ICRC

independent competition and regulatory commission

Water demand forecasting

Final technical paper

Report 2 of 2015, April 2015

The Independent Competition and Regulatory Commission is a Territory Authority established under the *Independent Competition and Regulatory Commission Act 1997* (the ICRC Act). The Commission is constituted under the ICRC Act by one or more standing commissioners and any associated commissioners appointed for particular purposes. Commissioners are statutory appointments and the current Commissioners are Senior Commissioner Malcolm Gray and Commissioner Mike Buckley. We, the Commissioners who constitute the Commission, take direct responsibility for delivery of the outcomes of the Commission.

We have responsibilities for a broad range of regulatory and utility administrative matters. We have responsibility under the ICRC Act for regulating and advising government about pricing and other matters for monopoly, near-monopoly and ministerially declared regulated industries, and providing advice on competitive neutrality complaints and government-regulated activities. We also have responsibility for arbitrating infrastructure access disputes under the ICRC Act. In discharging our objectives and functions, we provide independent robust analysis and advice.

Our objectives are set out in section 7 of the ICRC Act and section 3 of the *Utilities Act 2000*.

Correspondence or other inquiries may be directed to the Commission at the following addresses:

Independent Competition and Regulatory Commission
PO Box 161
Civic Square ACT 2608

Level 8
221 London Circuit
Canberra ACT 2601

We may be contacted at the above addresses, by telephone on (02) 6205 0799, or by fax on (02) 6207 5887. Our website is at www.icrc.act.gov.au and our email address is icrc@act.gov.au.

Contents

List of tables

List of figures

1 Introduction

1.1 Background and purpose of this paper

Given the importance of the water volumes forecast for setting water prices and Icon Water's revenue stream, the Commission published a technical paper on 20 January 2015 setting out its working conclusions on forecast water sales volumes.[1] The Commission did not receive any submissions on its January 2015 technical paper.

The Commission's development of an improved water sales forecasting model may inform the regulator's determination of forecast water sales volumes for the next regulatory period. For this reason the Commission has published this follow-up paper which describes a number of refinements to the modelling approach presented in the first paper and presents and set of forecast volumes from the revised forecasting model.

1.2 Paper structure

The remainder of the paper is structured as follows:

- Chapter 2 presents the Commission's preferred model specification, describes the refinements to the releases forecasting approach and presents updated releases and billed consumption forecasts.
- Chapter 3 discusses the refinements to the Commission's method of apportioning total sales volume into tier 1 and tier 2 sales and presents updated tier proportions based on the revised billed consumption forecasts.
- Chapter 4 concludes the paper.

[1] ICRC (2015). Available at http://www.icrc.act.gov.au/water-and-sewerage/inquiries-and-investigations/.

2 Forecasting aggregate water volumes

2.1 January 2015 technical paper

2.1.1 Introduction

In its January 2015 technical paper, the Commission provided evidence of a new and stable relationship between water sales and climate variables post July 2006. The paper also noted that starting from that point, we now have data series spanning some eight and a half years, including years presenting a variety of climate experience. On the strength of this, the Commission developed a seasonal ARIMA model that utilises daily releases and climate data to provide daily release forecasts.[2]

2.1.2 ARIMA model

The paper presented the preferred ARIMA model, utilising data to 5 January 2015, with the following characteristics:

- seasonal ARIMA (1,0,2)(1,1,1)[7] specification;
- maximum temperature entered at lags of 0, 1, 5 and 7 days and rainfall entered at lags of 0 through 6 days; and
- no water restrictions dummy variables as they were not significant, either in combination or separately.

The estimated model parameters are shown Figure 2.1. All parameters in the model were significant and the model returned an Akaike Information Criterion (AIC) of 22,753.

[2] The terminology to describe the various model forms that can utilise an ARIMA specification is not standardised, see Hyndman (2010a). A more precise description of the January model would be a regression equation relating dam releases to climate variables with an error process described by an ARIMA model.

Figure 2.1 January 2015 preferred ARIMA equation results

```
Series: Releases11TS
ARIMA(1,0,2)(1,1,1)[7]

Coefficients:
         ar1      ma1      ma2     sar1     sma1  Temp11  Temp10   Temp6
      0.8546  -0.1711  -0.0557  0.0938  -0.8628  0.6089  1.6586  0.2964
s.e.  0.0169   0.0249   0.0241  0.0272   0.0187  0.0542  0.0546  0.0548
       Temp4   Rain11   Rain10    Rain9    Rain8    Rain7    Rain6    Rain5
      0.1359  -0.4028  -0.4964  -0.3248  -0.2664  -0.1767  -0.1311  -0.0733
s.e.  0.0549   0.0313   0.0347   0.0353   0.0358   0.0353   0.0346   0.0309

sigma^2 estimated as 91.93:  log likelihood=-11359.4
AIC=22752.8   AICc=22753   BIC=22855.43

Training set error measures:
                      ME     RMSE      MAE        MPE     MAPE      MASE
Training set 0.001038763 9.566532 6.812764 -0.4255724 5.392724 0.7764008
                      ACF1
Training set -6.768085e-05
```

Source: ICRC, 2015: 38.

The fitted dam release values using the preferred model showed a very close match to the observed dam releases over the period from July 2006 to December 2014, as shown in Figure 2.2.

Figure 2.2 January 2015 ARIMA model fitted versus actual dam releases observations

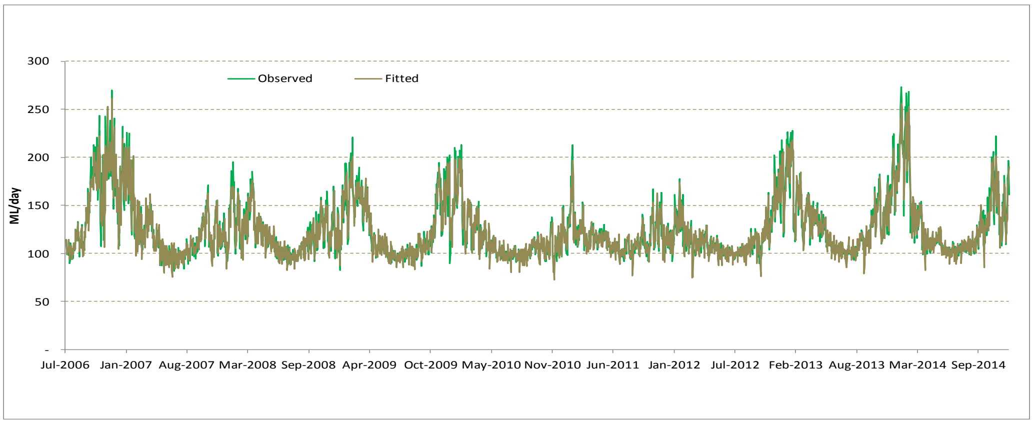

Source: ICRC, 2015: 44.

Similarly, the in sample forecasts for billed consumption showed a very close match, lying within a range of ±2 per cent of actual annual billed consumption with many within one per cent, as shown in Figure 2.3.

Figure 2.3 January 2015 billed consumption observed versus modelled values

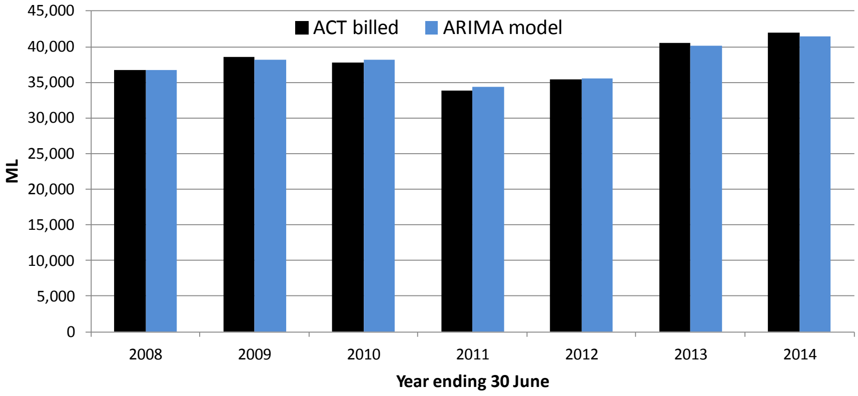

Source: ICRC, 2015: 43.

