

Adult prison population size in New South Wales: comparative forecasts

Neil Donnelly and Wai-Yin Wan

Aim: To compare the accuracy of Autoregressive Integrated Moving Average (ARIMA) model and Holt-Winters additive exponential smoothing method for forecasting the size of the total NSW adult prison population.

Method: NSW adult prison population data was obtained up until July 2016. A rolling origin approach was used with 20 estimation periods (of increasing length) and 20 validation periods (12-months length). ARIMA model and Holt-Winters additive method were applied to each rolling estimation period. Shorter term (e.g. 1 to 3 months) through longer term (e.g. 6 to 12 months) forecasts were made for the relevant 12-months validation period. These forecasts were compared with the actual monthly number of prisoners using mean absolute error (MAE), root mean square error (RMSE) and mean absolute percentage error (MAPE) to assess accuracy. The average of each accuracy measure was calculated for the one-step through 12-steps lead times across the 20 estimation and validation periods.

Results: For shorter forecast lead times (e.g. one-step through three-steps) the ARIMA model and Holt-Winters additive method gave similar accuracy measures for MAE, RMSE and MAPE. In each case, the accuracy of the forecasts decreased as the lead time increased. ARIMA was more accurate than Holt-Winters additive at longer lead times (e.g. six-steps through 12-steps) with smaller forecasting errors.

Conclusion: The ARIMA model provided more accurate forecasts compared with the Holt-Winters additive method for longer periods such as six-months to 12-months.

Keywords: prison population, forecasting accuracy, ARIMA, Holt-Winters additive; exponential smoothing, trend, seasonality, time series cross-validation.

INTRODUCTION

The aim of this study is to compare ARIMA model and exponential smoothing methods in the accuracy of their prospective forecasts of the monthly total adult prison population in New South Wales (NSW).

Prior research conducted in NSW has mostly used ARIMA models with both trend and seasonal components (Wan, Moffatt, Xie, Corben, & Weatherburn, 2013; Donnelly, Halstead, Corben & Weatherburn, 2015) to forecast prison population numbers. Wan et al. (2013), for example, used NSW Corrective Services data to compare different ARIMA models to forecast the NSW remand and sentenced prisoner populations. Monthly Offender Integrated Management Systems (OIMS) data was used from January 1998 to March 2013 (183 months). For the remand population, the estimation period was defined as January

1998 to December 2010 (156 months). Two different ARIMA models were fit for this period. In addition to terms to control for autocorrelation, these models also contained time dependent covariates such as prior monthly numbers of breaches of bail and imprisonment penalties. Each of the models produced forecasts for the validation period January 2011 to March 2013 (27 months or 15% of the total number of observations). The actual data for the validation period was compared with the forecasts from the two ARIMA models. The accuracy of the forecasts from each ARIMA model was then compared using the mean absolute percentage error (MAPE) and the root mean square error (RMSE). Lower values of MAPE and RMSE indicate that the forecasts are more accurate. Wan et al. (2013) selected the most accurate ARIMA model for the validation period and used this to forecast the remand prison population for a future period.

A similar approach was adopted by Wan et al. (2013) to forecast the sentenced prison population; the only difference being that the estimation period on which the ARIMA model forecasts were estimated was longer and the validation period was shorter. The validation period for the sentenced population was 15 months or eight per cent of the total observation period. Wan et al. (2013) used a longer estimation period to take account of a change in the trend which occurred in mid-2009. Again, the best ARIMA model was selected on the basis of the MAPE and RMSE measures during the validation period and used to forecast the sentenced population for a future period.

