-0.66 (0.47)	-0.74 (0.53)	-0.76 (0.54)	-0.69 (0.50)
All	-0.62 (0.39)	-0.65 (0.46)	-0.63 (0.40)	-0.69 (0.47)	-0.74 (0.54)	-0.68 (0.46)

Table 2: Slopes (correlations) and R-squared values (in brackets) by season for the best fitting regression lines of the standardised real South Island weekly average electricity spot prices versus standardised Waitaki weekly average storage levels for no transformation and after transformation (shifted logarithm transformation for prices and Johnson S_B transformation for storage, both with constant thresholds).

As before, Table 2 shows that the linear relationship is always strongest after transformation and the results for the pre-2009 data are in good agreement with those given in Thomson (2013). For the most part, differences between pre-2009 and post-2009 slopes within seasons are not statistically different. The exceptions are Autumn and Spring. In Autumn the pre-2009 and post-2009 slopes (correlations) for the transformed data are significantly different (1% level) with a weaker relationship post-2009. The reverse is true in Spring which has a stronger relationship post-2009 although the difference is only just significant at the 6% level. The two shoulder seasons (Spring and Autumn) appear to have swapped roles pre-2009 and post-2009.

For each transformed data set (all, pre-2009 and post-2009), the pairwise differences in slopes across seasons are generally not significantly different with the exception of all and

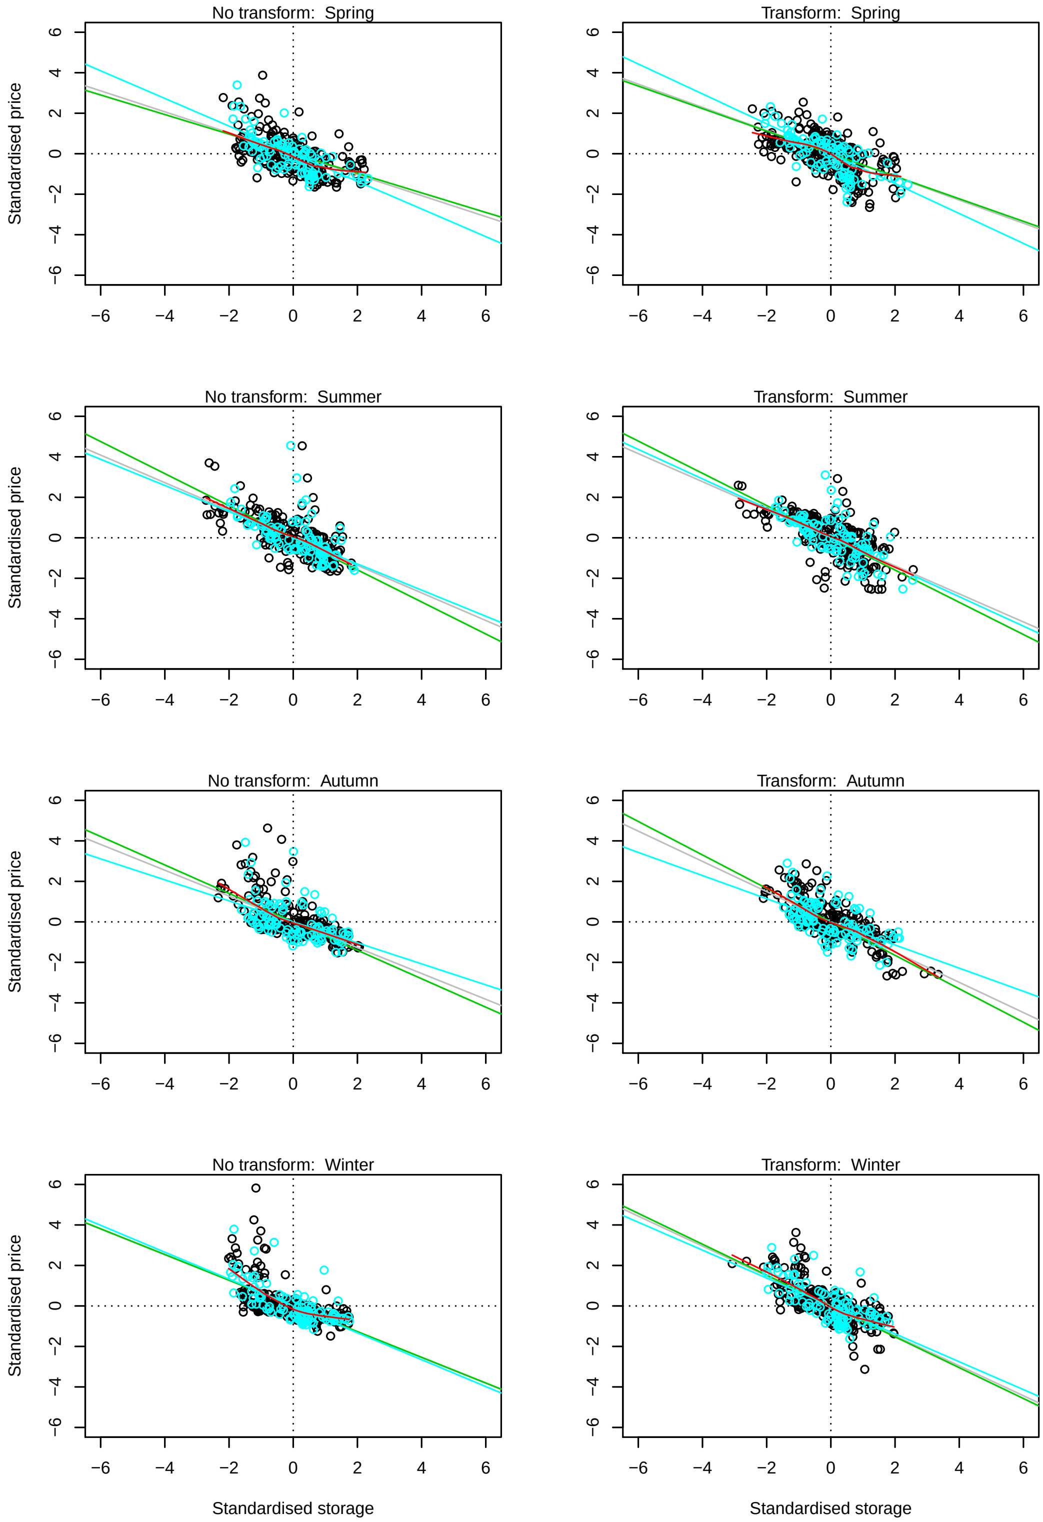

Figure 12: Scatterplots by season of the standardised real South Island weekly average electricity spot prices versus standardised Waitaki weekly average storage levels for no transformation and after transformation (shifted logarithm for prices and Johnson S_B transformation for storage, both with constant thresholds). All data points are shown with post-2009 data points highlighted (cyan). Least squares regression lines for all (grey), pre-2009 (green) and post-2009 (cyan) data are superimposed as is a *loess* local regression functions (red) for all the data.

pre-2009 data sets where Spring differs significantly from the other seasons. In particular, there appear to be no significant differences between slopes (correlations) across seasons post-2009. These findings suggest that standardised transformed spot prices and standardised transformed storage levels are approximately bivariate Gaussian with constant correlation for the most part, especially post-2009.

As in Thomson (2013), these results and those of Sections 3.1 and 3.2 support a general regression model of the form

$$Y_t = \log(P_t - \theta_t) = a_t + b_t X_t + \epsilon_t \tag{6}$$

where transformed spot price Y_t has seasonal price thresholds $\theta_t = \theta_{t+52}$ and transformed storage $X_t = \log R_t$ has seasonal storage thresholds $\alpha_t = \alpha_{t+52}$, $\beta_t = \beta_{t+52}$ with R_t the storage ratio (4). The residual error process ϵ_t has zero mean and seasonal standard deviation $\sigma_t = \sigma_{t+52}$. From (2) and (5) the seasonal regression coefficients $a_t = a_{t+52}$, $b_t = b_{t+52}$ satisfy

$$a_t = \mu_t^Y - b_t \mu_t^X, \qquad b_t = \rho_t \frac{\sigma_t^Y}{\sigma_t^X}$$

where $\rho_t = \rho_{t+52}$ is the seasonal correlation between the standardised spot price V_t and standardised storage levels U_t. *For the remainder of this section we focus on the important special case of constant (non-seasonal) transformation thresholds ($\theta_t = \theta$, $\alpha_t = \alpha$, $\beta_t = \beta$) and constant (non-seasonal) correlation ($\rho_t = \rho$).*

Figure 13 gives a view of the quality and accuracy of the constant correlation ($\rho_t = \rho$), seasonal regression of transformed real South Island weekly average electricity spot prices Y_t against transformed Waitaki weekly average storage levels X_t for all data, pre-2009 data and post-2009 data. The root-mean-squared-error (RMSE) plots of the regression residuals indicate that, in each case (all, pre-2009 and post-2009), regression on transformed storage provides a significant improvement on no regression (predicting transformed prices with just their seasonal means). While this improvement is greatest for the pre-2009 data (around 66% reduction with the regression explaining around 34%), the improvement for the post-2009 data (around 73% reduction with the regression explaining around 27%) is still worthwhile and, in particular, leads to the lowest RMSE overall. Moreover, the RMSE of the post 30 September 2009 residuals for regression using all data is always greater than the post-2009 RMSE, with the exception of winter (weeks 23 - 35) when the differences are marginal and accuracy is best. Note that the post-2009 RMSE (an estimate of σ_t) is also more constant and less variable than the pre-2009 RMSE supporting the case for constant (non-seasonal) standard deviation ($\sigma_t = \sigma$) for the residual error process ϵ_t. The change in the seasonal pattern for storage post the 2009 Ministerial Review of Electricity Market Performance evidently needs to be accounted for.

Notched boxplots of the regression residuals by four-week period of the year are also shown in Figure 13. These show little, if any, seasonality and vary about a zero mean. Indeed, the boxplot notches indicate that there are no significant differences between pre-2009 and post-2009 median residuals and almost all notch intervals include zero, especially for the post-2009 data. For each data set (all, pre-2009, post-2009), the boxplots appear to have approximately constant (non seasonal) standard deviation with the exception of the regression residuals for the pre-2009 data where the interquartile ranges suggest seasonal

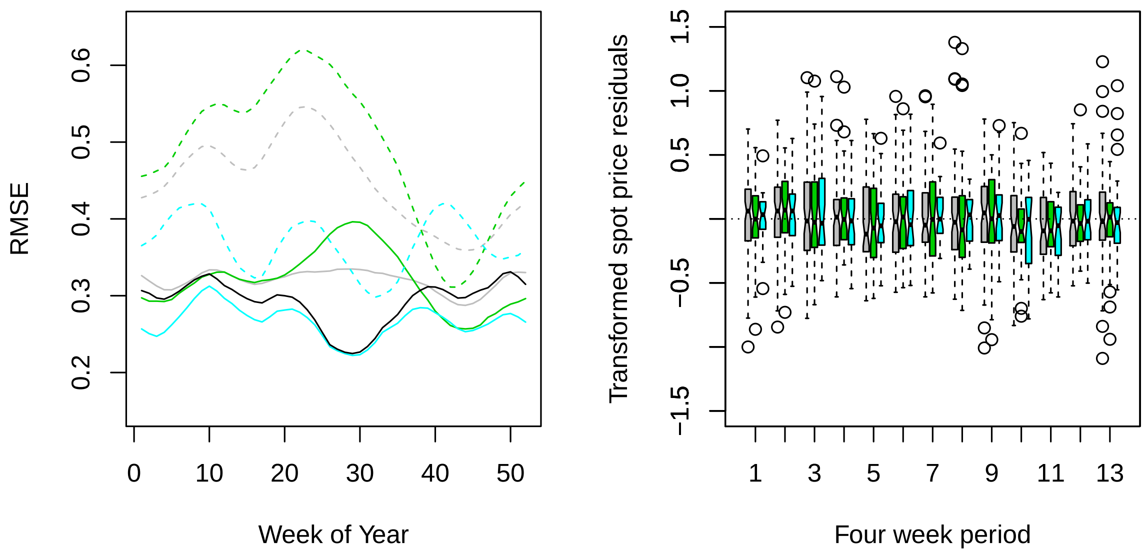

Figure 13: Plots of the residuals and their RMSE for the constant correlation ($\rho_t = \rho$), seasonal regression of transformed real South Island weekly average electricity spot prices Y_t on transformed Waitaki weekly average storage levels X_t for all (grey), pre-2009 (green) and post-2009 (cyan) data. The transformations are the shifted logarithm for prices and the Johnson S_B transformation for storage, both with constant thresholds. The **left panel** shows RMSE estimates of the residuals by week of the year for no regression (dotted), after regression (solid) and for the regression residuals (all data) post 30 September 2009 (black). Notched boxplots of the regression residuals by four-week period of the year are shown in the **right panel**.

variation in standard deviation across the year. As expected, the general scale order of the boxplot interquartile ranges is also consistent with the RMSE plots shown in the left panel of Figure 13 with pre-2009 residuals generally greater in magnitude than those post-2009.