The January 2015 paper presented a three step forecasting process:

- forecast daily releases from assumed climate conditions using the ARIMA model;
- aggregate daily data into monthly totals comparable in coverage to the monthly billed consumption date; and
- apply the historical average ratio of billed consumption to releases to annual releases to calculate forecast billed consumption.

In order to create estimates of annual billed consumption, the Commission first aggregated daily releases to monthly figures in a way that reflects the logic of Icon Water's 91 day billing cycle. Using historical data over the period 1999 to 2014, the Commission applied the overall average ratio of billed consumption to releases to the sum of monthly releases. Utilising this approach, the Commission presented the forecasts set out in Table 2.1.

Table 2.1 January 2015 ARIMA model annual forecasts (ML)

Financial year	Dam releases	Billed consumption
2015	47,341	39,848
2016	46,981	39,800
2017	46,846	39,636

ICRC, 2015: 44.

2.2 April 2015 position

2.2.1 Introduction

When another month of releases data became available the Commission re-estimated the model with the releases and climate dataset updated by 29 days to 4 February 2015.

The model produced slightly different parameter estimates, as would be expected with the addition of even a small number of new data points. Nonetheless, the releases forecasts it produced were substantially different, up to 1.8 GL per year, from those presented in January 2015.

The volatility in the model forecasts prompted the Commission to revisit the three step model building process, with a view to first identifying the reason for the volatility, and then developing a more satisfactory model.

2.2.2 Model identification

Stationarity and cointegration

A review of the statistical literature around stationarity and the application of ARIMA methods to equations of the kind used in the January 2015 model, suggests a potential explanation for the divergent forecasts obtained from that model. When Box and Jenkins introduced their ARIMA methodology for forecasting and control, it was predominately a method of using a time series' own past to predict its future from the correlations of the current value with past values.[3] In order to be able estimate the parameters in the postulated relationship between the series and its own past, the series needed to be stationary. A stationary time series has the property that its statistical characteristics such as the mean and the autocorrelation structure are constant over time.[4] The Box-Jenkins methodology relied on differencing non-stationary series until stationarity was achieved and then estimating the parameters of the postulated relationship using the differenced series.

This approach works well as long as the series to be forecast is related only to itself and an error term that is assumed to be stationary. In this case the only possible source of any non-stationarity is in the structure of the postulated model itself. The Box-Jenkins methodology modelled this source of non-stationarity and used differencing to eliminate it from the series before estimation was attempted.[5] When another series is introduced into the model, any non-stationarity in the series to be forecast can come from another source, one to which differencing may not be the appropriate response.

---

[3] See Box and Jenkins (1970). Although the book contains a treatment of relationships between different times, this was not a focus of early attention and the implications of extending the approach in this way were not fully appreciated until the work of Engle and Granger, described below.

[4] In simple terms, statistical estimation relies on averaging out the random disturbances in data to focus on the underlying relationships. This averaging will only be effective if the process generating the random disturbances is stable. Consider trying to estimate the proportion of black balls in a jar by using the proportion of black balls in a sample taken from the jar, when an unknown number of balls are being added to or removed from the jar in between drawings.

[5] The non-stationarity was modelled by introducing an appropriate number of unit roots into the distributed lag function relating the series to its own past. The unit roots were eliminated by redefining the relationship as being between the appropriate difference of the series and the past of those differences.

Granger (1986) put forward the theory that certain pairs of economic variables may not diverge from each other by too great an extent, at least in the long-run, even though they may drift apart in the short-run or according to seasonal factors.[6] Alternatively:

> ... a vector of time series, all of which are stationary only after differencing, may have linear combinations which are stationary without differencing. In such a case, those variables are said to be [cointegrated] ...[7]

Consider the simple equation:

$$y_t = \beta x_t + z_t$$

If y and x are non-stationary but cointegrated then z, the stochastic part of the equation, may be stationary. If this is so, then there is no need to difference the equation. In fact, doing so is inappropriate.

Engle and Yoo (1987) coined the phrase 'over-differenced' to describe the situation where cointegrated variables are unnecessarily differenced, concluding that:

> When a forecasting model is needed for [a cointegrated] time series, a vector autoregressive model in differences is inappropriate. This is because, even though the residuals may appear to be white, such a model suffers misspecification and the forecasts will diverge from each other.[8]

Looking at it from another angle, Hyndman and Athanasopoulos (2012) state that in the case where non-stationary variables are co-integrated but the data has not been differenced 'then the estimated coefficients are correct'.[9]

The ARIMA modelling process, as applied by the Commission in the January 2015 paper, fits a regression model with ARIMA errors, that is, z in the simple model above is generated by an ARIMA process. If cointegration is present, it is possible that while the variable being modelled, y in the above simple equation, may require differencing to be stationary, the error process may be stationary without the need for any differencing. In other words, even though individual variables may be non-stationary without differencing, the variables when modelled in combination can produce a system with stationary errors.

Granger's theory was primarily aimed at macroeconomic variables such as short and long term interest rates, where typically economic theory proposes forces that tend to keep such series together. Even though we are not dealing with similar economic variables here, the Commission considered it worthwhile to test whether the first-order

---

[6] Granger, 1986: 213.

[7] Engle and Yoo, 1987: 143.

[8] Engle and Yoo, 1987: 158.

[9] Hyndman and Athanasopoulos, 2012: 9.1 Dynamic regression models.

seasonally differenced component applied in the January 2015 model was a case of 'over-differencing' and the cause of the divergent forecasts.

Assessing the need for differencing

The first step was a closer examination of stationarity of the releases and climate data series to establish whether differencing is actually required for stationarity of the system rather than just stationarity of the individual time series data.

An examination of the ordinary least squares (OLS) residuals of the dam releases time series (Releases) regressed against various lags of maximum temperature (Temp) and rainfall (Rain) showed no indication of a random walk process. Rather, the residuals appeared stationary with no sign of any trend or drift, as shown in Figure 2.4. This suggests that differencing may not be required.

Figure 2.4 OLS residuals Releases against Temp and Rain

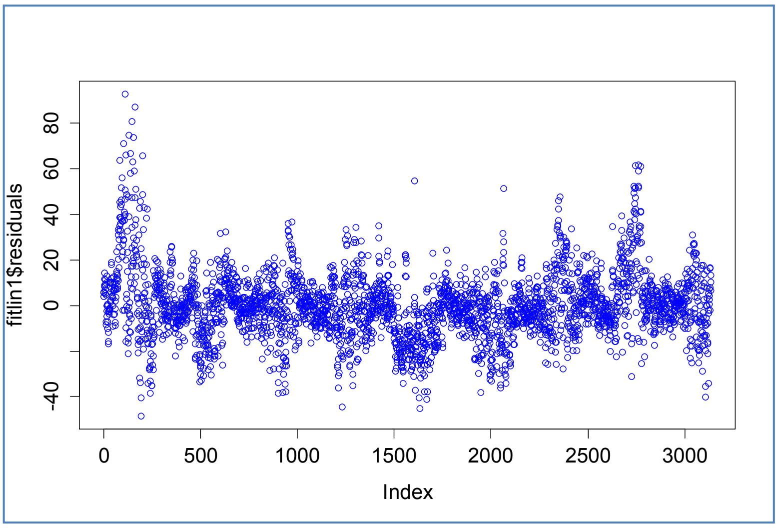

Source: ICRC analysis R Studio output.