ARIMA models are not the only means by which forecasts can be generated. Exponential smoothing methods such as simple exponential smoothing, Holt's linear trend and Holt-Winters seasonal methods have been used in other forecasting contexts such as the market place or government expenditure agencies (Chatfield, 2004; Hyndman & Athanasopoulos, 2014). In these models a very large number of forecasts are typically needed and often the forecast period is shorter (e.g. the next quarter or year). An advantage of exponential smoothing methods is that they give greater weight to more recent observations when forecasting the value of a new observation. The weight given to an observation declines exponentially with the distance between it and the most recent observation (Makridakis, Wheelwright & Hyndman, 1998). The degree of exponential decay depends on the smoothing parameters. Values of the smoothing parameters closer to one have less of a smoothing effect and give greater weight to the recent changes in the time series, whereas values closer to zero have a greater smoothing effect and are less responsive to recent changes. A number of exponential smoothing methods have been developed so it is important that the correct type of approach be identified. If only a weight for level is required then simple exponential smoothing (SES) can be used. If both level and trend weights are required then Holt's linear (two-parameter) method can be used. Where weights are required for level, trend and seasonality Holt-Winters seasonal method can be used (Goodwin, 2010).

In the United Kingdom, forecasts from Holt-Winters methods and ARIMA models were compared across a large number of public service areas (UK Centre for the Measurement of Government Activity, 2008). This was done for public service areas that had either annual or quarterly data to forecast. A one-step-ahead forecast (single quarter or year) was used where there were between six and 51 series to calculate the MAPE from. There were differences across the public service areas in terms of whether the ARIMA model or Holt-Winters method produced the smaller (and more accurate) MAPE. A notable difference between the UK study and the NSW corrective services prison studies (Wan et al. 2013) is that, in the former, the MAPE is calculated for a single time point using multiple series. For the NSW prison population forecasts the MAPE (and RMSE) is

calculated across multiple lead times (months) using a single series (NSW ault prisoner numbers).

Makridakis and Hibon (2000) reported findings from the M3-Competition¹ which included 3,003 series analysed using different forecasting approaches. Different approaches were then compared in terms of how close the forecast values (from the estimation period) were to the actual values. These series were subdivided into different content areas, such as micro-economic, industry, finance and demographic. The series also varied in terms of the units in which they were measured (yearly, quarterly, monthly) and also what the length of the forecast period was (one-step-ahead forecasts or longer lead times such as six-steps-ahead, or 12-steps-ahead). Makridakis and Hibon (2000) drew a large number of conclusions about forecasting methods based on their results but there are two which are of particular interest to the current study. Firstly, exponential smoothing methods often provided more accurate forecasts compared with specific ARIMA modelling². Secondly, the accuracy of a particular forecast methodology can vary depending on the length of the forecast period.

Many exponential smoothing methods such as simple exponential smoothing, Holt's linear and Holt-Winters methods have an equivalent ARIMA model (Fomby, 2008; Hyndman, 2014). Identifying an alternative and adequate ARIMA model to make specific forecasts, however, can be time consuming. Exponential smoothing methods, by contrast, can be quickly applied to a large number of series.

CURRENT NSW PRISON FORECAST STUDY

The first aim of this research is to determine whether a specific ARIMA model or an exponential smoothing method (such as Holt-Winters additive) is more accurate in providing 12-monthly forecasts for the NSW adult prison population. An important feature of this comparison is that the ARIMA model does not contain any exogenous predictors. A time series cross-validation approach was used to assess their accuracy with a large number of estimation periods and 12-month validation periods (Hyndman & Athanasopoulos, 2014; Hyndman, 2011). For each estimation period, forecasts from the ARIMA model and Holt-Winters methods were made for the associated 12-month validation period. Actual monthly data from the validation period were compared with the forecast values from the ARIMA model and Holt-Winters methods using mean absolute error (MAE), RMSE and MAPE measures to assess their accuracy. For each step of the 12-monthly forecast, the average of each accuracy measure was calculated across the 20 validation periods. The advantage of using this rolling origin approach, rather than a single split between the estimation and validation period is that the former is based on more than one validation period. This makes it more robust to outliers or unexpected values.