Plots of the fitted values and residuals are shown in Figure 14 for the seasonal regression of transformed real South Island weekly average electricity spot prices against transformed Waitaki weekly average storage levels. The fit to the transformed prices is reasonable, but there are still periods when the regression relationship persistently over or under estimates the transformed price. As a consequence the residuals exhibit strong positive autocorrelation (their lag one autocorrelation is around 0.8) and show evidence of time-varying (evolving) volatility. These results replicate the findings of Thomson (2013) although the post-2009 fits and RMSE of the residuals are slightly better.

As in Thomson (2013), the model (6) can be written in terms of the original weekly average spot prices P_t as

$$P_t = \theta_t + c_t R_t^{b_t} e_t \tag{7}$$

where $c_t = \exp a_t$ and $e_t = \exp \epsilon_t$ is now multiplicative error. In particular, the conditional mean of P_t given R_t is

$$E(P_t|R_t) = \theta_t + c_t R_t^{b_t} e^{0.5\sigma_t^2}$$

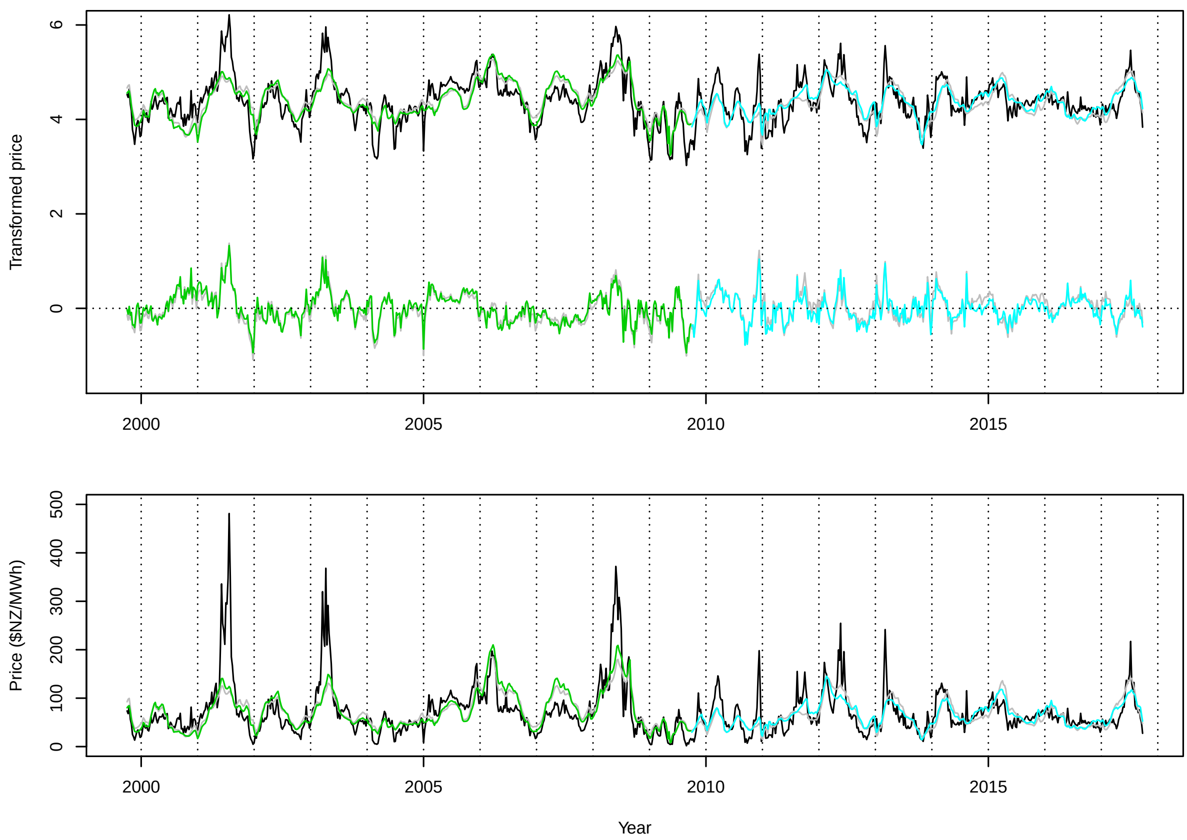

Figure 14: Plots of the real South Island weekly average electricity spot prices (bottom panel) and their transforms (top panel) together with fitted values and residuals (top panel) from the constant correlation ($\rho_t = \rho$), seasonal regression of transformed price against transformed Waitaki weekly average storage for all (grey), pre-2009 (green) and post-2009 (cyan) data. The transformations are the shifted logarithm for prices and the Johnson S_B transformation for storage, both with constant thresholds.

and an estimate of this quantity is plotted in the lower panel of Figure 14. In terms of the original weekly average spot prices, the agreement of the fitted regression is, as expected, much the same as for the transformed prices, but with large positive deviations amplified and large negative deviations compressed. The periods where the regressions perform worst typically occur before 30 September 2009 (the 2001 winter and 2003 autumn are obvious examples) and the quality of the fits has generally improved post 30 September 2009, especially over more recent years. However these static regression models do not capture the dynamic structure of the residuals whose persistence (positive autocorrelation) reflects other stochastic variation such as evolving seasonality.

3.4 Summary

The findings in Sections 3.1, 3.2 and 3.3 largely confirm the general statistical framework proposed in Thomson (2013) where weekly average spot prices P_t depend on weekly average hydro storage levels H_t through the model

$$\log(P_t - \theta_t) = a_t + b_t \log R_t + \epsilon_t \tag{8}$$

and R_t is the storage ratio

$$R_t = \frac{H_t - \alpha_t}{\beta_t - H_t}.$$

The thresholds θ_t, α_t, β_t and regression coefficients a_t, b_t are periodic functions with period 52 weeks and the residual component ϵ_t is a zero mean stochastic process that captures the non-systematic and dynamic components of this general relationship.

However, this model now needs to account for a structural break in seasonality following the 2009 Ministerial Review of Electricity Market Performance. In particular, real prices and especially hydro storage show significant changes in their long-term seasonal mean levels post 30 September 2009 compared to those before. In Autumn and early winter post-2009 mean real prices are lower than those pre-2009, but otherwise pre-2009 and post-2009 seasonal means are much the same. For hydro storage, the post-2009 seasonal mean levels are much the same as pre-2009 in late summer and autumn, but exceed those for the pre-2009 data over the rest of the year. The latter is particularly evident in late winter and spring.

As in Thomson (2013), the shifted logarithm transformation for prices and Johnson S_B transformation for storage, both with constant (non-seasonal) thresholds ($\theta_t = \theta$, $\alpha_t = \alpha$, $\beta_t = \beta$), make the respective marginal distributions more Gaussian and the joint distribution more amenable to linear regression modelling. After transformation correlations measuring strength of linear association between standardised transformed prices and standardised transformed storage were found to be much the same as those reported in Thomson (2013). In particular, this correlation looks to be constant (non-seasonal), at least as a first approxination.

These exploratory results update and generally confirm those of Thomson (2013). Further detailed analysis is needed to refine the general framework (8) and, in particular, incorporate a structural break to account for the impact of the 2009 Ministerial Review of Electricity Market Performance. The nature of the non-systematic error component ϵ_t and its dynamic structure need to be better determined and suitable stochastic models developed. Nevertheless, the systematic general relationship (8) appears to provide a relatively simple, readily interpretable, framework in which to embed stochastic dynamic models for electricity spot prices influenced by seasonal hydro storage levels whose variation is subject to operational capacity constraints.

4 SH model update

The static seasonal regression analysis carried out in Section 3 does not properly account for the dynamic structure of either price or storage. In particular, it is likely that evolving seasonality, in one form or another, will be present in the storage data and, in turn, be reflected in prices. Evolving seasonality can occur in may ways, from changing smoothly over years to exhibiting more abrupt, episodic behaviour with seasons starting earlier or later than expected and varying in length from year to year. The latter describe stochastic seasons, or seasonal regimes, which are in direct contrast to the conventional, three month, deterministic seasons (December, January, February denoting summer, March, April, May denoting Autumn, etc).

Thomson (2014) developed and fitted a non-homogeneous hidden Markov model (NHMM) for Waitaki weekly average storage over the period 1 November 1996 to 30 September 2012. This seasonal switching model was informed by previous studies of New Zealand weekly hydro catchment inflows (see Harte and Thomson, 2007, for example) and builds on a hidden seasonal switching model for daily rainfall developed by Carey-Smith, Sansom and Thomson (2014). Instead of the standard fixed seasons, this model allows the seasons to occur earlier or later than expected and have varying duration while maintaining the usual seasonal precession. The model dynamically classifies weekly storage into seasons whose onsets vary from year to year and within which the model parameters are assumed to be time homogeneous.

Following Thomson (2014), weekly storage X_t is assumed to follow a hidden Markov switching model of the form

$$X_t = \mu_{S_t} + \sigma_{S_t} Z_t \qquad (t = 1, 2, \ldots) \tag{9}$$

where the states S_t form an unobserved Markov chain that takes on the values $1, \ldots, 4$. If S_t is known, the conditional mean and variance of X_t are given by

$$E(X_t|S_t) = \mu_{S_t}, \qquad \mathrm{Var}(X_t|S_t) = \sigma^2_{S_t}.$$

so that the mean level and standard deviation of X_t switch between the pairs of values $(\mu_1, \sigma_1), \ldots, (\mu_4, \sigma_4)$ according to the state or regime specified by S_t. The time series Z_t is assumed to be a zero-mean stationary Gaussian process that is independent of S_t and has unit variance. It is modelled as Gaussian white noise or as a low order autoregressive moving-average (ARMA) process, depending on the serial correlation present within regimes.

The model for S_t is specified by two simple 2-state Markov chains C_t and V_t which each take on the values 0 and 1 with the mapping between S_t and C_t, V_t given by Table 3. Thomson (2014) used C_t to model low ($C_t = 0$) and high ($C_t = 1$) storage seasons that have stochastic onsets and durations, but otherwise occur on an annual basis as expected. The Markov chain V_t describes a secondary storage state. Within each stochastic season C_t, weekly hydro storage is assumed to follow a conventional HMM switching model

S_t	C_t	V_t	μ_{S_t}	σ_{S_t}
1	0	0	μ_1	σ_1
2	0	1	μ_2	σ_2
3	1	0	μ_3	σ_3
4	1	1	μ_4	σ_4

Table 3: Mapping of S_t state labels to those for C_t and V_t.

with $V_t = 0$ modelling variation about a normal or intermediate hydro storage level, and $V_t = 1$ modelling variation about a more extreme level (a higher level in the case of the high storage season and a lower level in the case of the low storage season).

Simple procedures for fitting and predicting the model are developed in Thomson (2014) based on a mix of statistical theory, such as maximum likelihood, and more empirical procedures. The seasonal switching model is readily simulated which allows a variety of simulation-based methods to be considered, for estimation and prediction. These procedures, collectively, allow for improved risk forecasting and a better understanding of the seasonal dynamics of New Zealand weekly hydro storage, particularly when storage is low.

In the following sections, the model (9) is fitted to Waitaki weekly average storage over the period 1 October 1996 to 30 September 2017 using the same analysis and methods as Thomson (2014). This data set contains exactly 21 years of weekly data compared to the almost 16 years used in Thomson (2014). It is noted that, over the time points in common, the two data sets are not quite identical with the new measurements typically greater than the old measurements, although the differences are very small with a mean difference of 0.05 TWh.

Here, as in Thomson (2014), the analysis is confined to the untransformed Waitaki, weekly average, storage levels ($X_t = H_t$) since any skewness in the data is likely to be reasonably well explained by the mixture distributions inherent in HMMs such as (9). Using all the data, a fitted model is determined that is dominated by the pre-2009 data and which is close to the model fitted in Thomson (2014). This model is then used to classify seasonal states and better understand the impact of the 2009 Ministerial Review of Electricity Market Performance. Have the dynamics of the seasonal states changed post 30 September 2009 and, if so, how? Are price-storage relationships within states much the same for pre-2009 and post-2009 data? These and related issues are explored in the following sections.

Following Thomson (2014), Section 4.1 explores the structure of the data using a simple non-seasonal HMM and, in the light of this, the seasonal switching model (9) is fitted in Section 4.2. Section 4.3 explores the relationship between price and storage within seasonal states and Section 4.4 draws conclusions.