The Commission then ran the 'ndiffs' and 'nsdiffs' functions in the R 'forecast' package which use unit root tests to determine the number of differences (seasonal for 'nsdiffs') required for a particular time series to be made stationary. Applying these functions to the releases and climate data, including squares of the climate data, showed that:

- no series requires seasonal differencing, either weekly or annually;
- the releases data (Releases) requires differencing once; and
- the temperature data (Temp) does not appear to require differencing in the estimation period, but seems to require differencing once over the 47 year period from 1967 suggesting that the stochastic process generating temperature is non-stationary of order 1.

Testing for cointegration

Since both Releases and Temp appear to be non-stationary of order 1, the next step was to check for cointegration. The Commission used the Phillips and Ouliaris unit root test function ('co.pa') in the 'urca' R package to test for any evidence of cointegration between the Temp and Releases time series. In all the cases run, the null hypothesis of no cointegration is clearly rejected for cointegration between Releases and both Temp and the square of Temp (see Table 2.2 for example results). This indicates that regression analysis should be undertaken without differencing the cointegrated series.

Table 2.2 Cointegration test results Releases against Temp

Coefficient	Estimate	Standard error	t-ratio	p-value
zr1	0.837198	0.008075	103.67	<0.00001
zr2	0.936635	0.046886	19.98	<0.00001

Source: ICRC analysis.

Prewhitening to identify lags

Cryer and Chan (2008) note that with strongly autocorrelated data, such as that we are dealing with here, it is difficult to distinguish between the linear associations between x and y, in the simple equation above, and their autocorrelation. In these circumstances, the cross-correlation function (CCF) between the variables can be misleading. Cryer and Chan recommend using prewhitening to check whether a relationship actually exists and as an aid to identifying what lags of x should be used in the regression.[10]

Prewhitening involves fitting an ARIMA model for the x series sufficient to reduce the residuals to white noise. The x series are then filtered with this model to get the white noise residual series. The y series is then filtered with the same model and the filtered y result is cross-correlated with the x filtered series. The CCF produced is then used to identify possible lag terms to use in the regression.

Before moving to the regression stage, a prewhitening analysis was undertaken using the 'prewhiten' function in the 'TSA' R package in order to identify what lags of Temp and rainfall (Rain) to include in the regression against Releases. Figure 2.5 shows the CCF graphs produced for Temp, Temp squared, Rain and Rain squared after the prewhitening process.

The possible lags indicated by this analysis are:

- Temp (0, 1, 3, 5, 12)
- Temp squared (0, 1, 3, 5, 12)
- Rain: (0 - 8)
- Rain squared: (0 - 3, 6).

[10] Cryer and Chan, 2008: 265.

Figure 2.5 Prewhitening CCF graphs

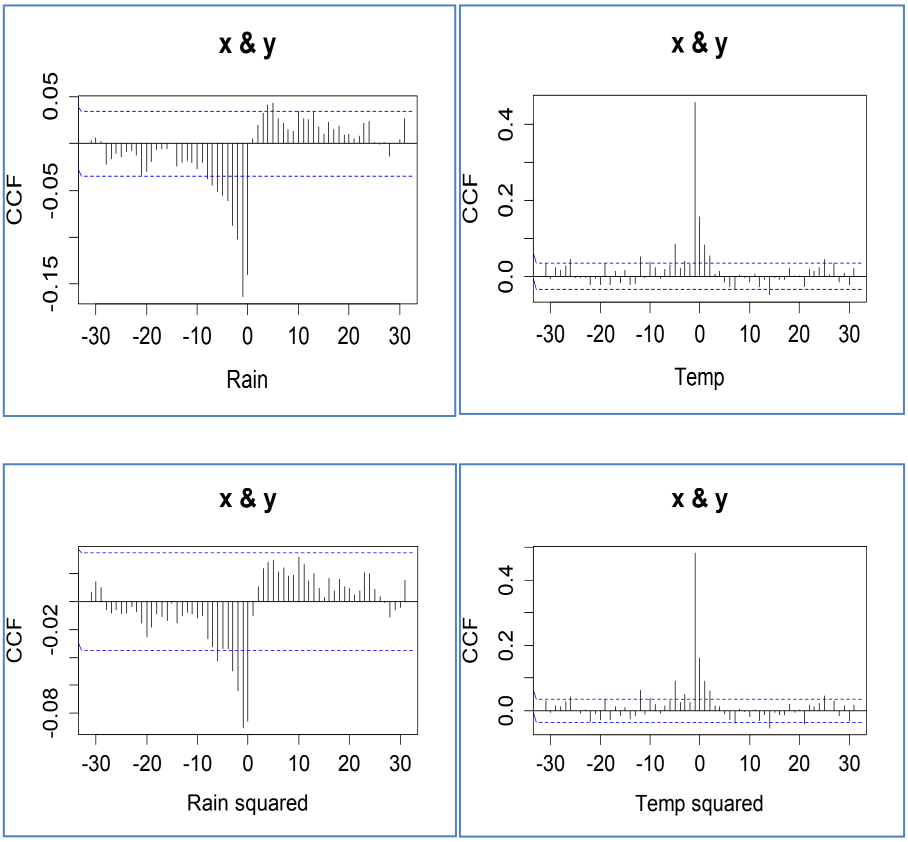

Source: ICRC analysis, R Studio output.

Running alternative model specifications

Having identified potential lags, the next step was running multiple alternative model specifications, all of them without differencing, to identify the preferred model. It is important to note that prewhitening process described above was only used for identification purposes. The unadjusted time series was used to estimate the preferred regression equation.

The preferred equation was identified with reference to minimising the Akaike Information Criterion (AIC) and root mean square error (RMSE), the significance of the equation coefficients, unit root tests to check for stationarity and the stability of the model forecasts.

The AIC is a statistical measure for model selection. All else being equal, the model with the lower AIC value is to be preferred as the better model. This statistic rewards goodness of fit and also includes a penalty for increasing parameter numbers.

The RMSE is the square root of the mean squared error and, when adjusted for degrees of freedom for error, is also referred to as the estimated white noise standard deviation as shown below:

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(e_i^2)}$$

where:

- n is the number of explanatory variables; and
- e is the error term.

For reasons explained in the January 2015 technical paper, the Commission elected to use the time series data from 1 July 2006 onwards, in estimating the preferred equation. In addition to the Releases, Temp and Rain data, the Commission also included daily, monthly and water restrictions dummies in the analysis. The initial work was carried out using data up to 4 February 2015. The data set was subsequently updated to 28 February 2015 and all the results reported here relate to that more recent dataset.

Preferred model

The preferred model was identified as a having an error term generated by a seasonal ARIMA (1,0,2)(1,0,1)[7] model of the general specification shown below:

$$(1-\phi_1 B)(1-\Phi_1 B^7)z_t = (1+\theta_1 B+\theta_2 B^2)(1+\Theta_1 B^7)\varepsilon_t$$

where ε_t is white noise. The releases data exhibited a strong recurring weekly pattern so the ARIMA seasonality cycle was set to seven days. The following climate lags and dummy variables are included in the specification:

- Temp (0, 1, 3, 12), Temp squared (0, 1, 3, 5, 12), Temp square root (1);
- Rain (0, 1), Rain squared (2, 3), Rain square root (0 - 8), Rain cubed (0 - 2), Rain square root (6);
- daily dummies (Sunday, Monday, Tuesday, Wednesday, Thursday, Friday);
- water restrictions dummy (Stage 3); and
- a Fourier series term.

The stage 3 water restrictions dummy was included to ensure that any impact restrictions had was captured. The dummy variable was not significant at the 10 per cent level, but did give a very marginal improvement in the AIC. Because, as Hyndman (2011) points out, AIC is the more reliable guide to the value of a variable in a forecasting equation, the restrictions dummy was retained in the forecasting equation.