The second research aim is to compare the 12-monthly forecasts made from ARIMA model and Holt-Winters method using all the NSW prison population collected to date. While this comparison provides no further information on the accuracy of the ARIMA and Holt-Winters forecasts, it is of interest to compare their predictions for future growth in the NSW prison population.

METHOD

DATA

Data on monthly prisoner numbers were provided by Corrective Services NSW. These data were extracted from OIMS and covered the period July 2001 through July 2016 (181 months). The total adult prison population per month was obtained by combining the sentenced prisoner and remand population numbers from OIMS. The remand population also includes a small number of individuals who are held in police and court cells.

Estimation and validation periods

A rolling origin approach was implemented in which 20 estimation periods were used to produce one-step-ahead to 12-steps-ahead forecasts for 20 forecast periods. The first estimation period was July 2001 to December 2013 with a 12-month validation period from January 2014 to December 2014. The second estimation period had one extra estimation data point (July 2001 to January 2014) and the corresponding 12-month validation period was February 2014 to January 2015. The final estimation period was July 2001 to July 2015 and the associated 12-month validation period was August 2015 to July 2016.

STATISTICAL ANALYSIS

Analyses were conducted to identify the best ARIMA model and Holt-Winters method for fitting the monthly total prisoner numbers in NSW using the full period from July 2001 to July 2016. These analyses were conducted using SAS (Version 9.4). ARIMA modelling was undertaken using PROC ARIMA and Holt-Winters methods were undertaken using PROC ESM.

ARIMA model

As the size of the adult prison population contained both increasing and decreasing trend and a seasonal component, augmented Dickey-Fuller (ADF) tests were conducted to determine if the series needed to be seasonally differenced and/or first-order differenced to make it stationary (Box, Jenkins, & Reinsel, 1994; Enders, 2015). It was found that seasonal differencing was not required but first-order (serial) differencing was. Autocorrelation function (ACF) and partial autocorrelation function (PACF) plots were examined to determine the nature of the autocorrelation in the first-differenced stationary data. This included assessing the need for autoregressive (AR) and moving

average (MA) terms at particular lags (such as 1, 12 or some other value). The AR and MA terms at relevant lags were tested. The final ARIMA model contained those AR and MA terms that were statistically significant and had the best goodness of fit (Box et al., 1994; Chatfield, 2004; Hyndman & Athanasopoulos, 2014).

Residuals from the final ARIMA model were examined to ensure all significant autocorrelation had been adjusted for. This included using the Ljung-Box Q test at lags of six, 12, 18 and 24 (Ljung & Box, 1978). A non-significant p -value indicates that there is no autocorrelation within that lag period. The value of the Akaike Information Criteria (AIC) was also used to compare the fit of different ARIMA structures. A lower AIC indicates a better fit to the data (Akaike, 1974). This final ARIMA model was used later in the cross-validation.

Exponential smoothing method

Given the presence of both trend and seasonality, the Holt-Winters additive method and Holt-Winters multiplicative method were used (Makridakis et al., 1998). Both methods include three smoothing equations for level (overall smoothing), trend and season. Each equation contains a smoothing weight to be estimated. Using SAS 9.4, the smoothing weights were optimized so as to minimize the sum of squared, one-step-ahead within-sample forecast errors. Additive and multiplicative Holt-Winters methods were compared on the goodness of fit to the complete prison population data.

Accuracy of the forecasts from the estimation period to the validation period

In the cross-validation, 20 sets of 12-steps-ahead forecasts were obtained from the ARIMA model and the Holt-Winters additive method respectively from the 20 validation periods. Three forecast accuracy measures were used to compare the one-step-ahead to 12-steps-ahead forecasts (\hat{y}_t) with the actual observation (y_t) averaged across the 20 validation sets ($n=20$).