4.1 Exploratory analysis with non-seasonal switching model

Here we model Waitaki weekly average storage X_t by the non-seasonal hidden Markov switching model (9) where C_t, V_t are now independent homogeneous 2-state Markov chains specifying S_t through Table 3. Although non-seasonal, this model is sufficiently flexible to be a useful exploratory tool for determining the nature of the seasonal regime structure within the historical data. The model was fitted by maximum likelihood using the Expectation Maximisation (EM) algorithm and the choice of model was guided by the Akaike Information Criterion (AIC). This trades model fit against model complexity by selecting the model order p that minimises

$$AIC = -2 \text{ maximised log-likelihood} + 2p$$

	Old		New	
S_t	$\hat{\mu}_{S_t}$	$\hat{\sigma}_{S_t}$	$\hat{\mu}_{S_t}$	$\hat{\sigma}_{S_t}$
1	1.24	0.14	1.61	0.11
2	0.80	0.14	1.08	0.24
3	1.68	0.14	1.96	0.11
4	2.24	0.20	2.36	0.14

Old

C_t	0	1
0	0.95	0.05
1	0.02	0.98

V_t	0	1
0	0.96	0.04
1	0.05	0.95

New

C_t	0	1
0	0.97	0.03
1	0.04	0.96

V_t	0	1
0	0.94	0.06
1	0.06	0.94

Table 4: Estimated parameters (Old and New) of the fitted non-seasonal switching model. The left panel gives the estimates of the state means and standard deviations. The remaining panels give the estimated transition probabilities for the Markov chains C_t and V_t.

where p denotes the total number of parameters in the model. In practice the EM algorithm is a robust and secure method for exploring the surface of the likelihood and its approximations to obtain suitable parameter estimates. Where necessary these can be further refined by numerical maximisation of the log-likelihood. Full details of the procedures used and approximations made is given in Thomson (2014).

The likelihood calculations were initiated from a variety of starting points for the parameters and two local maxima or turning points identified. One, the absolute maximum, estimated parameters that were strongly influenced by the post-2009 data and, among other differences, produced estimates of the state mean levels that were considerably higher than those determined in Thomson (2014) and more in keeping with the contracted scale of the post-2009 data. The other local maximum produced parameters that were very similar to those determined in Thomson (2014) based on data to 30 September 2012 and produced very similar results. In keeping with these associations we refer to the two solutions as the **New** and **Old** parameter estimates respectively. However it should be noted that both are determined from all the storage data and their differences reflect their ability to handle the structural break introduced by the 2009 Ministerial Review of Electricity Market Performance. The Old and New parameter estimates are given in Table 4.

First consider the Old parameter estimates given in Table 4 and the mapping given in Table 3. The two lowest state means correspond to $C_t = 0$ and the two highest to $C_t = 1$ which leads us to identify C_t as indexing a storage state or season with $C_t = 0$ denoting the low storage season and $C_t = 1$ as the high storage season. Within each of the storage states, $V_t = 0$ corresponds to intermediate storage levels and $V_t = 1$ to extreme storage levels (lowest in the case of the low storage state $C_t = 0$ and highest in the case of the high storage state $C_t = 1$). As in Thomson (2014) we can interpret C_t as the the primary storage regime with two levels (low when $C_t = 0$, high when $C_t = 1$) and, within each of these regimes, the secondary storage regime V_t differentiates between an intermediate ($V_t = 0$) or extreme ($V_t = 1$) storage level. The transition probabilities of the Markov chains C_t and V_t show that all self-transition probabilities are highly persistent (all exceed 0.95) with low storage $C_t = 0$ less persistent than high storage $C_t = 1$ and intermediate storage $V_t = 0$ more persistent than extreme storage $V_t = 1$.

The New parameter estimates show similar interpretations. The most significant difference is that now, by comparison to the Old parameter estimates, all state means have been shifted upwards with the two lowest state means ($S_t = 1$, $S_t = 2$ or $C_t = 0$) shifted up the most and the two highest state means ($S_t = 3$, $S_t = 4$ or $C_t = 1$) shifted up the least. Indeed, the $S_t = 1$ state mean for the New parameters is now very similar to the $S_t = 3$ state mean for the Old parameters, yet the $S_t = 4$ state means for both parameter sets are not far apart. This contraction is consistent with the changes observed in Section 3.2 and due to the 2009 Ministerial Review of Electricity Market Performance. The transition probabilities of the Markov chains C_t and V_t are similar to the Old parameter estimates, but are slightly less persistent in general.

In the following analysis, attention is restricted to the solution provided by the Old parameter estimates. This choice preserves the classification and analysis of Thomson (2014) which was based on data to 30 September 2012. In particular, it allows us to check for any differences in dynamics that may have occurred since then.

The EM algorithm depends on the classification probabilities

$$\gamma_t(j,k) = P(S_{t-1} = j, S_t = k|\mathbf{X}), \qquad \gamma_t(j) = P(S_t = j|\mathbf{X}) = \sum_{k=1}^{4} \gamma_t(j,k). \tag{10}$$

where $\mathbf{X}$ denotes the data $X_1, \ldots, X_T$ and T denotes the number of weekly observations available. The classification probabilities are useful in their own right and can also be used to extract estimates of stochastic quantities such as μ_{S_t} and σ_{S_t} among many other possibilities. For example, an estimate of μ_{S_t} is given by

$$E(\mu_{S_t}|\mathbf{X}) = \sum_{j=1}^{4} \mu_j \gamma_t(j) \tag{11}$$

which is an estimate of the mean level or trend of X_t over time called the HMM trend. Such quantities, together with the classification probabilities $\gamma_t(j)$, are particularly useful for diagnostic purposes to assess the quality of the fitted model.

Figure 15 shows the Waitaki weekly average storage with the HMM trend superimposed as well as plots of the classification probabilities $P(C_t = 0|\mathbf{X})$ and $P(V_t = 0|\mathbf{X})$ where the latter are calculated using the mapping given in Table 3. The HMM trend closely follows the general movement of the 25 week triangular moving average (also shown in Figure 15). By contrast to the 25 week moving average, the HMM trend is more adaptive and able to accommodate the sharp changes in level of the storage data and, in addition, provide trend estimates at the ends of the series. Despite the fact that it is just a weighted average of the four estimated mean levels, the HMM trend provides a good summary of the mean storage over time. The classification probabilities are generally very definite in their classifications with most probabilities being close to 0 or 1. Since $C_t = 0$ denotes the low storage regime, Figure 15 shows that the primary storage regimes indexed by C_t are highly persistent with varying durations. They also appear to be seasonal although some seasons are absent in some years. The secondary regimes indexed by V_t are also persistent, but not as persistent as those for C_t.

As noted in Thomson (2014), the fidelity of the HMM trend, the definiteness of the classification probabilities and their persistence provide strong empirical evidence in support

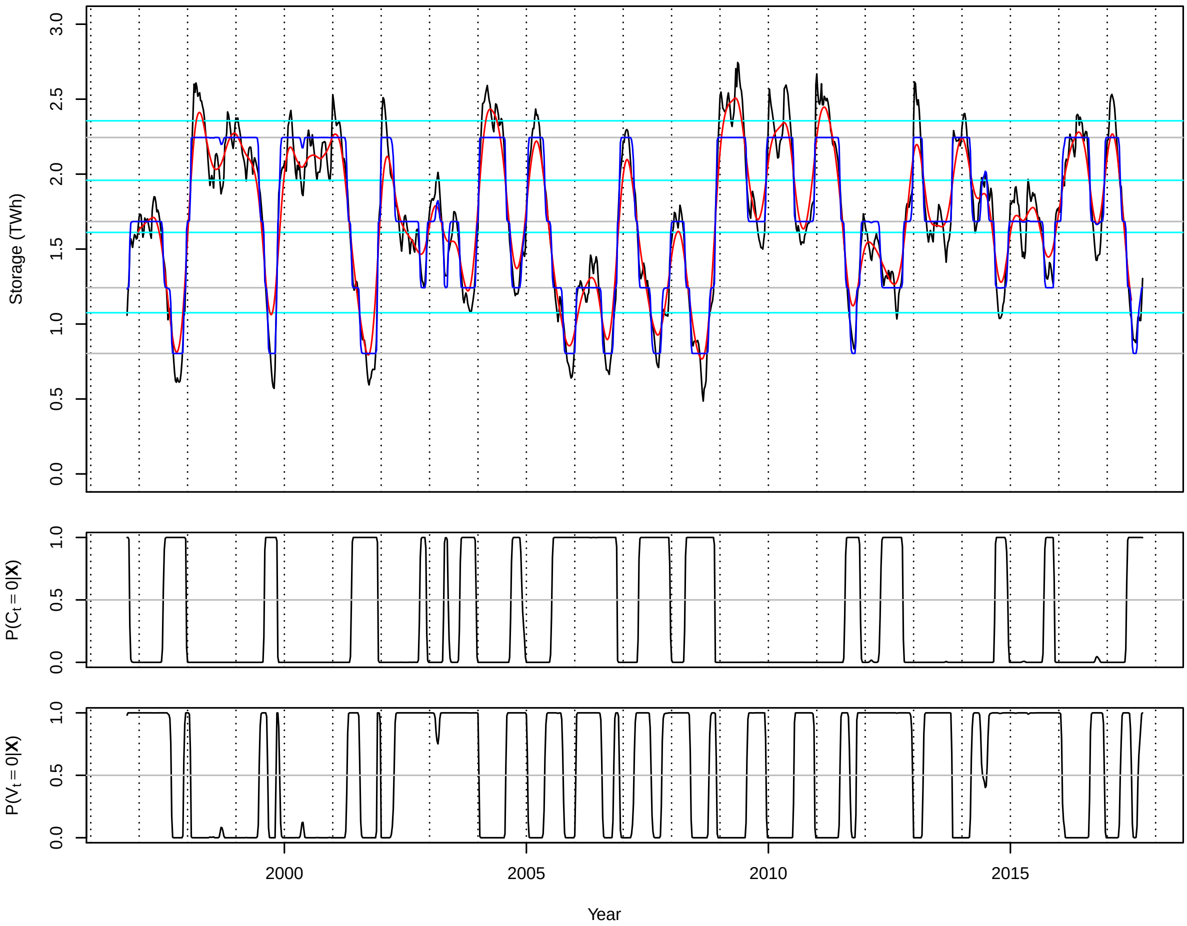

Figure 15: Waitaki weekly average storage (black, top panel) with the HMM trend (blue) and estimates of the state mean levels μ_j (horizontal grey) for the Old parameter estimates. A 25 week triangular moving average (red) of weekly storage and estimates of the state mean levels for the New parameters (horizontal cyan) are shown for reference. The two lower panels give the classification probabilities $P(C_t = 0|\mathbf{X})$ and $P(V_t = 0|\mathbf{X})$ respectively.

of a regime switching model such as (9). Since regime onsets and durations are readily identified, they can be used to examine the nature of state transitions and to check for evidence of seasonality in the regime dynamics as well as any changes to the dynamics pre and post 30 September 2009. They are also used to estimate the standardised process Z_t and identify its stochastic properties.

An informative view of the seasonal dynamics is given by Figure 16 which shows the boxplots of the classification probabilities $P(C_t = 0|\mathbf{X})$ and $P(V_t = 0|\mathbf{X})$ with mean classification probabilities by four week period of the year superimposed. The latter estimate the proportion of visits to the low storage regime $C_t = 0$ and the intermediate storage regime $V_t = 0$ by four week period of the year. Evidently $P(C_t = 0|\mathbf{X})$ is strongly seasonal with the low storage regime typically occurring in Spring as expected. If values of $P(C_t = 0|\mathbf{X})$ greater than 0.5 are used to classify a particular week as in the low storage

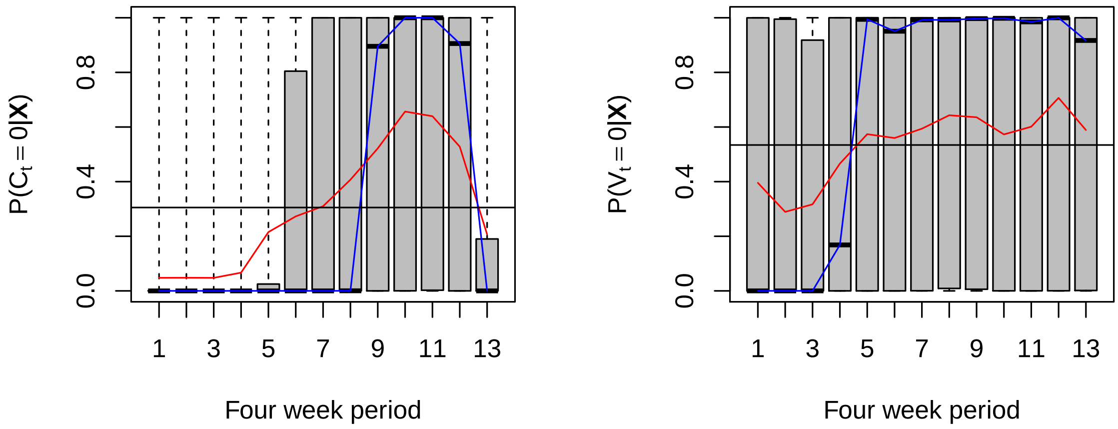

Figure 16: Boxplots of the classification probabilities $P(C_t = 0|\mathbf{X})$ (left) and $P(V_t = 0|\mathbf{X})$ (right) for Waitaki weekly average storage by four week period of the year with median (blue) and mean (red) classification probabilities superimposed. In each case the overall mean classification probability (horizontal black) is shown for reference.

regime or not, then low regimes typically occur in the four week periods 9–12 (mid August to the end of November), but can start as early as four week period 5 (late April to late May) or finish as late as four week period 13 (December). By comparison $P(V_t = 0|\mathbf{X})$ shows much less evidence of seasonality and the proportion of visits to the intermediate storage regime is more constant across the year. In this case the intermediate storage regimes typically occur in the four week periods 5–13 (late Autumn to early Summer). These results confirm and strengthen the findings of Thomson (2014). Strong seasonality is present in C_t with a lack of seasonality present in V_t.