Even though the monthly dummies were excluded as the analysis progressed, the Commission remained concerned about the ARIMA model not having the capability to account for more than one seasonal pattern in the equation. Following Hyndman (2010), the Commission adopted a Fourier series approach where the seasonal pattern (annual in this case) is modelled using Fourier terms and the short-term series dynamics (weekly in this case) is modelled in the ARIMA error term.

A Fourier series is a means of representing a periodic function as a sum of sine and cosine waves, and is modelled as shown below using the 'fourier' function in the R 'forecast' package:

$$y_t = a + \sum_{k=1}^{K} \left[\alpha \sin\left(\frac{2\pi kt}{m}\right) + \beta \sin\left(\frac{2\pi kt}{m}\right) \right] + N_t$$

where:

- N_t summarises all the other variables in the model, including the ARIMA error term.

The value of k was set at 1 and was chosen by minimising the AIC. The introduction of the Fourier series significantly improved the AIC of the preferred model.

The Commission chose not to include a population variable because there is no trend in the Releases data over the estimation period. As the Commission noted in the January 2015 paper and as is illustrated in Figure 2.7, over the period 1998 to 2014, there has been a significant, downward trend in Releases. The recent work by Cardno for the Industry Panel illustrates that adopting a population based approach can make the formulation of a satisfactory forecasting equation more rather than less difficult.[11]

2.2.3 Model estimation

The estimated model parameters, using data from 1 July 2006 to 31 March 2015, are shown in Figure 2.6. The parameters in the model are significant at the 95 per cent level with the exception of Rain6 cubed (Rain6cbe), stage 3 restrictions dummy (dumS3) and S1-365. The model returns an AIC of 22,729.33 and a RMSE of 8.48.

[11] Cardno (2014).

Figure 2.6 Preferred ARIMA equation results

```
> summary(fit)

Call:
arima(x = Releases13TS, order = ord, seasonal = sord, xreg = xM)

Coefficients:
         ar1      ma1      ma2     sar1     sma1  intercept    Temp0    Temp1    Temp3
      0.8477  -0.2595  -0.0606   0.8749  -0.6732   114.0238  -1.1414  -6.2346  -0.4559
s.e.  0.0185   0.0263   0.0246   0.0171   0.0270    18.1614   0.1929   1.8172   0.1955
       Temp12  Temp1sqrt  Temp0sq  Temp1sq  Temp3sq  Temp5sq  Temp12sq   Rain0   Rain1
      -0.5674    23.6927   0.0323   0.1095   0.0102   0.0055    0.0155  0.4250  0.5700
s.e.   0.1903    10.7243   0.0040   0.0146   0.0040   0.0010    0.0039  0.0627  0.0678
      Rain0sqrt  Rain1sqrt  Rain2sqrt  Rain3sqrt  Rain4sqrt  Rain5sqrt  Rain6sqrt
        -4.8039    -5.7933    -2.1850    -2.0667    -1.2677    -1.0521    -0.5754
s.e.     0.3066     0.3258     0.1719     0.1718     0.1335     0.1254     0.1378
      Rain7sqrt  Rain8sqrt  Rain2sq  Rain3sq  Rain6cbe  dumDM1   dumDM2  dumDM3  dumDM4
        -0.5562    -0.5625   0.0036   0.0022         0  6.9819  12.9180  7.0624  5.6953
s.e.     0.1191     0.1248   0.0008   0.0008       NaN  1.1080   1.2983  1.3563  1.3583
      dumDM5  dumDM6    dumS3  S1-365     C1-365
      4.6262  4.9172  -4.6341  0.9450   -14.3627
s.e.  1.3011  1.1111   2.7901  1.9172     1.8527

sigma^2 estimated as 71.86:  log likelihood = -11324.66,  aic = 22729.33

Training set error measures:
                        ME     RMSE      MAE       MPE     MAPE      MASE         ACF1
Training set -0.006455886 8.476998 6.088727 -0.407996 4.825358 0.6862742 0.0005632147
```

Source: ICRC analysis R Studio output.

The coefficient of the stage 3 restrictions dummy (dumS3) suggests that it has a negative effect on releases of about 4.6 megalitres (ML) per day. This amounts to about 3.7 per cent of average daily releases of about 125 ML per day over the estimation period. Lest the small size of the estimated coefficient should be interpreted as implying that restrictions had no effect, it needs to be remembered that Stage 3 restrictions came into effect within six months of the beginning of the estimation period and that all, except permanent water conservation measures, were removed three and a half years before the end of the estimation period. This means the coefficient is measuring the degree of bounce-back following the ending of restrictions and not the impact of imposing restrictions. That this effect is small and poorly defined statistically supports the Commission's hypothesis that water consumption behaviour has been stable since 1 July 2006.

As shown in Figure 2.7, the estimation period starts after a paradigm shift in the relationship between water sales and climate variables as discussed in the January 2015 technical paper.

Figure 2.7 Observed annual releases 1999 to 2014

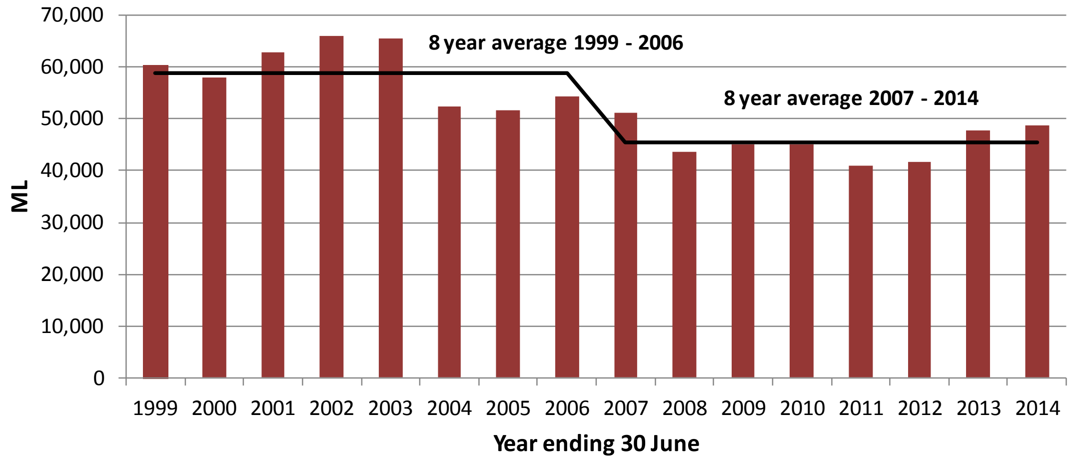

Source: ICRC analysis.

2.2.4 Diagnostic checking

The residuals of the preferred model were then examined to see if they appear to be white noise – that is, they have no remaining autocorrelations.

Visual inspection

Figure 2.8 shows the autocorrelation function (ACF) of the residuals against Temp and Rain. As was the case in January 2015, the majority of the spikes in the top left chart for both variables are within the significance limits, suggesting that the residuals are white noise.