These measures include:

- (1) the mean absolute error (MAE) defined as:

$$MAE = \frac{\sum_{t=1}^n |y_t - \hat{y}_t|}{n} \quad (1)$$

- (2) the root mean square error (RMSE) defined as:

$$RMSE = \sqrt{\frac{\sum_{t=1}^n (y_t - \hat{y}_t)^2}{n}} \quad (2)$$

- (3) the mean absolute percentage error (MAPE) defined as:

$$MAPE = \frac{1}{n} \sum_{t=1}^n \left| \frac{y_t - \hat{y}_t}{y_t} \right| * 100 \quad (3)$$

The size of the three measures was compared for the one-step-ahead to 12-steps-ahead forecasts between the ARIMA model and the chosen Holt-Winters method. Smaller values of MAE, RMSE and MAPE indicate that the forecast values were closer to the actual total prison population (Hyndman & Koehler, 2006).

Future forecasts from the complete dataset

In the final section, the ARIMA model and the chosen Holt-Winters method were applied to the total adult prison population from July 2001 until July 2016 to generate forecasts for the 12-month period up until July 2017. The 12-month forecasts from each model were compared graphically and reported in terms of monthly values, monthly averages and the final forecast value in July 2017.

RESULTS

TRENDS IN NSW ADULT PRISON POPULATION

Figure 1 shows the size and trend of the adult sentenced, remand and total prison populations from July 2001 until July 2016. Between mid-2001 and mid-2009 the total prison population increased by 34 per cent. This was followed by a decrease of eight per cent up until September 2012 and then an increase of 13 per cent by March 2014 (Weatherburn, Wan, & Corben, 2014). There was a very small decline in the total prison population between May and December 2014 (by 2.5 per cent).

This decline was mostly due to a transient drop in the remand population of around 10 per cent between May and December 2014 (Weatherburn & Fitzgerald, 2015). Since January 2015 the total prison population size has increased by 16 per cent.

IDENTIFYING THE BEST ARIMA MODEL & EXPONENTIAL SMOOTHING METHOD FOR LATEST PRISON POPULATION DATA

ARIMA model

The monthly prison population data did not require seasonal differencing as the augmented Dickey-Fuller test rejected the null hypothesis that there was a stochastic seasonal unit root. In terms of the first level it was confirmed that the series was difference-stationary rather than trend-stationary by conducting augmented Dickey-Fuller tests (Enders, 2015).³ The data were differenced. Examination of the autocorrelation function (ACF) and partial autocorrelation function (PACF) plots showed that the difference-stationary series contained autocorrelation. Autoregressive parameters at lags of one (AR1), eight (AR8)

Figure 1. Total, sentenced and remand prison population in NSW




Table 1. ARIMA time series model of NSW total prison population Estimation period (July 2001 - July 2016)

	Estimate	Standard Error	p value
Constant	29.079	11.957	= .015 *
Autoregressive (AR1)	0.414	0.068	< .001 *
Lag 1			
Autoregressive (AR8)	-0.235	0.074	= .001 *
Lag 8			
Autoregressive Seasonal (AR12)	0.466	0.068	< .001 *
Lag 12			

* Significant at the 0.05 level

and 12 (AR12) were added to the ARIMA model and provided the best fit to the data as indicated by the AIC. AR1 shows that adjoining months are autocorrelated while AR12 shows seasonal autocorrelation (e.g. March prison number is correlated with previous March prison number). The AR8 parameter provided a better fit to the data compared with the MA8 parameter (moving average of lag 8). The parameters from the final ARIMA model are shown in Table 1.

For diagnostic checking, examination of the residuals using ACF and PACF plots found that this ARIMA model had removed all statistically significant autocorrelation (including those at lags one, eight and 12). This was confirmed by the Ljung-Box Q tests conducted for the first six lags ($p = .264$), 12 lags ($p = .435$), 18 lags ($p = .667$) and 24 lags ($p = .861$). The AIC for this final model was 2024.7. Examination of the residuals from the final ARIMA model showed them to be normal, which was confirmed by the straight line shown in the normal QQ-plot. This final ARIMA model is later used in the cross-validation to produce 20 sets of 12-monthly forecasts for 20 validation periods.