The simple non-seasonal HMM adopted assumes that C_t and V_t are independent Markov chains. However it is possible that they are dependent with V_t depending on whether C_t is in the low or high storage regime as well as V_{t-1}. More generally, the dynamics of C_t and V_t may well have changed post 30 September 2009. A more detailed analysis of the estimated state classifications and classification probabilities (10) is needed to assess such issues. To this end, the simple *moment estimates* described in Thomson (2014) are used to estimate suitable transition probabilities and related quantities for all, pre-2009 and post-2009 periods.

Table 5 gives the moment estimates of key probabilities and transition probabilities for all, pre-2009 and post-2009 periods. The estimates of the long-run or stationary probability distributions of the Markov chains C_t and V_t show that, while the results for all and pre-2009 periods are very similar, differences have emerged over the post-2009 period. In particular, over the post-2009 period the low storage regime $C_t = 0$ occurs much less frequently (almost half as frequently as over the pre-2009 period) and the intermediate storage regime $V_t = 0$ has become more prevalent. The moment estimates of the transition probability matrices of C_t and V_t are similar for all periods with the exception of C_t over the post-2009 period. Here the low storage regime $C_t = 0$ is less persistent and the high

	All: 0	All: 1	Pre-2009: 0	Pre-2009: 1	Post-2009: 0	Post-2009: 1
$P(C_t)$	0.31	0.69	0.35	0.65	0.17	0.83
$P(V_t)$	0.53	0.47	0.50	0.50	0.65	0.35

C_t	All: 0	All: 1	Pre-2009: 0	Pre-2009: 1	Post-2009: 0	Post-2009: 1
0	0.95	0.05	0.96	0.04	0.93	0.07
1	0.02	0.98	0.02	0.98	0.01	0.99

V_t	All: 0	All: 1	Pre-2009: 0	Pre-2009: 1	Post-2009: 0	Post-2009: 1
0	0.96	0.04	0.96	0.04	0.97	0.03
1	0.05	0.95	0.04	0.96	0.06	0.94

Table 5: Estimates of the unconditional probability functions and transition probability matrices of the 2-state Markov chains C_t and V_t for all (left panel), pre-2009 (middle panel) and post-2009 (right panel) data.

regime more persistent than for the other periods.

To check the dependence of the secondary storage regime V_t on the primary storage regime C_t, moment estimates of the transition probability matrices of V_t conditional on the value of C_t are shown in Table 6. When $C_t = 0$ the estimated conditional transition probability matrices of V_t for all periods are very similar and the same is true when $C_t = 1$. The exception is the post-2009 period where the probability of a transition from extreme to intermediate storage ($V_t = 1$ to $V_t = 0$) is higher than the other periods indicating greater risk aversion over the post-2009 period. In general, for each period the transition from extreme to intermediate storage ($V_t = 1$ to $V_t = 0$) in the low storage regime $C_t = 0$ is at least twice as likely as that for the high storage regime $C_t = 1$. The latter finding provides strong evidence that C_t and V_t are dependent over all periods with greater risk aversion present post 30 September 2009.

Finally, we briefly examine the dynamic structure of the stationary process Z_t. As noted

$V_t \mid C_t = 0$	All: 0	All: 1	Pre-2009: 0	Pre-2009: 1	Post-2009: 0	Post-2009: 1
0	0.96	0.04	0.96	0.04	0.97	0.03
1	0.08	0.92	0.08	0.92	0.13	0.87

$V_t \mid C_t = 1$	All: 0	All: 1	Pre-2009: 0	Pre-2009: 1	Post-2009: 0	Post-2009: 1
0	0.96	0.04	0.96	0.04	0.96	0.04
1	0.04	0.96	0.03	0.97	0.06	0.94

Table 6: Estimates of the transition probabilty matrices for the 2-state Markov chain V_t conditioned on the value of the storage regime C_t for all (left panel), pre-2009 (middle panel) and post-2009 (right panel) data.

in Thomson (2014), the residuals from the HMM trend (11) are given by

$$X_t - E(\mu_{S_t}|\mathbf{X}) = X_t - \sum_{j=1}^{4} \mu_j \gamma_t(j)$$

where the $\gamma_t(j)$ are the state classification probabilities and the μ_j are replaced by their estimates. This residual time series estimates $\sigma_{S_t} Z_t$ and reflects time-varying volatility present in the weekly storage X_t due to σ_{S_t}. To correct for this time-varying volatility, estimates of the standardised residuals Z_t are given by

$$\hat{Z}_t = E(\frac{X_t - \mu_{S_t}}{\sigma_{S_t}}|\mathbf{X}) = \sum_{j=1}^{4} \gamma_t(j) \frac{X_t - \mu_j}{\sigma_j} \tag{12}$$

where, as before, the $\gamma_t(j)$ are the state classification probabilities and the μ_j, σ_j are replaced by their estimates.

Plots of the residuals, the standardised residuals and their autocorrelation functions are given in Figure 17. Adjusting the Waitaki, weekly average, storage X_t by its HMM trend reduces the variability of X_t by 73%. This is a significant reduction that underscores the importance and quality of the Markov switching level μ_{S_t}. The HMM trend adjustment also removes the dynamics of μ_{S_t} so that both residual series show markedly less autocorrelation structure than the original series X_t. Nevertheless, both residual series are very similar and show significant autocorrelation structure which will need to be modelled.

An autoregressive moving-average (ARMA) process with zero mean was fitted to $\hat{Z}_t$ using conventional time series techniques and the optimal model selected using AIC and other criteria. An AR(3) model was identified with $\hat{Z}_t$ satisfying

$$\hat{Z}_t = \alpha_1 \hat{Z}_{t-1} + \alpha_2 \hat{Z}_{t-2} + \alpha_3 \hat{Z}_{t-3} + \epsilon_t \qquad (t = 1, 2, \ldots) \tag{13}$$

where the parameters were estimated as

$$\hat{\alpha}_1 = 1.19 \ (0.03), \quad \hat{\alpha}_2 = -0.54 \ (0.04), \quad \hat{\alpha}_3 = 0.14 \ (0.03).$$

with standard errors in parentheses. This model should prove appropriate for short-term forecasting of the standardised residual process Z_t.

The simple non-seasonal HMM given by (9) has, once again, proved to be a suitable and flexible framework to identify the regime switching structure of Waitaki weekly average storage. Secure state classification probabilities have led to a better understanding of the seasonal dynamics of weekly storage. There is strong evidence of switching seasonal regimes with the dynamic structure within regimes well-modelled by an ARMA process. However the post-2009 dynamics have changed compared to pre-2009 with a greater prevalence of intermediate storage states and a lower prevalence of low storage states indicating greater risk aversion to extreme low storage. These results are consistent with the findings in Section 3.2 and are a consequence of the 2009 Ministerial Review of Electricity Market Performance.

The exploratory analysis of this section confirms, and further supports, the seasonal regime switching model developed in Thomson (2014).

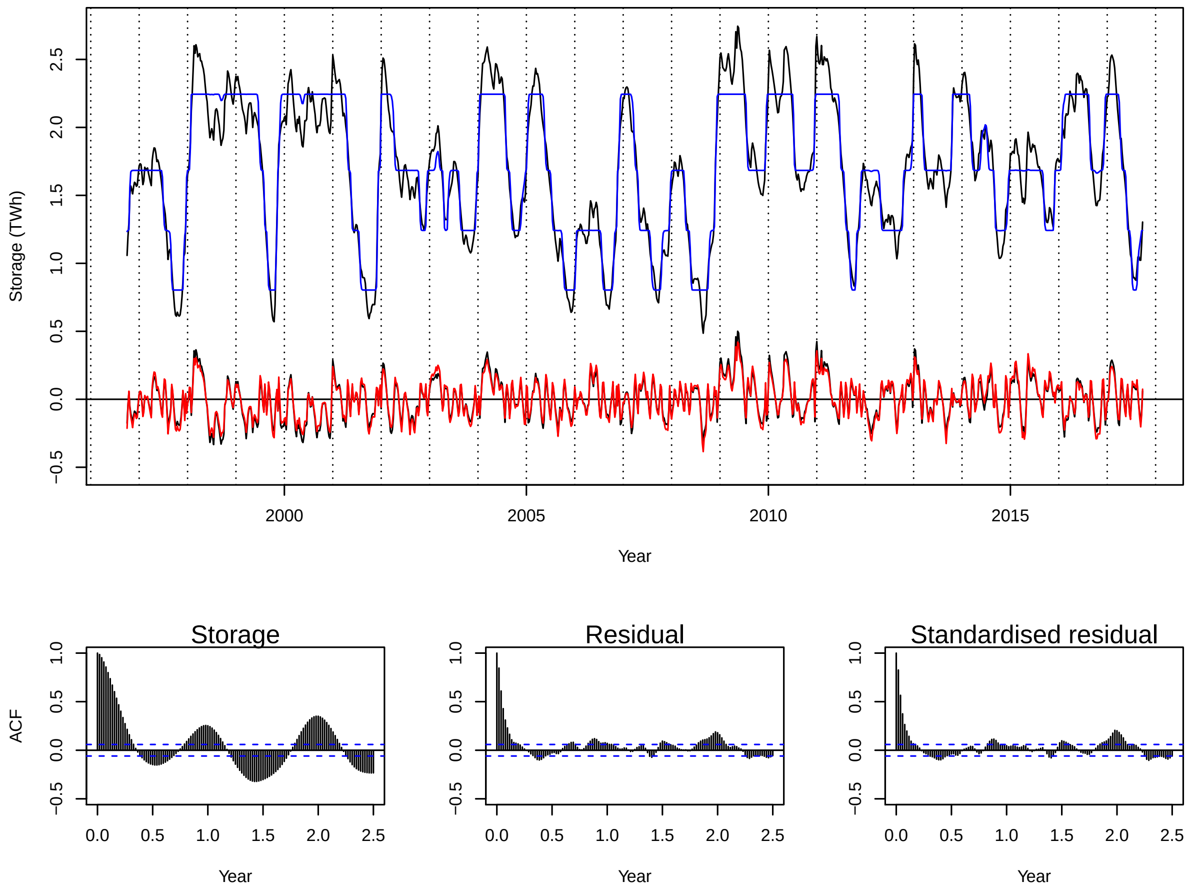

Figure 17: Waitaki weekly average storage (black, upper panel) with HMM trend (blue) superimposed. The upper panel also shows the residual series after HMM trend adjustment (black) as well as the standardised residuals (red) scaled to have the same standard deviation as the residual series. The lower panels show the autocorrelation functions of these three series.

4.2 Seasonal switching model

In this section we fit the seasonal switching model developed in Thomson (2014) to Waitaki weekly average storage over the period 30 September 1996 to 30 September 2017 and evaluate the results. This model builds on Carey-Smith et al. (2014), who developed such models for New Zealand daily rainfall, and Harte and Thomson (2007) who suggested similar models for New Zealand, hydro catchment, weekly inflows. By contrast to conventional fixed seasons and strictly periodic seasonality, these seasonal switching models have the key property that annual seasons can occur earlier or later than expected and have varying durations.