Figure 2.8 ACF temperature and rain

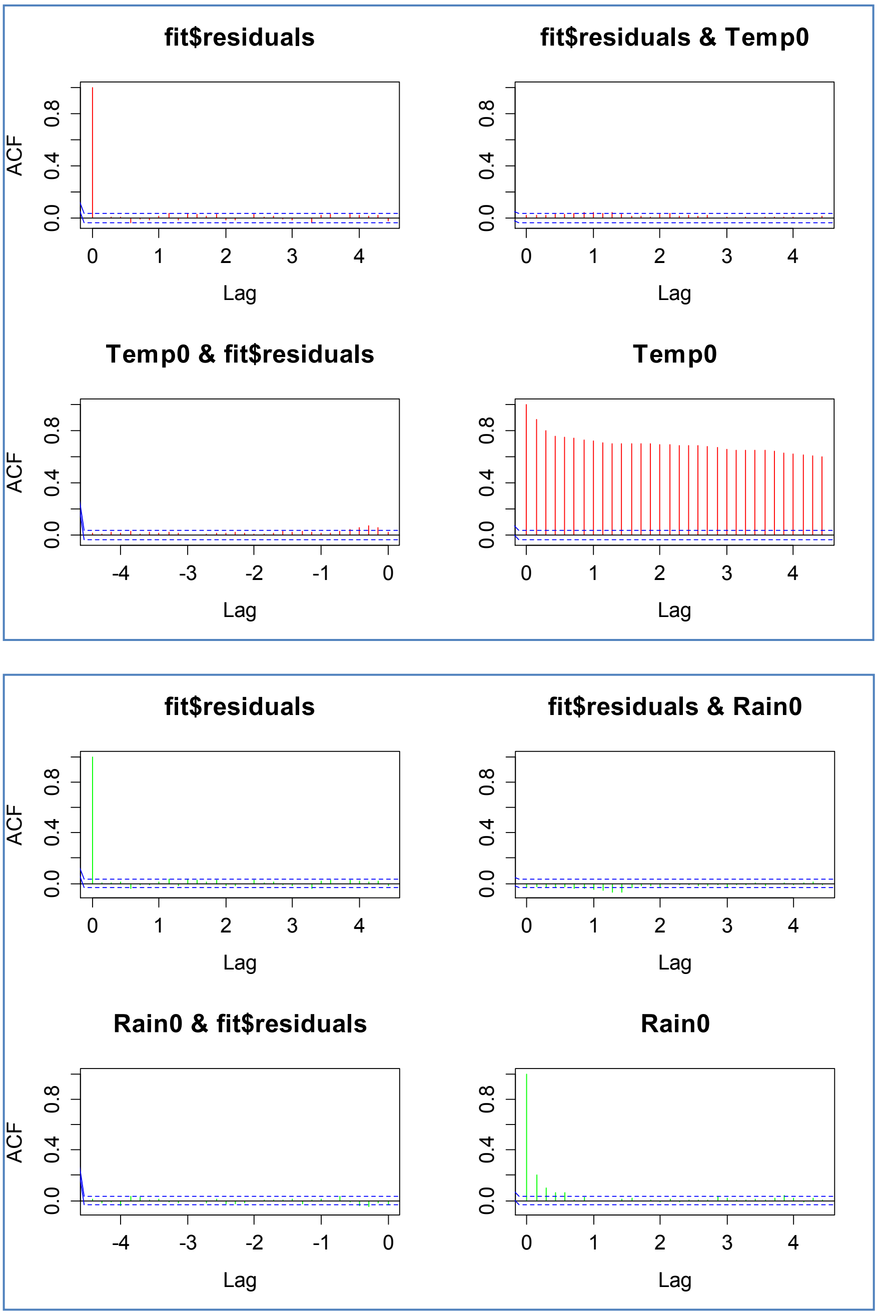

Source: ICRC analysis R Studio output.

Figure 2.9 shows the raw periodogram which represents an estimate of the spectral density of the residuals for the updated model. The figure is reflective of a Gaussian white noise signal, which is characterised by a normal distribution and zero mean, once again suggesting that there is no remaining residual autocorrelation.

Figure 2.9 Raw periodogram

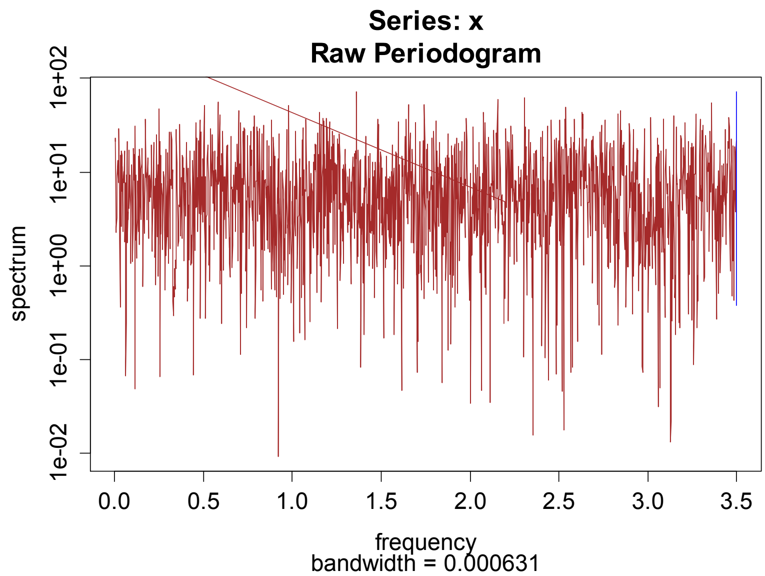

Source: ICRC analysis R Studio output.

Statistical tests

The methods used to estimate ARIMA models rely on the error term being serially uncorrelated or white noise. The Box-Ljung test returned a test statistic of 77.57 with a p-value of 0.0013. This suggests that the null hypothesis of white noise residuals should be rejected at the 95 per cent level. The KPSS test also resulted in the null hypothesis being rejected with a p-value of 0.0255.

As was the case for the Commission's January 2015 model, when interpreting these results it needs to be borne in mind that, with a sample size in excess of 3,000, the power of the test is going to be very high, that is, it will be capable of detecting even small deviations from pure, white noise residuals. This is seen in the autocorrelation function of the residuals shown in Figure 2.8 above. Here the confidence interval is very small and autocorrelation estimates as low as 0.05 are showing as significant.

2.2.5 Fitted results

Dam releases

The fitted dam release values using the preferred model are a very close match to the observed dam releases over the period from July 2006 to March 2015 as shown in Figure 2.10.

Figure 2.10 April 2015 ARIMA model fitted versus actual dam releases observations

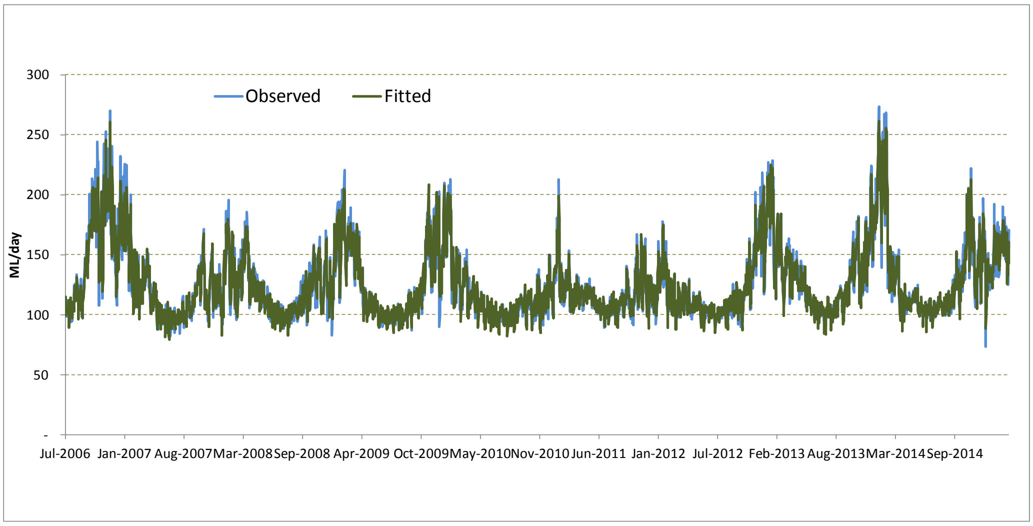

Source: ICRC analysis.