Exponential smoothing method

Holt-Winters additive method was compared with Holt-Winters multiplicative method. The additive form assumes that the seasonal component remains constant over time while the multiplicative form assumes that the seasonality either increases or decreases in magnitude. The Holt-Winters additive method provided a better fit to the data compared with the multiplicative method. This was indicated by smaller values of the AIC (1524.7 vs. 1593.6), MAE (50.6 vs. 61.0), RMSE (66.4 vs. 80.3) and

MAPE (0.53% vs. 0.63%). The Holt-Winters additive method was used in the time series cross-validation analyses reported below and compared with the ARIMA model specified above.

COMPARISON OF FORECASTING ACCURACY BETWEEN ARIMA MODEL AND HOLT-WINTERS ADDITIVE METHOD

The forecasting accuracy of each method was compared using time series cross-validation. Figure 2 shows the three measures of forecasting accuracy averaged over the 20 validation sets for the ARIMA model and Holt-Winters additive method. Accuracy is shown for one-step through 12-steps lead times. Values of the MAE and RMSE are shown on the left hand axis. The MAPE is shown on the right hand axis. All curves rise as the forecast extends from one to 12-months, indicating that the accuracy of the forecasts declines as the forecast horizon increases. Both ARIMA model and Holt-Winters additive method give similar MAE, RMSE and MAPE values at earlier lead times; however the differences increase from six-steps-ahead onwards. The ARIMA model outperforms Holt-Winters additive method at later lead times with smaller forecasting errors.

The better performance of the ARIMA model is probably due to the fact that it captures the change point at the second quarter of 2014 better than the Holt-Winters additive method and hence gives a more accurate forecast especially at later lead times. In the next section we re-run the ARIMA model and Holt-Winters additive method using data for the full observation period (181 observations) in order to obtain a 12-monthly forecast for the period from August 2016 to July 2017.

Figure 2. Comparison of forecasting accuracy measures between ARIMA model and Holt-Winters additive method




Figure 3. Total prison population in estimation period: ARIMA 12 - monthly forecast





Figure 4. Total prison population in estimation period: Holt-Winters additive 12-monthly forecast



FORECASTS FROM LATEST PRISON POPULATION DATA

The ARIMA model as shown in Table 1 was used to forecast the total prisoner numbers for the period August 2016 through July 2017. The black line in Figure 3 shows the actual prison population values. The green line shows the predicted values from the ARIMA model. The red line shows the 12-monthly forecasts of total prisoner numbers based on the ARIMA model. The graph shows a monotonic increasing trend throughout the forecasting period.

Figure 4 shows the corresponding Holt-Winters additive 12-monthly forecasts. Again, the red line shows the 12-month forecasts of total prisoner numbers based on this method. The Holt-Winters forecast shows an increasing trend over the 12-month period.

Comparing estimates from the two forecasting methods we see some interesting differences. The average monthly prison population over the 12-monthly forecast period, based on the ARIMA model, is 13,098, while the average monthly prison population forecast based on the Holt-Winters additive method

is slightly larger at 13,134. For the ARIMA model, the short-term forecast is that, if all else remains constant, NSW will have 13,354 prisoners by July 2017, while the 12-month estimate from the Holt-Winters additive forecast is slightly higher at 13,542 prisoners.

DISCUSSION

Reliable estimates of future trends in the size of the prison population are essential for correctional services policy and planning. The short-term forecasts of prisoner numbers currently provided by BOCSAR are based on ARIMA modelling. The current study was undertaken to see whether Holt-Winters methods provide more accurate forecasts of prisoner numbers over 12-monthly periods than those derived from ARIMA models.

The two forecasting methods were compared by conducting time series cross-validation analyses in which each of the 20 estimation periods were used to make 12-month forecasts for the corresponding validation period. Using a rolling origin approach the estimation periods got progressively larger while the validation periods remained 12-months in length (i.e. one-step-ahead through 12-steps-ahead forecasts were calculated).