Following Thomson (2014), the stochastic storage season C_t (low and high) is now modelled by a non-homogeneous Markov chain and, within each storage season C_t, the secondary storage state V_t (intermediate and extreme) follows a homogeneous Markov chain so that V_t can be dependent on C_t. To specify C_t, the weeks of the year need to be blocked into two mutually exclusive *season change intervals* $(\tau_0, \tau_1]$ and $(\tau_1, \tau_0]$ with the convention that these intervals are wrapped circularly around the 52 weeks of the year. The *season anchor points* τ_0, τ_1 are fixed with $C_{\tau_0} = 0$ and $C_{\tau_1} = 1$ so that week of the year τ_0 is always in the low storage season and week of the year τ_1 is always in the high storage season. Over the interval $(\tau_0, \tau_1]$ the storage regime C_t is assumed to change once from low storage ($C_t = 0$) to high storage ($C_t = 1$) and, over the interval $(\tau_1, \tau_0]$, it is assumed to change once from high storage ($C_t = 1$) to low storage ($C_t = 0$). These conditions guarantee an orderly succession of seasons with each season occuring once each year. Although the season change intervals are fixed, the onset of each season and its duration can vary from year to year. In essence, this model replaces fixed annual seasons by fixed seasonal change intervals.

The stochastic seasons C_t are assumed to follow a 2-state Markov chain with non-homogeneous transition probability matrix

$$\mathbf{Q}(w) = \left[\begin{array}{cc} Q_{00}(w) & Q_{01}(w) \\ Q_{10}(w) & Q_{11}(w) \end{array} \right] \tag{14}$$

where w denotes week of the year,

$$\mathbf{Q}(w) = \left[\begin{array}{cc} 1-q(w) & q(w) \\ 0 & 1 \end{array} \right] \quad (w \in (\tau_0, \tau_1]), \qquad \mathbf{Q}(w) \quad = \left[\begin{array}{cc} 1 & 0 \\ q(w) & 1-q(w) \end{array} \right] \quad (w \in (\tau_1, \tau_0])$$

and $q(\tau_0) = q(\tau_1) = 1$. Here $q(w)$ is called the *season change probability* since it gives the probability of a switch in the storage season for week w in each of the two season change intervals. This function is defined over all weeks of the year and reflects the stochastic properties of storage season onsets and durations.

A simple example of $q(w)$ is shown in Figure 18 together with a realisation of the storage season C_t over a year. For each season change interval, Carey-Smith et al. (2014) show that $q(w)$ is the hazard function of the distribution of the season onset time. This allows $q(w)$ to be specified directly, or in terms of the distributions of the season onset times. For example, if season onsets were equally likely to occur at any point in their respective season change intervals then

$$q(w) = \left\{ \begin{array}{ll} \frac{1}{\tau_1 - w + 1} & (w = \tau_0 + 1, \ldots, \tau_1) \\ \frac{1}{\tau_0 - w + 1} & (w = \tau_1 + 1, \ldots, \tau_0) \end{array} \right. . \tag{15}$$

This simple model involves no parameters, apart from the season anchor points, a fact that is useful for exploratory in-sample analysis. The example of $q(w)$ shown in Figure 18 is for uniform season onset times. In general, any suitable family of distributions can be chosen for the onset distributions.

If the storage season is $C_t = c$, then the secondary storage level V_t has transition probability matrix

$$\mathbf{P}^{(c)} = \left[\begin{array}{cc} P_{00}^{(c)} & P_{01}^{(c)} \\ P_{10}^{(c)} & P_{11}^{(c)} \end{array} \right] = \left[\begin{array}{cc} 1 - p_0^{(c)} & p_0^{(c)} \\ p_1^{(c)} & 1 - p_1^{(c)} \end{array} \right] \qquad (c = 0, 1) \tag{16}$$

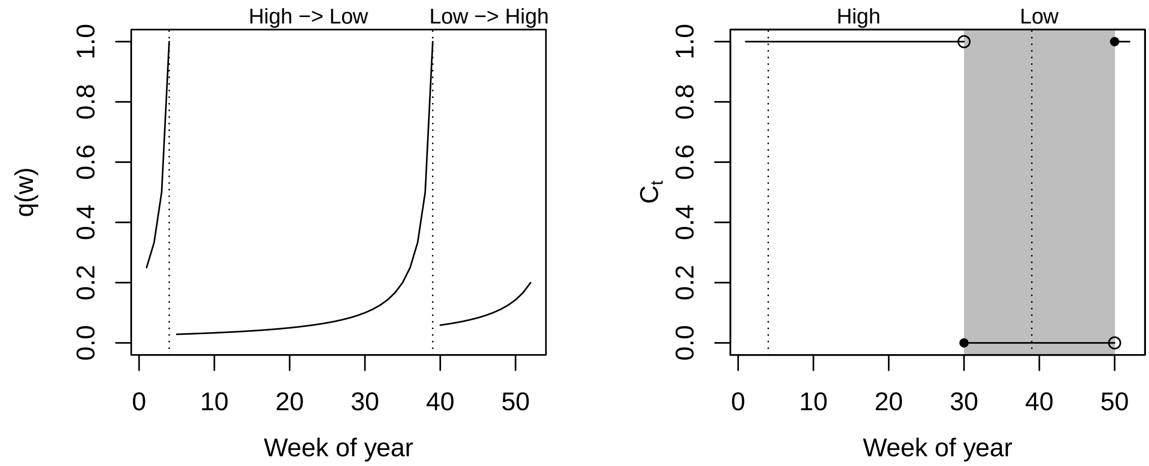

Figure 18: The left panel shows an example of the season change probability $q(w)$ with the season change intervals defined by the season anchor points (vertical dotted lines). The right plot shows a realisation of the storage season C_t over a year ($C_t = 0$ denotes the low season and $C_t = 1$ the high season) with the shaded period showing the low storage season for that year.

so that the dynamics of V_t depend only on the current season C_t and the previous week's secondary storage level V_{t-1}, while the dynamics of C_t depend only on the previous storage season C_{t-1}. A summary of the dynamic structure of the various states with their possible transitions and transition probabilities is shown in Figure 19. Further details can be found in Thomson (2014).

Following the same procedures as those used in Thomson (2014), the stochastic seasonal switching model (9) with uniform season onsets was fitted to Waitaki weekly average storage over the period 30 September 1996 to 30 September 2017. Here the low season anchor point was identified as week $\tau_0 = 39$ (end of September) while the high season anchor point was identified as week $\tau_1 = 4$ (late January). These differ only very slightly

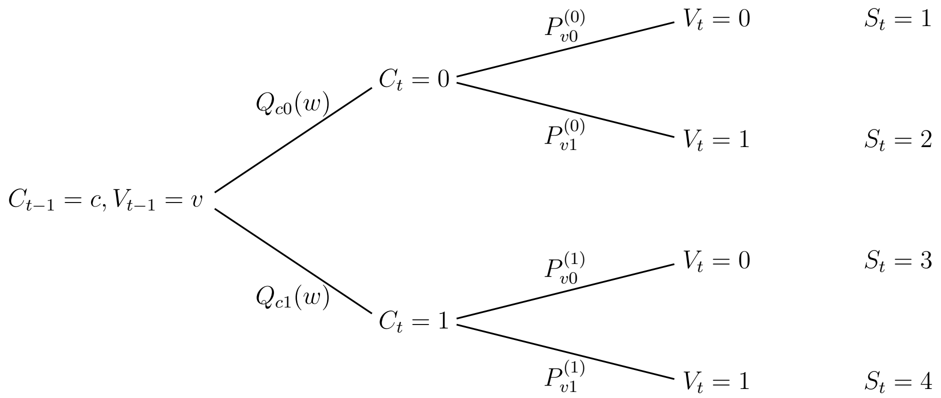

Figure 19: A transition diagram showing the possible transitions and transition probabilities for the storage seasons C_t, the secondary storage levels V_t and the states S_t.

S_t	$\hat{\mu}_{S_t}$	$\hat{\sigma}_{S_t}$
1	1.30	0.16
2	0.83	0.16
3	1.72	0.16
4	2.27	0.18

$V_t \vert C_t = 0$	0	1
0	0.96	0.04
1	0.08	0.92

$V_t \vert C_t = 1$	0	1
0	0.96	0.04
1	0.04	0.96

Table 7: Parameter estimates for the seasonal switching model fitted to Waitaki weekly average storage with uniform season onsets and season anchor points at weeks 4 and 39. The left panel gives the estimates of the state means and standard deviations. The remaining two panels give the estimated transition probability matrices for the 2-state Markov chain V_t conditioned on the storage regime C_t.

from the values $\tau_0 = 39$ and $\tau_1 = 5$ used by Thomson (2014). As in the case of the non-seasonal switching model, there were two possible models identified corresponding to the Old and New parameters of Section 4.1. Again we focus on the solution corresponding to the Old parameters which is given in Table 7.

The parameter estimates given in Table 7 are very similar to those reported in Thomson (2014), especially for the dynamics of the secondary storage regime V_t, and similar comments apply. The estimated state means in Table 7 are also in good agreement with the state means for the non-seasonal HMM fitted in Section 4.1 and reported in Table 4. Furthermore, the transition probability matrices of V_t conditional on the storage season C_t given in Table 7 are essentially the same as the moment estimates given in Table 6 for the non-seasonal HMM. As in Thomson (2014), the conditional transition probability matrices in Table 7 confirm the dependence of V_t on C_t with transitions from $V_t = 1$ (extreme storage state) to $V_t = 0$ (intermediate storage state) twice as likely when $C_t = 0$ (low storage season) than when $C_t = 1$ (high storage season). This indicates, in general, a greater risk aversion to extremely low storage as compared to extremely high storage.

Figure 20 shows the Waitaki weekly average storage and the HMM trend from the seasonal switching model with uniform season onsets and season anchor points $\tau_0 = 39$, $\tau_1 = 4$. The associated classification probabilities $P(C_t = 0|\mathbf{X})$ and $P(V_t = 0|\mathbf{X})$ are also shown. All the plots are very similar to those for the non-seasonal HMM given in Figure 15 and similar comments apply. As before, the HMM trend closely follows the general movement of the 25 week triangular moving average with very few exceptions. Since the non-seasonal HMM is very flexible, the close agreement of the two fitted models (non-seasonal and seasonal) suggests that little is lost by adopting the more constrained seasonal switching model.

However, as in the earlier analysis Thomson (2014), there are important points of difference. In the four years 1998, 2000, 2009 and 2010 the low storage season does not appear to have occurred and, in 2006, there doesn't appear to have been a high storage season. For these years the requirement that the season anchor points τ_0 and τ_1 must always be in the low and high storage seasons respectively has led to large trend deviations at the season anchor points. These large trend deviations have the potential to bias the analysis of the residuals and should be removed either by changing the model so that it can accommodate missing seasons, or by empirical methods such as censoring. As in Thomson

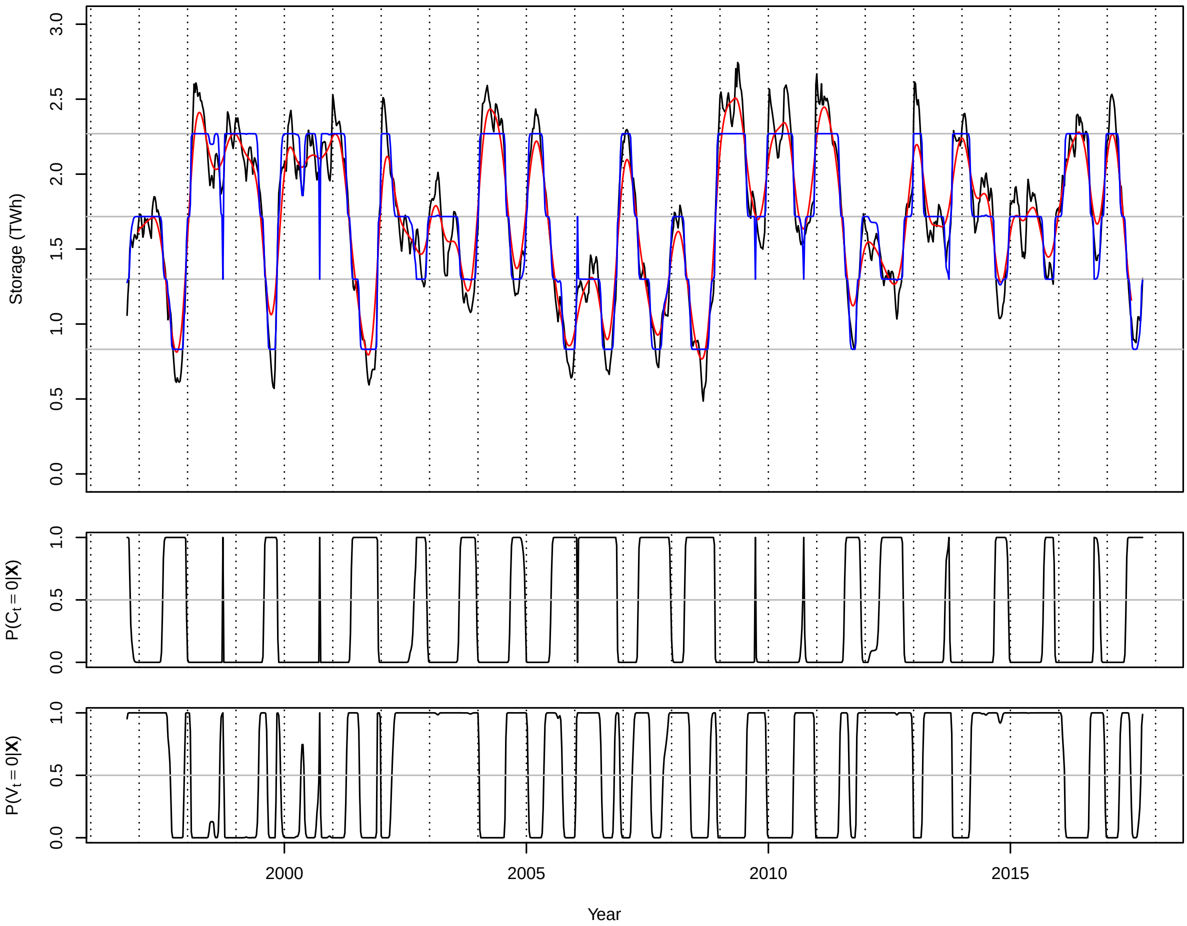

Figure 20: Fit of the seasonal switching model to Waitaki weekly average storage (black, top panel) with uniform season onsets and season anchor points at weeks 4 and 39. The HMM trend (blue) and estimates of the state mean levels μ_j (horizontal grey) are superimposed and a 25 week triangular moving average (red) of Waitaki weekly storage is shown for reference. The lower panels give the classification probabilities $P(C_t = 0|\mathbf{X})$ and $P(V_t = 0|\mathbf{X})$ respectively.