Billed consumption

In order to create estimates of annual billed consumption, the daily releases are first aggregated to monthly figures in a way that reflects the logic of Icon Water's 91 day billing cycle. A regression estimate of the ratio of annual billed consumption to annual releases over the period 1999 to 2014 is then calculated and applied to the sum of monthly releases. The estimated ratio over this period is approximately 85 per cent.

Before using this procedure to generate forecasts, the Commission tested the in sample performance using daily fitted values from the ARIMA model. Figure 2.11 shows the results of this exercise over the 7 years from 2007–08 to 2013–14. The in sample forecasts lie in a range of ±2 per cent of actual annual billed consumption with many within one per cent, a very close match. The in sample forecast for 2013–14 was 41.6 GL, a little lower than the observed billed consumption of 42.0 GL.

Figure 2.11 April 2015 billed consumption observed versus modelled values

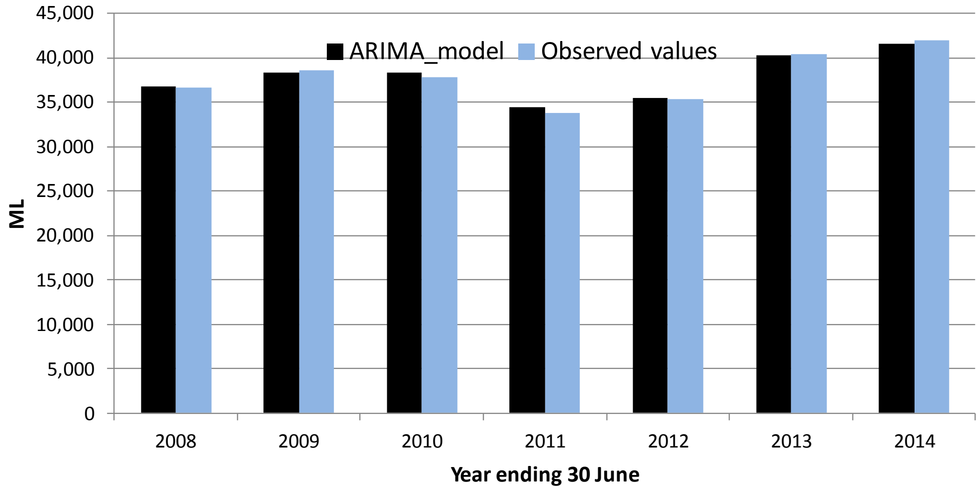

Source: ICRC analysis.

2.2.6 Forecasts

The process of forecasting annual water sales over the forecast period from 1 April 2015 to 30 June 2017 involves three steps:

- first, forecast daily releases from assumed climate conditions using the ARIMA model;
- second, aggregate this data into monthly totals comparable in coverage to the monthly billed consumption date; and
- third, apply the regression estimate of the ratio of billed consumption to releases to annual releases to calculate forecast billed consumption.

A number of changes have been made to the first step from that presented in the January 2015 paper. The climate forecast is now obtained by averaging 73 separate climate scenarios derived from actual climate over a succession of 2½ year periods over a period of 75 years. Given the nonlinearities now present in the model, this procedure provides a better estimate of expected releases over the forecast period. In addition, a temperature trend factor is now applied to the climate scenarios to take into account the rising trend in maximum temperature. No temperature adjustment was made to the forecast climate scenario applied in the January 2015 paper.

Figure 2.12 shows the annual average maximum daily temperature at Canberra Airport over the period 1939–40 to 2013–14. The chart suggests a declining trend in temperature from the start of this period until 1955–56, followed by an increasing trend from then until 2013–14.

Figure 2.12 Canberra Airport annual average maximum daily temperature

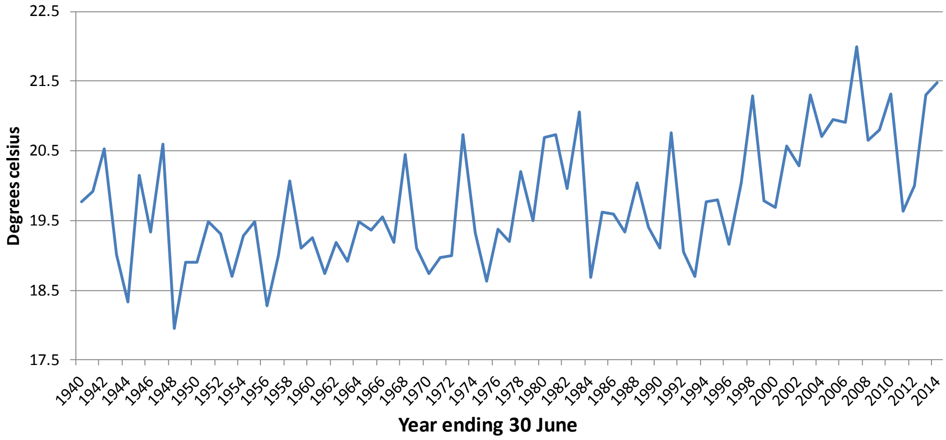

Source: BOM data.

Running linear regressions of time against temperature confirms a significant downward trend over the first period and a highly significant upward trend over the second period. Table 2.3 shows the regression results.

Table 2.3 Regression results Temp against time

	Estimate	Standard error	t-ratio	p-value
1940 to 1956				
Intercept	19.6600	0.1779	110.5000	< 0.0001
time	-0.000129	0.0000	-2.6000	0.0094
1956 to 2014				
intercept	18.8800	0.0940	200.9700	< 0.0001
time	0.00009166	0.0000	11.9300	< 0.0001

Source: ICRC analysis.

In order to utilise the full data set from 1940 to 2014 and to provide the maximum number of historical climate scenarios, the deviations from trend in the early period have been measured and added to the trend estimated for the later part of the period and applied in the forecast period.

The dam releases and billed consumption forecasts from the preferred ARIMA specification using the forecasting process described above are presented in Table 2.4 and Figure 2.13.

Table 2.4 April 2015 ARIMA model annual forecasts (ML)

Financial year	Dam releases	Billed consumption
2015	47,297	39,964
2016	46,260	39,301
2017	46,084	39,109

Source: ICRC analysis.

Figure 2.13 April 2015 ARIMA model dam releases forecasts

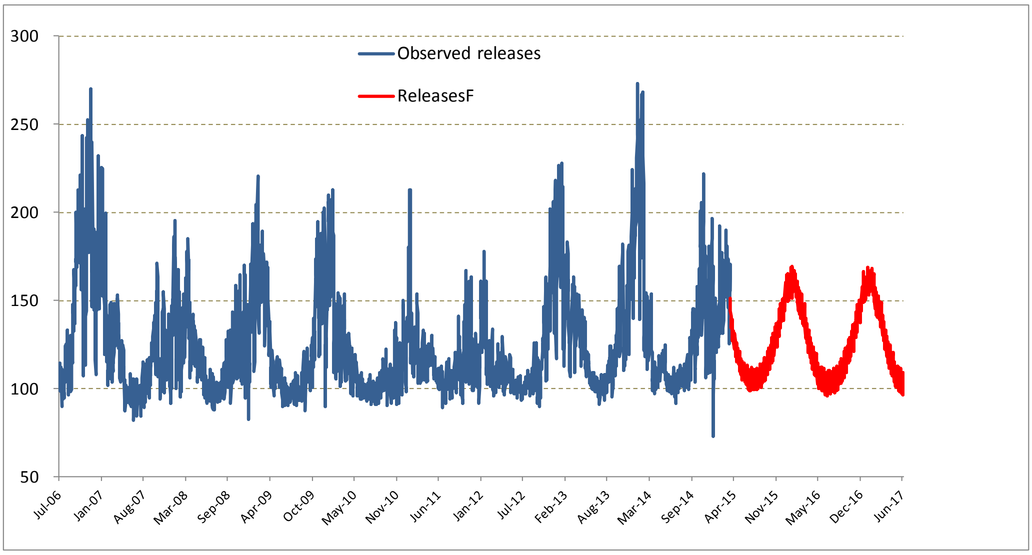

Source: ICRC analysis.