The results of the cross-validation analyses revealed that shorter term forecasts (e.g. 1 month, 2 months, 3 months ahead) were more accurate than longer term forecasts (e.g. 6 months, 7 months or more). Moreover, the accuracy of the ARIMA model and Holt-Winters additive method were similar when estimating over the shorter-term. However for longer term forecasts, such as six-months through to 12-months, the ARIMA model was found to be the more accurate method.

The ARIMA model and Holt Winters additive method estimate an increasing trend but differ in how much they anticipate the population will grow. If all else remains constant, the ARIMA model predicts NSW will have 13,354 prisoners by July 2017, while the Holt-Winters additive method predicts the prison population will be slightly larger at 13,542. The caveat 'if all else remains constant' is very important here. Forecasts based on unadjusted ARIMA models or Holt-Winters additive methods are a sophisticated extrapolation from past trends. These methods ignore changes in the determinants of prison population growth, such as crime rates, arrest rates, bail policy and sentencing policy. It is possible that models which incorporate these determinants (e.g. in multivariate ARIMA models) would give more reliable forecasts than the univariate models. Any benefits derived in terms of accuracy would, however, need to be weighed against any delays incurred as a result of having to source these additional data and re-estimate the models accordingly. Future research could consider whether the inclusion of other time-dependant covariates results in more robust longer-term prison population forecasts. In the meantime, the current work has established reliable methods for forecasting prison population over periods of up to 12-months.

ACKNOWLEDGEMENTS

We would like to acknowledge Simon Corben and Jennifer Galouzis from Corrective Services NSW for providing the prison population data. We thank Dr Suzanne Poynton and Dr Don Weatherburn for their review of an earlier version of this bulletin. The very helpful comments and statistical advice provided by an external reviewer is greatly appreciated.

NOTES

1. The M-Competitions (or Makridakis Competitions) are empirical studies that have compared the performance of a large number of major time series methods in terms of their forecast accuracy. Forecasts from experts are compared with each other and with other some simple methods as benchmarks. The M3-competition is the latest of these studies and includes more methods, experts and series than the earlier two attempts.
2. Makridakis and Hibon (2000, p. 459) did point out, however that more specific ARIMA models can be better than exponential smoothing methods by adequately controlling for auto-correlated errors when using "available historical data".
3. In augmented Dickey-Fuller test the null hypothesis is that the series has a stochastic unit root. The alternative hypothesis is different depending on which version of the test is used. There are three versions of the test: (a) test for unit root ($\Delta y_t = y_t - y_{t-1} = \gamma y_{t-1} + e_t$); (b) test for unit root with drift ($\Delta y_t = b_0 + \gamma y_{t-1} + e_t$); and (c) test for unit root with drift and deterministic time trend ($\Delta y_t = b_0 + b_1 t + \gamma y_{t-1} + e_t$) (Enders, 2015). The null hypothesis assumes that γ equals zero ($\gamma = 0$). Failure to reject this null hypothesis indicates that the series needs to be differenced to attempt to remove the stochastic unit root. Augmented Dickey-Fuller tests for unit root with drift and deterministic trend were applied to the total adult prison population data from the estimation period. The deterministic intercept (b_0) and deterministic trend terms (b_1) were not significant and removed. The augmented Dickey-Fuller test for unit root found that the γ term in the random walk model was not statistically significant ($Z(t) = 1.81, p = .998$). This means the total prison population size series contains a stochastic unit root and first-differencing is required to make it stationary. It was then transformed to be a first difference series (each observation is subtracted from the previous observation) and augmented Dickey-Fuller tests were conducted again. It was found that the first differenced series did not contain a stochastic unit root as the null hypothesis that γ equals zero was rejected ($Z(t) = -6.99, p < .001$). This first difference total prison population size series was used in the ARIMA analyses. These findings were replicated using Phillips-Perron tests.