(2014), we adopt the latter approach and impose the condition that storage seasons have a minimum length (3 weeks or more). This simple *censoring rule* can be achieved here by adjusting the values of the estimated state classification probabilities at the anchor points for the years concerned. See Thomson (2014) for further details.

Adjusting the state classification probabilities in this way yields the censored seasonal HMM trend for the Waitaki weekly average storage shown in Figure 21. Also shown in Figure 21 are the residuals after adjusting for the censored HMM trend, the standardised residuals calculated using (12) with the adjusted state classification probabilities, and the autocorrelation functions of weekly storage, residuals and standardised residuals. As in the case of the non-seasonal HMM, the residuals and standardised residuals show much less autocorrelation (dynamic) structure by comparison to the original series X_t.

An ARMA model with zero mean was fitted to the standardised residuals with the optimal

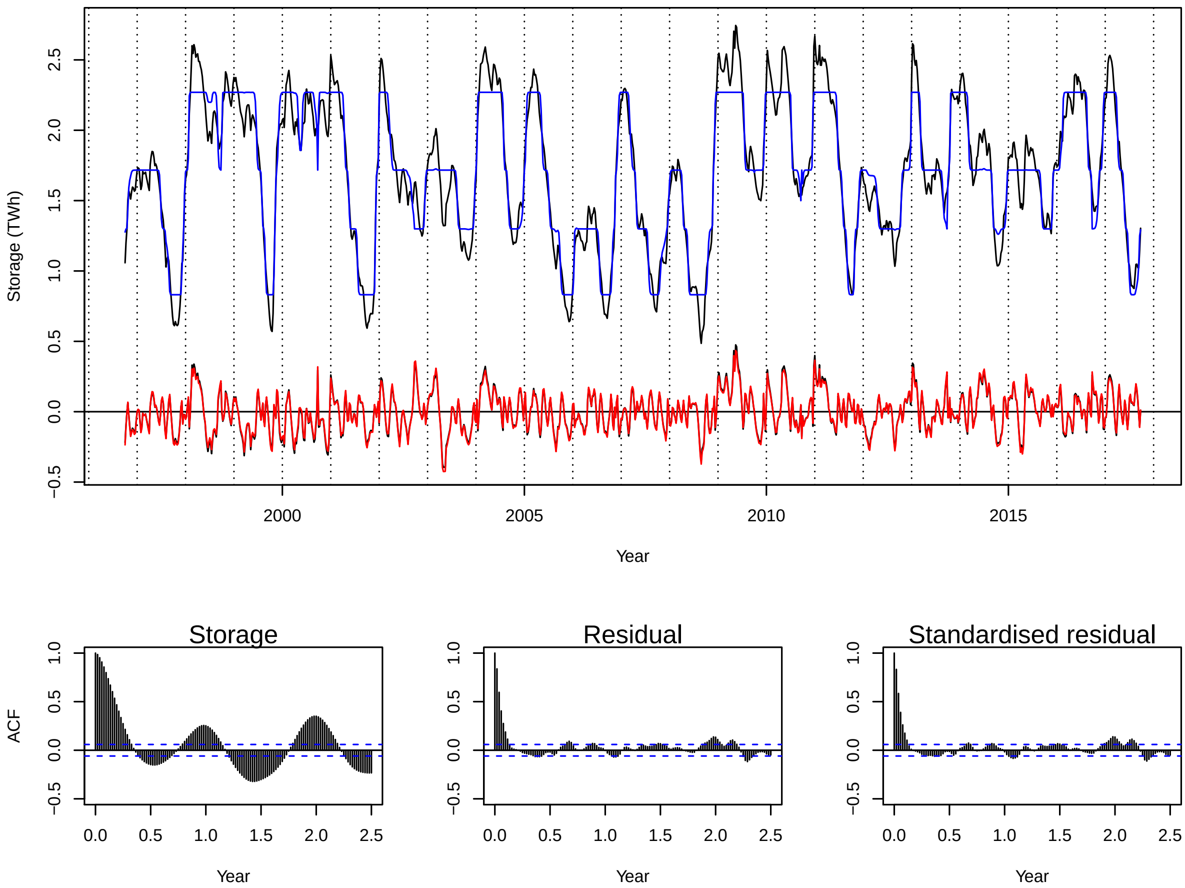

Figure 21: Waitaki weekly average storage (black, upper panel) with the censored HMM trend (blue) from the seasonal switching model superimposed. The upper panel also shows the residual series after trend adjustment (black) and the standardised residuals (red). The latter are calculated using the adjusted state classification probabilities and scaled to have the same standard deviation as the residual series. The various autocorrelation functions are shown in the lower panels.

model selected using AIC and other criteria. An ARMA(1,1) model was identified for $\hat{Z}_t$ which satisfies

$$\hat{Z}_t = \alpha \hat{Z}_{t-1} + \epsilon_t + \beta \epsilon_{t-1} \qquad (t = 1, 2, \dots) \tag{17}$$

with the parameters estimated as

$$\hat{\alpha} = 0.73 \ \ (0.02), \qquad \hat{\beta} = 0.44 \ \ (0.03)$$

and standard errors given in parentheses. A competing, less parsimonious, AR(3) model satisfying (13) has estimated parameters

$$\hat{\alpha}_1 = 1.18 \ \ (0.03), \quad \hat{\alpha}_2 = -0.50 \ \ (0.04), \quad \hat{\alpha}_3 = 0.11 \ \ (0.03).$$

Either model should prove appropriate for short-term forecasting of the standardised residual process Z_t.

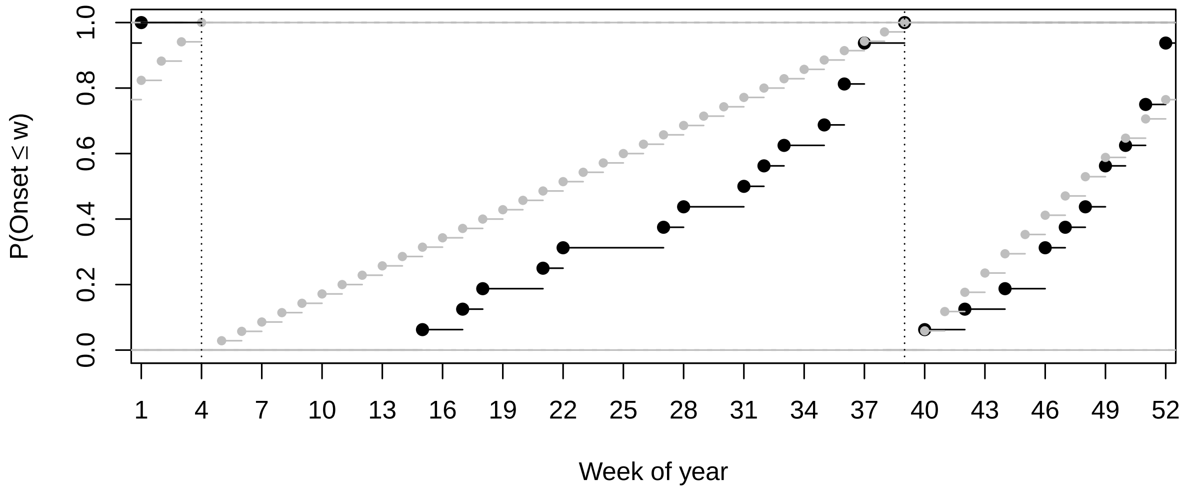

Figure 22: Empirical distribution functions of the onsets of low and high seasonal storage regimes for Waitaki weekly average storage. The season change intervals are indicated by vertical dotted lines with the low season onset occurring in the interval from week $\tau_1 = 4$ up to and including week $\tau_0 = 39$, and the high season onset occurring in the remaining weeks of the year. Uniform onset distribution functions (grey) are shown for reference.

Now consider the storage season onsets and durations. These are obtained from the adjusted classification probabilities $P(C_t = 0|\mathbf{X})$ by defining a week as in the low storage season when $P(C_t = 0|\mathbf{X}) > 0.5$ and the high storage season otherwise. Plots of the empirical distributions of the low and high storage season onset times are given in Figure 22 together with the uniform season onset distributions assumed by the fitted model. It

Pre-2009

Low onset Year	Low onset Week	High onset Year	High onset Week	Low sojourn Weeks
		1996	44	
1997	27	1997	52	25
1999	32	1999	46	14
2001	21	2001	50	29
2002	37	2002	51	14
2003	33	2003	52	19
2004	36	2004	52	16
2005	28	2006	47	71
2007	18	2008	1	35
2008	15	2008	48	33

Post-2009

Low onset Year	Low onset Week	High onset Year	High onset Week	Low sojourn Weeks
2011	31	2011	49	18
2012	17	2012	42	25
2013	35	2013	40	5
2014	36	2014	51	15
2015	37	2015	49	12
2016	39	2016	46	7
2017	22			

Table 8: Onsets of low and high seasonal storage regimes for Waitaki weekly average storage and sojourns of the low storage season. Low storage regimes were missing in 1998, 2000, 2009 and 2010 and a high storage regime was missing in 2006.

would seem that season onsets are not uniformly distributed over their respective season change intervals, especially the low season onset. Both season onsets are distributed over more concentrated ranges with low season onsets occurring between mid March and the end of September, and high season onsets occurring from early October to early January. Although the assumption of uniform season onsets is reasonable and practical for in-sample analysis, Figure 22 shows that more appropriate distributions will be needed for any predictive or forward-looking study.

The low and high storage season onsets and low storage season sojourns are reported in Table 8. The low storage season has median onset at week 28 pre-2009 and week 35 post-2009, whereas the high storage season has median onset at week 50.5 pre-2009 and week 47.5 post-2009. Furthermore, the low season sojourns have a median of 25 weeks pre-2009 and 13.5 weeks post 2009. Although the samples pre-2009 and post-2009 are too small to give rise to statistically significant differences, they do, nevertheless, all point to later low season onsets, earlier high season onsets, and shorter low season sojourns post the 2009 Ministerial Review of Electricity Market Performance. This observation is consistent with our earlier findings.