The 2014–15 dam releases forecast of 47.3 GL includes 9 months of observed data, from July 2014 to March 2015. It is instructive to compare the observed values of dam releases and billed consumption for 2014–15 to 2013–14. It is clear from Table 2.5 that both releases and billed consumption in the first 9 months of 2014–15 are significantly down on volumes observed over the same period in 2013–14, with the summer months in particular contributing to this result. On this evidence alone, and with only the 3 cooler months of 2014–15 to go, actual annual 2014–15 water sales are unlikely to greatly exceed the 40.0 GL forecast.

Table 2.5 2013–14 and 2014–15 observed releases and billed consumption (ML)

Month	Releases 2013–14	Releases 2014–15	Difference	Billed 2013–14	Billed 2014–15	Difference
Jul	3,113	3,169	2%	3,084	2,713	-12%
Aug	3,174	3,327	5%	2,514	2,157	-14%
Sep	3,397	3,366	-1%	2,704	3,254	20%
Oct	4,167	3,972	-5%	3,179	2,846	-10%
Nov	4,228	5,081	20%	2,961	2,969	0%
Dec	5,270	4,231	-20%	3,217	3,674	14%
Jan	6,543	4,267	-35%	4,531	3,471	-23%
Feb	5,023	4,253	-15%	4,282	3,841	-10%
Mar	4,034	4,976	23%	4,589	3,660	-25%
Total	**38,949**	**36,643**	**-6%**	**31,331**	**28,584**	**-9%**

Source: Icon Water data.

2.2.7 Model stability

The preferred ARIMA model was tested for stability as part of the model identification process.

First, the ARIMA specification was tested by running multiple variants of the specification, and in particular using a wide range of different Temp and Rain lags and lag transformations. The resulting ARIMA coefficient fit was relatively stable across the range of variants tried.

Second, the preferred ARIMA specification was tested for forecasting stability by running a range of different estimation and forecast periods. The preferred specification provided relatively stable forecasts when subjected to this test.

2.2.8 A note on the forecasts

The forecasts presented in the section are point or mean forecasts. Hyndman (2012) notes that:

> A point forecast is (usually) the mean of the distribution of a future observation in the time series, conditional on the past observations of the time series.[12]

In our case, the point forecast is also conditional on the average forecast climate scenario applied in the model. Importantly, the actual water sales volumes that eventuate in the forecast period will depend on the actual weather patterns experienced. If the weather is hotter and dryer than average, water sales are likely to be higher than forecast and the converse under cool and wet conditions.

[12] Hyndman (2012).

3 Forecasting tier 1 and tier 2 proportions

3.1 Introduction

The final procedure in forecasting water sales is apportioning the forecast annual billed consumption volumes into tier 1 and tier 2 sales. Icon Water applies two volumetric water prices; a tier 1 price from 1–200 kL/year and a tier 2 price for water sold in excess of 200 kL/year.[13] As such, for pricing and revenue purposes, demand forecasts need to be split accordingly.

3.2 January 2015 technical paper

The tier 1 and tier 2 consumption split that the Commission applied in 2013 of 56:44 per cent, which was based on the average consumption profile over the period 2008–09 to 2011–12, was relatively close to the reported 2013–14 actual split.

In its 2015 biennial information return, Icon Water calculated the tier 1 and 2 proportions for its total sales volume forecasts using the following exponential formula:

$$\text{Propn tier 1} = 0.75264 - 0.00811e^{\frac{0.012254 \text{ x total annual consumption (ML)}}{\text{number of annual supply charges}}}$$

In its January 2015 technical paper, the Commission noted that applying Icon Water's formula to the 2013–14 annual water sales and using a total a supply charge number of 162,951, results in a tier 1 proportion of 56 per cent, the same as the Commission's consumption split. The Commission's working conclusion in the technical paper was to retain the current 56:44 consumption split.

3.3 April 2015 position

The Commission undertook further analysis of Icon Water's preferred approach to apportioning total sales volume into the two tiers. Using the data shown in Table 3.1, the Commission analysed the relationship between the average amount of water consumed by each customer per year and the observed proportion of sales falling into the tier 1 category.

[13] The Commission's daily pricing regime allows for 548 L/day, which is approximately 50 kL/quarter and 200 kL/year. This means that tier 2 prices can be incurred even when 200 kL of water has not been consumed over a 12 month period.

Table 3.1 ACT tier 1 and 2 observed water sales volumes and customer numbers

Year	Total sales (ML)	Tier 1 sales (ML)	Tier 2 sales (ML)	No. customers (#)	ML/customer/ year	Observed tier 1 proportion
2008–09	38,179	20,448	17,731	144,165	0.265	**53.56**
2009–10	37,744	21,485	16,259	146,853	0.257	**56.92**
2010–11	33,780	20,906	12,874	149,794	0.226	**61.89**
2011–12	35,393	21,851	13,541	153,256	0.231	**61.74**
2012–13	40,428	23,032	17,396	158,258	0.255	**56.97**
2013–14	41,928	23,759	18,169	162,951	0.257	**56.67**

Source: Icon Water (2014).

A visual examination of the observed data suggests an exponential relationship between the tier 1 proportion and average customer consumption, as shown in Figure 3.1. The Commission estimated a number of equations using the nls (non linear least squares), lm (linear model) and poly (polynomial) functions in R in order to identify the preferred equation that:

- shows the best fit between observed and modelled values;
- displays significant equation parameters; and
- provides sensible modelled values across the range of average consumption values contemplated over the near-term forecasting period, including values that fall outside of the observed range.

Figure 3.1 Observed tier 1 proportion: ML per customer, 2008–09 to 2013–14

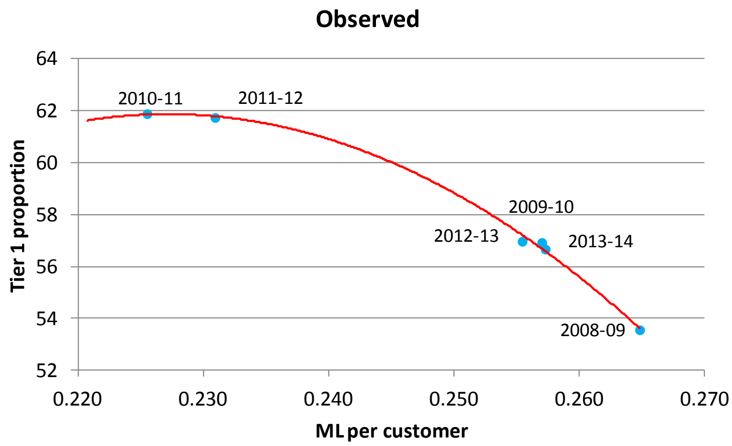

Source: ICRC analysis.

The Commission's preferred equation is of the form $y = \mathrm{c} + \mathrm{ae}^{\mathrm{bx}}$ as follows:

$$y = 63.41564 - 0.00001960964\mathrm{e}^{49.5675\mathrm{x}}$$

where:

- y is the tier 1 proportion of total ACT water sales measured as a proportion of 100 units; and
- x is the average ACT customer consumption per year in ML.