REFERENCES

- Akaike, H. (1974). A new look at the statistical model identification. *Automatic Control, IEEE Transactions on Automatic Control*, 19, 716-723.
- Box, G.E.P., Jenkins, G.M., & Reinsel, G.C. (1994). *Times series analysis: Forecasting and control. (3rd Ed.)*. Englewood Cliffs, NJ: Prentice Hall.
- Chatfield, C. (2004). *The analysis of time series: An introduction. (6th Ed.)*. London: Chapman & Hall/CRC.
- Donnelly, N., Halstead, I., Corben, S., & Weatherburn, D. (2015). The 2015 NSW prison population forecast. (Bureau Brief No. 105). Retrieved 19 August 2016 from NSW Bureau of Crime Statistics and Research website: <http://www.bocsar.nsw.gov.au/Documents/BB/BB105.pdf>
- Enders, W. (2015). *Applied econometric time series. (4th Ed.)*. Hoboken, NJ: John Wiley & Sons.
- Fomby, T.B. (2008). *Exponential smoothing models*. Retrieved 19 August 2016 from Southern Methodist University (SMU) website: http://faculty.smu.edu/efomby/eco5375/data/SMOOTHING%20MODELS_V6.pdf
- Goodwin, P. (2010). The Holt-Winters approach to exponential smoothing: 50 years old and going strong. *Foresight: The International Journal of Applied Forecasting*, 19, 30-33.
- Hyndman, R.J. (2011). *Time series cross validation: an R example*. Retrieved 19 August 2016 from Rob J. Hyndman website: <http://robjhyndman.com/hyndsight/tscvexample/>.
- Hyndman, R.J. (2014). *Forecasting using R: 10. Seasonal ARIMA models*. Retrieved 19 August 2016 from Rob J. Hyndman website: <http://robjhyndman.com/talks/RevolutionR/10-Seasonal-ARIMA.pdf>.
- Hyndman, R.J., & Athanasopoulos, G. (2014). *Forecasting: principles and practice*. Melbourne: OTexts.
- Hyndman, R.J. & Koehler, A.B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting*, 22, 679-688.
- Ljung, G.M., & Box, G.E.P. (1978). On a measure of lack of fit in time series models. *Biometrika*, 65, 297-303.
- Makridakis, S., & Hibon, M. (2000). The M3-Competition: results, conclusions and implications. *International Journal of Forecasting*, 16, 451-476.
- Makridakis, S., Wheelwright, S.C., & Hyndman, R.J. (1998). *Forecasting: Methods and Applications (3rd Ed.)*. New York: John Wiley & Sons.
- UK Centre for the Measurement of Government Activity. (2008). *From Holt-Winters to ARIMA modelling: Measuring the impact of forecasting errors for components of quarterly estimates of public service output*. London: Office for National Statistics.
- Wan, W-Y., Moffatt, S., Xie, Z., Corben, S., & Weatherburn, D. (2013). *Forecasting prison populations using sentencing and arrest data (Crime and Justice Bulletin No. 174)*. Retrieved 19 August 2016 from NSW Bureau of Crime Statistics and Research Website: <http://www.bocsar.nsw.gov.au/Documents/CJB/cjb174.pdf>
- Weatherburn, D., & Fitzgerald, J. (2015). *The impact of the NSW Bail Act (2013) on trends in bail and remand in New South Wales (Bureau Brief No. 106)*. Retrieved 19 August 2016 from NSW Bureau of Crime Statistics and Research Website: http://www.bocsar.nsw.gov.au/Documents/BB/Report_2015_Bail_and_Remand_bb106.pdf
- Weatherburn, D., Wan, W-Y., & Corben, S. (2014). *Why is the NSW prison population growing? (Bureau Brief No. 95)*. Retrieved 19 August 2016 from NSW Bureau of Crime Statistics and Research Website: <http://www.bocsar.nsw.gov.au/Documents/BB/bb95.pdf>