4.3 Price-storage relationship within seasonal state

In this section we explore the relationship between transformed price and storage within the four storage seasons identified by the seasonal switching model fitted in Section 4.2. Here the storage seasons are determined from the adjusted classification probabilities by defining week t to be in state $S_t = s$ when $P(S_t = s|\mathbf{X}) > 0.5$ with storage seasons identified using Table 3. In essence, the analysis undertaken assumes that transformed weekly average spot prices $Y_t = \log(P_t - \theta_t)$ and weekly average storage levels X_t follow the joint model

$$\begin{aligned} Y_t &= \mu^Y_{S_t} + \sigma^Y_{S_t} Z^Y_t \\ X_t &= \mu^X_{S_t} + \sigma^X_{S_t} Z^X_t \end{aligned} \tag{18}$$

where Z^X_t and Z^Y_t are jointly stationary with zero means, unit standard deviations and contemporaneous correlation ρ_{S_t} that may depend on state. As in (9), the parameters $\mu^X_{S_t}$, $\sigma^X_{S_t}$ denote the conditional mean and standard deviation of X_t given S_t, and $\mu^Y_{S_t}$, $\sigma^Y_{S_t}$ denote the corresponding parameters for transformed prices Y_t. We can now write

$$Y_t = \mu^Y_{S_t} + \rho_{S_t} \frac{\sigma^Y_{S_t}}{\sigma^X_{S_t}} (X_t - \mu^X_{S_t}) + \epsilon_t \tag{19}$$

where, as in (8), the residual error process ϵ_t has zero mean. Within this framework, the conditional mean of Y_t given the storage data ($E(Y_t|\mathbf{X})$) is approximated by the linear regression of Y_t against known X_t, S_t with S_t estimated as above.

Figure 23 shows notched boxplots of transformed weekly average spot prices Y_t with constant threshold and Waitaki weekly average storage levels X_t by storage season C_t and secondary storage state V_t within C_t. Consider first the case of all data. As expected, the storage levels are well-differentiated by storage regime and state, with the state means

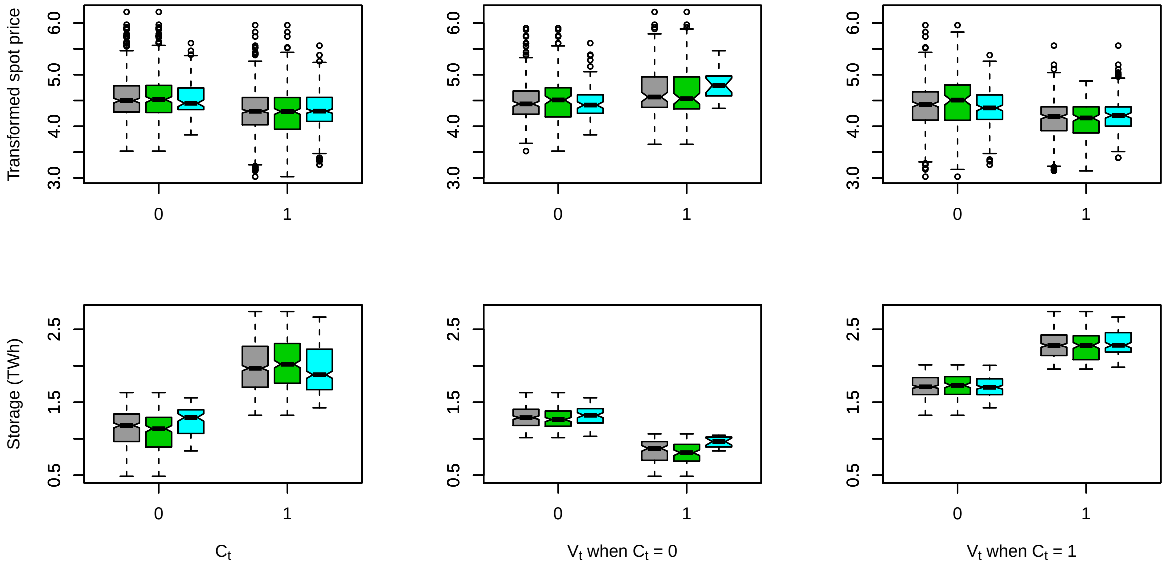

Figure 23: Notched boxplots of transformed real South Island weekly average spot prices Y_t (**upper panels**) and Waitaki weekly average storage levels X_t (**lower panels**) by storage season C_t and secondary storage state V_t within C_t. The price transformation is the shifted logarithm with constant shift and the component boxplots are for all (grey), pre-2009 (green) and post-2009 (cyan) data periods.

for all data almost exactly the same as those given in Table 7. The transformed spot prices are of more interest. In this case it is clear that the state medians when $V_t = 0$ (intermediate secondary storage) are not significantly different (the boxplots for all data have overlapping notches) so that prices would seem to be much the same when $V_t = 0$ regardless of the storage season. By contrast and as might be expected, when $V_t = 1$ (extreme secondary storage) transformed spot prices are generally higher in the low storage season ($C_t = 0$) and, in particular, lower in the high storage season ($C_t = 1$).

Now consider the pre-2009 and post-2009 boxplots in Figure 23 which reflect any changes following the 2009 Ministerial Review of Electricity Market Performance. In general, the pre-2009 medians are always close to those for all the data. This is also largely true for pre-2009 and post-2009 comparisons. The notable exception is the case of extreme secondary storage in the low storage season ($V_t = 1$ when $C_t = 0$) when both the post-2009 price and storage medians are significantly higher than their pre-2009 counterparts. While the increase in extreme low storage medians is consistent with earlier findings and storage risk aversion, the reasons for the corresponding increase in prices are less clear. Finally, given the general lack of differentiation between the price boxplots in Figure 23 by comparison to storage, one might ask how well the state means describe the overall mean level or trend of the transformed spot prices Y_t. Trend correction of Y_t using the state means yields residuals with standard error that is essentially the same as the corresponding quantity derived from the PH model fitted to all data. This suggests that the fitted mean levels or trends from both models, while not the same, have similar explanatory power.

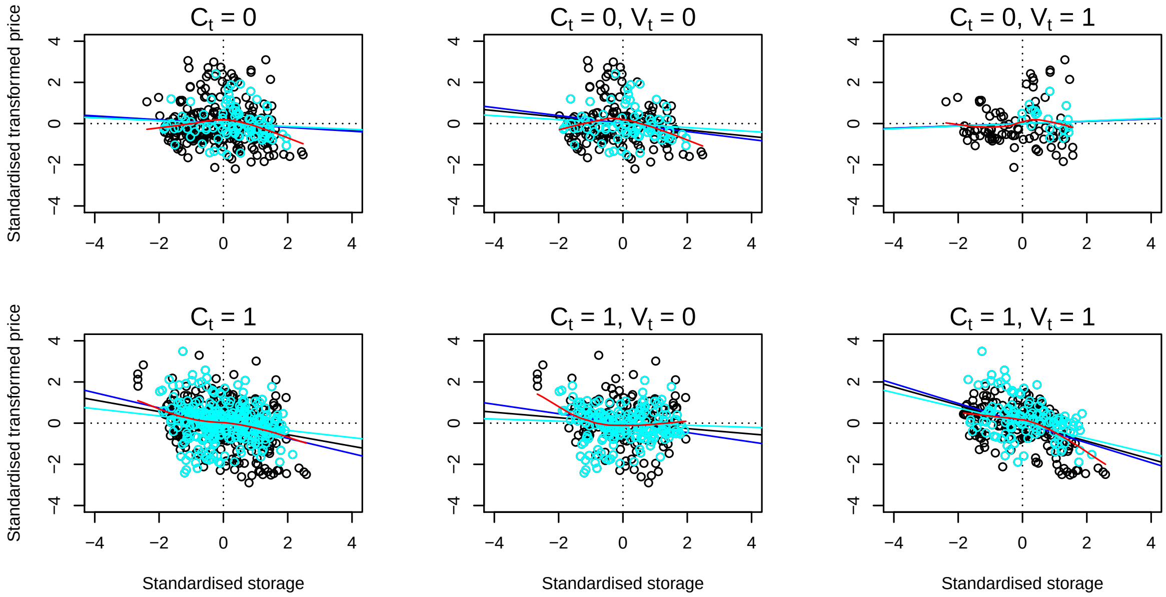

Figure 24: Scatterplots of standardised transformed real South Island weekly average electricity spot prices versus standardised Waitaki weekly average storage levels by storage season C_t and secondary storage state V_t within C_t. The price transformation is the shifted logarithm with constant shift. All data points are shown with post-2009 data points highlighted (cyan). Least squares regression lines for all (grey), pre-2009 (green) and post-2009 (cyan) data periods are superimposed as is a *loess* regression function (red) for all the data.

To fit the regression relation (19), estimates of the standardised variables Z_t^Y and Z_t^X are first obtained by standardising the transformed weekly average spot prices Y_t and weekly average storage levels X_t using their respective state means and state standard deviations. The state dependent correlation coefficients ρ_{S_t} can then be estimated by linear regression. Figure 24 shows the scatterplots of standardised transformed weekly average spot prices versus standardised weekly average storage by storage season C_t and secondary storage state V_t within C_t. Least squares regression lines for all, pre-2009 and post-2009 periods are superimposed as is a *loess* regression function for all the data. A summary of the regression relationships is given in Table 9.

	All	Pre-2009	Post-2009
$C_t = 0$	-0.09 (0.06)	-0.09 (0.07)	-0.07 (0.08)
$C_t = 0, V_t = 0$	-0.16 (0.07)	-0.19 (0.09)	-0.10 (0.09)
$C_t = 0, V_t = 1$	0.06 (0.10)	0.06 (0.11)	0.06 (0.16)
$C_t = 1$	-0.28 (0.04)	-0.37 (0.06)	-0.18 (0.05)
$C_t = 1, V_t = 0$	-0.13 (0.06)	-0.23 (0.10)	-0.05 (0.06)
$C_t = 1, V_t = 1$	-0.44 (0.05)	-0.48 (0.07)	-0.37 (0.09)

Table 9: Slopes (correlations) and their standard errors (in brackets) for the best fitting regression lines of the standardised transformed real, South Island weekly average electricity spot prices versus standardised Waitaki weekly average storage by storage season C_t and secondary storage state V_t within C_t. The price transformation is the shifted logarithm with constant shift.

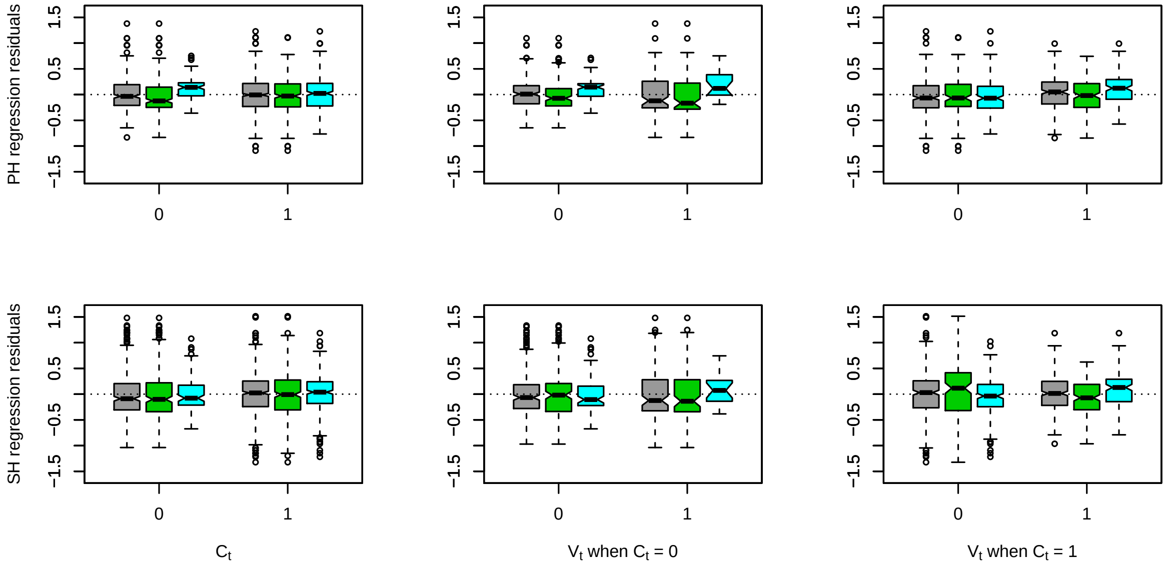

Figure 25: Notched boxplots of the regression residuals for the PH model (**upper panels**) and the switching regression model (19) (**lower panels**) by storage season C_t and secondary storage state V_t within C_t. The PH model (8) is applied to all data using the shifted logarithm transformation with constant shift, and constant correlation has been assumed across seasons. The component boxplots are for all (grey), pre-2009 (green) and post-2009 (cyan) data periods.

The regression relationships in Figure 24 are all much weaker than those for the PH model (see Figure 12 for example) since much of the price-storage relationship is now built into the state mean structure of the seasonal switching model (18). Most slopes (correlations) are close to, or not significantly different from, zero (no linear regression relationship). The exception is the extreme secondary storage state within the high storage season ($C_t = 1, V_t = 1$ or $S_t = 4$) where there is clear evidence of significant residual dependence not explained by the state means (from Table 7 this state is also the one with the highest state standard deviation). There is little evidence of significant differences in slopes (correlations) pre-2009 and post-2009. However, the apparent non-linearity of the *loess* regression functions fitted to all the data does suggest that transforming the storage data may yet have advantages.