This equation provides a better fit than Icon Water's proposed equation over the observed period, with a total absolute residual of 0.77 compared to 2.85, as shown in Table 3.2.[14]

Table 3.2 Observed versus modelled tier 1 proportions and residuals

Year	Observed	Commission modelled	Icon Water modelled		Commission residual	Icon Water residual
2008–09	53.56	53.56	54.45		0.01	0.89
2009–10	56.92	56.73	56.35		0.19	0.57
2010–11	61.89	62.01	62.41		0.13	0.52
2011–12	61.74	61.58	61.52		0.16	0.22
2012–13	56.97	57.23	56.71		0.26	0.26
2013–14	56.67	56.63	56.28		0.03	0.38
				Total	**0.77**	**2.85**

Source: ICRC analysis.

Two out of the three parameter estimates for the preferred equation are significant at the 99 per cent level, as shown in Table 3.3. In evaluating these results, it should be noted that the equation is being used as an interpolation formula not a forecasting equation.

Table 3.3 Commission equation parameter significance

	Coefficient	Standard error	t-value	p-value
a	0.00001960964	0.00003	-0.5910	0.5962
b	49.5675	6.26600	7.9110	0.0042
c	63.41564	0.50600	125.3230	0.0000

Source: ICRC analysis.

Applying this equation to the total water sales forecasts presented in the previous chapter and estimated customer numbers for the 2014–15 to 2016–17 period is shown in Table 3.4.

[14] The Commission rejected a better fitting quadratic equation on the basis of the third assessment criterion as this equation produced results at lower levels of customer consumption that are counterintuitive and unsupported by the observed data.

Table 3.4 ACT tier 1 and 2 forecast water sales volumes and customer numbers

Year	Total sales (ML)	Tier 1 sales (ML)	Tier 2 sales (ML)	No. customers (#)	ML/customer/ year	Forecast tier 1 proportion
2014–15	39,964	24,252	15,712	167,244	0.239	60.68
2015–16	39,301	24,280	15,021	171,915	0.229	61.78
2016–17	39,109	24,356	14,753	176,717	0.221	62.28

Source: ICRC analysis.

Figure 3.2 Observed versus Commission modelled tier 1 proportion

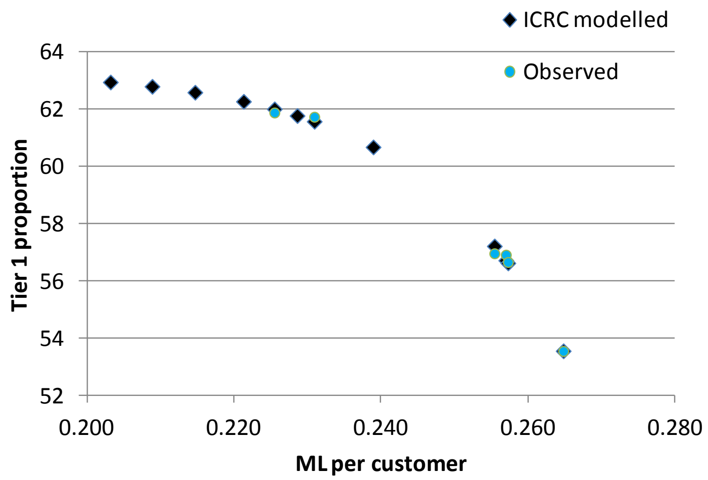

Source: ICRC analysis.

4 Summary

4.1 Preferred ARIMA model

The Commission's preferred model is a seasonal ARIMA (1,0,2)(1,0,1)[7] model, estimated over the period starting 1 July 2006, with the following climate lags and dummy variables:

- Temp (0, 1, 3, 12), Temp squared (0, 1, 3, 5, 12), Temp square root (1);
- Rain (0, 1), Rain squared (2, 3), Rain square root (0 - 8), Rain cubed (0 - 2), Rain square root (6);
- daily dummies (Sunday, Monday, Tuesday, Wednesday, Thursday, Friday);
- water restrictions dummy (Stage 3); and
- a Fourier series term.

4.2 Aggregate water sales forecasts

The dam releases and billed consumption forecasts from the preferred ARIMA specification are summarised in Table 4.1.

Table 4.1 April 2015 ARIMA model annual forecasts (ML)

Financial year	Dam releases	Billed consumption
2015	47,297	39,964
2016	46,260	39,301
2017	46,084	39,109

Source: ICRC analysis.

4.3 Tier 1 and tier 2 consumption split

The Commission's preferred equation for apportioning total water sales into tier 1 and tier components is of the form $y = \mathrm{c} + \mathrm{a}\mathrm{e}^{\mathrm{bx}}$ as follows:

$$y = 63.41564 - 0.00001960964\mathrm{e}^{49.5675\mathrm{x}}$$

where:

- y is the tier 1 proportion of total ACT water sales measured as a proportion of 100 units; and
- x is the average ACT customer consumption per year in ML.

Applying this equation to the total water sales forecasts and estimated customer numbers for the 2014–15 to 2016–17 period is shown in Table 4.2.

Table 4.2 ACT tier 1 and 2 forecast water sales volumes and customer numbers

Year	Total sales (ML)	Tier 1 sales (ML)	Tier 2 sales (ML)	No. customers (#)	ML/customer/ year	Forecast tier 1 proportion
2014–15	39,964	24,252	15,712	167,244	0.239	60.68
2015–16	39,301	24,280	15,021	171,915	0.229	61.78
2016–17	39,109	24,356	14,753	176,717	0.221	62.28

Source: ICRC analysis.

4.4 Final word

The modelling efforts undertaken by the Commission and reported in this and the January 2015 technical paper, provide a basis for regulator's determination of forecast water sales volumes for the next regulatory period.

Should anyone wish to examine the ARIMA model presented in this paper further, please contact the Commission for more detail on the data utilised and the R script developed to run the model.

Abbreviations and acronyms

ACF	Autocorrelation function
ACT	Australian Capital Territory
AIC	Akaike Information Criterion
ARIMA	Autoregressive integrated moving average
BOM	Bureau of Meteorology
CCF	Cross correlation function
Commission	Independent Competition and Regulatory Commission
GL	gigalitre (1,000 ML)
ICRC	Independent Competition and Regulatory Commission
kL	Kilolitre (1,000 litres)
L	Litre
ML	Megalitre (1,000 kL)
OLS	Ordinary least squares
PACF	Partial autocorrelation function
RMSE	Root mean square error

References

Box, GEP and Jenkins, GM 1970. Time Series Analysis: Forecasting and Control, Wiley.

Cardno, 2014. Independent review of ICRC price direction: Technical report - Prepared for the Industry Panel. Cardno (QLD) Pty Ltd. Available at: http://apps.treasury.act.gov.au/industrypanel/industry-panels-draft-report.

Cryer, JD and Chan, K 2008. "Time series analysis with applications in R". Second edition. Springer.

Engle, RF and Yoo, BS 1987. "Forecasting and testing in cointegrated systems." Journal of Econometrics. 35, 143-159.

Granger, CWJ 1986. "Developments in the study of cointegrated variable." Oxford Bulletin of Economics and Statistics. 48, 3, 213-228.

Hyndman, RJ 2010. "Forecasting with long seasonal periods". http://robjhyndman.com/hyndsight/longseasonality/.

Hyndman, RJ 2010a. "The ARIMAX model muddle". http://robjhyndman.com/hyndsight/arimax/.

Hyndman, RJ 2011. " Statistical tests for variable selection", http://robjhyndman.com/hyndsight/tests2/.

Hyndman, RJ 2012. "Flat forecasts". http://robjhyndman.com/hyndsight/flat-forecasts/.

Hyndman, RJ and Athanasopoulos, G 2012. "Forecasting: principles and practice: Online textbook on forecasting." https://www.otexts.org/fpp.

Icon Water 2014. "2015 Biennial recalibration - Statement of changes, Attachment E1: Water volumes forecasts." Canberra: Icon Water Limited.

ICRC, 2015. "Technical paper: Water demand forecasting." Report 1 of 2015. January 2015. Canberra: Independent Competition and Regulatory Commission.