Figure 25 shows notched boxplots of the the regression residuals for the PH model (PH residuals) and the switching regression model (19) (SH residuals) by storage season C_t and secondary storage state V_t within C_t. The PH model (8) is applied to all data using the shifted logarithm transformation with constant shift, and constant correlation has been assumed across seasons. All boxplots are distributed about zero, as expected, with similar spreads, although the latter are slightly less for the PH residuals. Indeed, the RMSE of the all data PH residuals (0.32) is 26% less than the RMSE of the SH residuals (0.43). Over the post-2009 period these figures become 0.29 and 0.38 respectively, an almost 23% reduction.

The all data boxplots for the SH residuals in Figure 25 generally have notch intervals

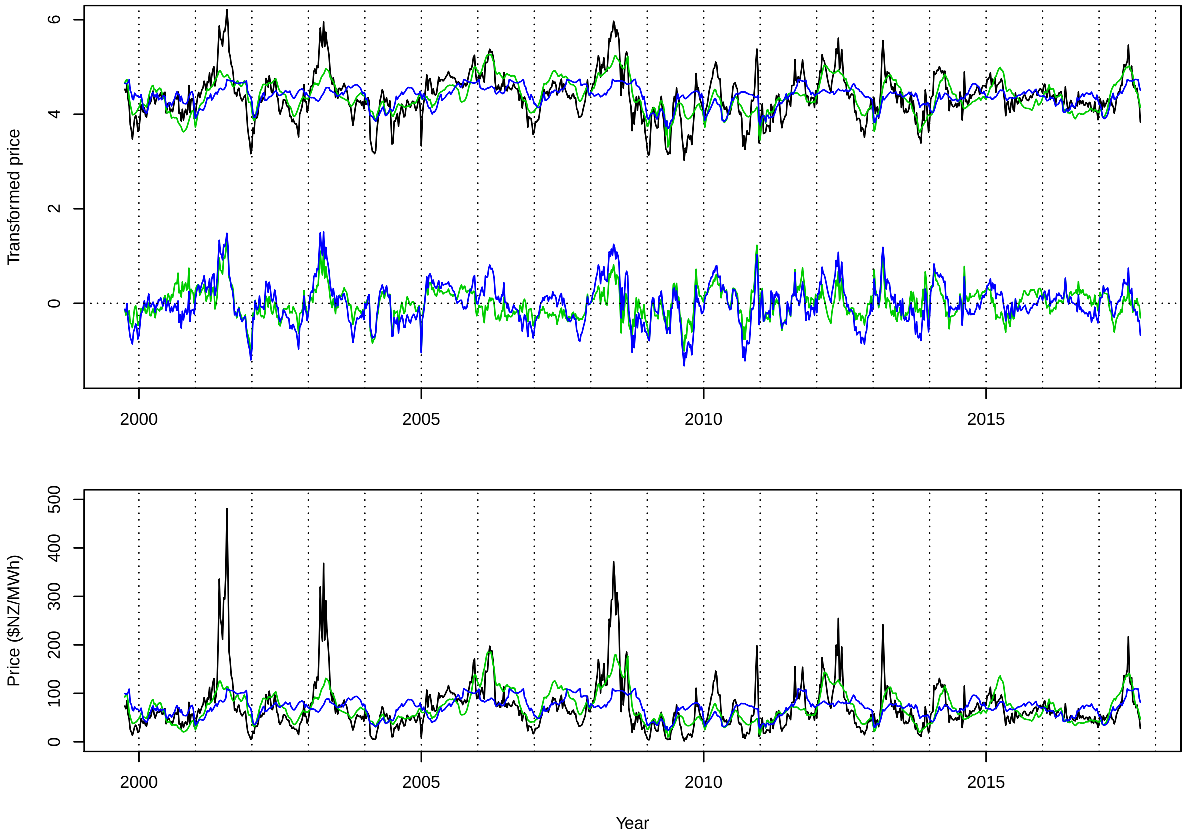

Figure 26: Plots of the real South Island weekly average electricity spot prices (bottom panel) and their transforms (top panel) together with fitted values (both panels) and residuals (top panel) from the switching regression model (18) (blue) and from the PH model (constant correlation) applied to all data (green). The spot prices have been transformed using the shifted logarithm with constant shift.

that include zero (unbiased fits) with the exception of the low storage season where lack of fit is evident. In general, the all data boxplots of the PH residuals show biases or lack of fit (notch intervals don't include zero) for all states S_t with the exception of the intermediate secondary storage state in the low storage season ($C_t = 0, V_t = 0$ or $S_t = 1$). Note that these biases cancel when considering the boxplots of the PH residuals by storage season alone. In general the pre-2009 and post-2009 medians of the SH residuals are not significantly different with the exception of the case when $V_t = 1$ and $C_t = 0$ (extreme secondary storage in the high storage season). For the most part, the pre-2009 and post-2009 medians of the PH residuals are significantly different indicating that seasonal regression alone cannot explain all the relationship between transformed weekly average spot price and weekly average storage.

Figure 26 shows plots of the fitted values and residuals from the switching regression model (18) of transformed real South Island weekly average electricity spot prices against Waitaki weekly average storage levels. Also shown are the fitted values and residuals from the seasonal regression model (6) (PH model) for all data and the case of constant correlation. As before the spot prices have been transformed using the shifted logarithm with constant shift, and their fitted values obtained using the same general procedure as that described following (7). The fits of the two models to the transformed prices are reasonable, with the PH model performing slightly better in terms of RMSE as noted earlier. Although the two predictors are highly correlated (a correlation of 0.78), there are time periods when the PH model out-performs the switching regression model and vice versa. Moreover, the autocorrelation functions of the two sets of residuals show that the switching regression residuals have much stronger residual seasonality than the PH model residuals. This is not unexpected. The switching regression model is based on dynamic storage seasons that are a function of hydro storage levels alone, whereas the PH regression model is based on static seasonal patterns that reflect seasonal demand for electricity in addition to seasonal storage and other possible covariates. Despite this limitation, the switching regression model based on storage seasons manages to provide a competitive and informative view of the relationship between price and storage.

4.4 Summary

The seasonal switching model (SH model) developed in Thomson (2014) for Waitaki weekly average hydro storage has largely been revalidated on 21 years of data to 30 September 2017. This model, a non-homogeneous hidden Markov model (NHMM), more accurately reflects the stochastic nature of seasonal weekly storage with season onset times that can occur earlier or later than expected and storage seasons that can vary in length from year to year. The NHMM has two primary storage seasons (high and low) within which weekly storage switches between two secondary storage levels (intermediate and extreme). The historical onsets of storage seasons have been identified and their stochastic properties examined. As discussed in Thomson (2014), the seasonal regime switching model is readily simulated, allowing a variety of simulation-based methods to be considered for improved risk and scenario forecasting, and a better understanding of the seasonal dynamics of weekly hydro storage, particularly when storage is low.

Although the general structure of the SH model remains unchanged, the 2009 Ministerial Review of Electricity Market Performance has led to changes in the dynamics of the SH model post 30 September 2009. Two optimal models were identified by maximum likelihood with one, the absolute maximum, strongly influenced by the post-2009 data resulting in state mean levels that were higher than those determined in Thomson (2014) and more in keeping with the contracted scale of the post-2009 data. The other optimal model was dominated by the pre-2009 data and produced parameters that were very similar to those determined in Thomson (2014). The differences between the two models reflects their ability to handle the structural break and the contracted scale of the post-2009 storage data. In practice these differences might be alleviated by re-scaling the pre-2009 data appropriately (contracting its scale using the thresholds estimated in Section 3.2) with the aim of producing an optimal model with state means that are homogeneous over the entire data set. However, this re-scaling is unlikely to remedy any changes to the state dynamics (transition probabilities) caused by the structural break.

Subsequent analysis was based on the optimal model dominated by the pre-2009 data.

This choice preserves the classifications and analyses of Thomson (2014) and uses them to check for any change in dynamics post 30 September 2009. There is clear evidence of differences. In particular, the low storage season is less persistent in the post-2009 period with shorter sojourns and the probability of a transition from extreme to intermediate storage is higher in the post-2009 period. The analysis of storage season onsets also points to later low season onsets, earlier high season onsets and shorter low season sojourns post the 2009 Ministerial Review of Electricity Market Performance. These results are consistent with greater risk aversion to low storage in the post-2009 period. However these observations are at best indicative and will need to be reassessed following any adjustment to the pre-2009 data such as re-scaling.

A preliminary exploration was undertaken of the relationship between spot price and storage within the four storage seasons identified by the SH model. In essence, a switching regression model was fitted between transformed weekly average spot prices and weekly average storage levels. The regression relationships fitted were, in general, much weaker than those for the PH model since most of the price-storage relationship is now built into the mean structure of the switching regression model. Indeed, within most seasonal states there was little significant price-storage correlation with the exception of the extreme secondary storage state in the high storage season where there was significant negative correlation. However, the fitted regressions showed some non-linearity indicating that better results might be obtained using tranformed storage data. While the fit of the switching regresssion model is reasonable, it is not quite as good as the PH model which performs slightly better in terms of RMSE. The switching regression residuals also have much stronger residual seasonality than the PH model residuals. This is not unexpected since the switching regression model is based on dynamic storage seasons that are a function of hydro storage levels alone, whereas the PH regression model is based on static seasonal patterns that reflect seasonal demand for electricity in addition to seasonal storage and other possible covariates. Despite this limitation, the switching regression model based on storage seasons manages to provide a competitive and informative view of the relationship between price and storage.

The SH model provides a simple, yet flexible, stochastic framework within which to examine weekly hydro storage data and better understand its variability. As noted earlier, its open informative structure lends itself to forecasting and simulation-based scenario risk assessment. However modifications to the model are needed to account for the structural break caused by the 2009 Ministerial Review of of Electricity Market Performance and other shortcomings identified. The price-storage model will also need to be augmented to include conventional seasonality as well as the switching storage seasons. These and other issues remain topics for further research and development.

Acknowledgements

The reviewer is grateful to Doug Watt (Electricity Authority) and Anthea Jiang (Electricity Authority) for advice, guidance and the provision of background information provided over the course of the analysis.

References

Bruce, A., Jurke, S. and Thomson, P. (1994) Forecasting load-duration curves. *Journal of Forecasting* **13**, 545-559.

Carey-Smith, T., Sansom, J. and Thomson, P.J. (2013) A hidden seasonal switching model for multisite daily rainfall. Submitted for publication.

Cleveland, W.S., Grosse, E. and Shyu, W.M. (1992) Local regression models. Chapter 8 of *Statistical Models in S* (eds J.M. Chambers and T.J. Hastie), Wadsworth & Brooks/Cole, California.

Fenton, C.P., Nance, P.K., Speaker, S.C., Jansen, M.J. and Brown, J.G. (2011) Subzero commodity prices: why commodity prices fall through the zero bound and where and how it could happen in 2011. J.P. Morgan Global Commodities Research. Available at SSRN: http://ssrn.com/abstract=1922787.

Harte, D.S., Pickup, M.L. and Thomson, P.J. (2004) Stochastic models for hydro catchment inflows; an exploratory analysis. Report commissioned by the New Zealand Electricity Commission.

Harte, D.S. and Thomson, P.J. (2006) Development of stochastic models for hydro catchment inflows. Report commissioned by the New Zealand Electricity Commission.

Harte, D.S. and Thomson, P.J. (2007) Hidden Markov models for New Zealand hydro catchment inflows: a preliminary analysis. Report commissioned by the New Zealand Electricity Commission.

MBIE (2009) *Ministerial Review of Electricity Market Performance.* `http://www.mbie.govt.nz/info-services/sectors-industries/energy/previous-reviews-consultations/review-of-the-electricity-market-2009`

Paine, S. and McConchie, J. (2010) Lake generation potential history. Report prepared by Opus International Consultants Ltd for the New Zealand Electricity Authority. http://www.ea.govt.nz/industry/monitoring/cds/hydro-lake-storage

R Development Core Team (2004) *R: A language and environment for statistical computing.* R Foundation for Statistical Computing, Vienna. http://www.R-project.org

Thomson, P.J. (2013) An exploratory analysis of the relationship between electricity spot price and hydro storage in New Zealand (2013). Report commissioned by the New Zealand Electricity Authority.

Thomson, P.J. (2014) A seasonal regime switching model for South Island hydro storage (2014). Report commissioned by the New Zealand Electricity Authority.

Tipping, J.P., Read, E.G. and McNickle, D.C. (2004) The incorporation of hydro storage into a spot price model for the New Zealand electricity market. Presented at the Sixth European Energy Conference: Modeling in Energy Economics and Policy, Zurich.

Venables, W.N. and Ripley, B.D. (2002) *Modern Applied Statistics with S (Fourth Edition).* Springer, New York.