Seasonal Stochastic Rainfall Modelling Using Climate Indices: A Bayesian Hierarchical Model

Henley, B.J. [1], M.A. Thyer [1] and G. Kuczera[1]

[1] School of Engineering, The University of Newcastle, Australia, Email: ben.henley@newcastle.edu.au

Keywords: *Bayesian hierarchical, Interdecadal Pacific Oscillation, IPO, climate variability, drought risk, stochastic model*

EXTENDED ABSTRACT

It is well documented that large-scale climate mechanisms such as the Interdecadal Pacific Oscillation (IPO) and El Niño Southern Oscillation (ENSO) play a dominant role in the hydrological variability of Australia. The recent drought in many parts of Australia has highlighted the importance of reliably estimating drought risk. Stochastic models play a key role in the estimation of drought risk to water supply systems.

Much work has been done to improve the ability of stochastic models to capture the observed statistics of hydrological records, albeit sometimes with the somewhat contestable underlying assumption of a static long-term climate. A Bayesian Hierarchical framework is presented that can incorporate climate indices into a stochastic rainfall model. **Figure 1** compares the IPO index (11-year moving average) with the annual rainfall anomaly (11-year moving average) at Stroud in the Hunter Valley, NSW.

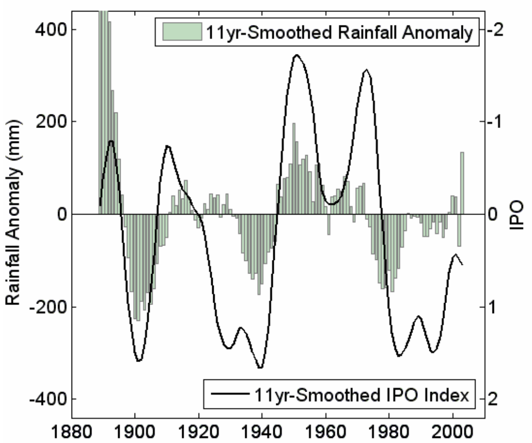

Figure 1 Stroud annual rainfall anomaly and IPO index: the influence of a multi-decadal climate mechanism on rainfall

The seasonal impact of the IPO on rainfall is investigated using distribution-free statistical tests. This is used to inform the structure of a seasonal stochastic rainfall model, termed the H4 model. The maximum IPO-impact seasons for several sites near to the case-study were summer-centred with lengths of two to four months. However, many other season definitions were statistically significant.

The distribution of run-lengths of the cumulative sum of annual rainfall anomaly was presented as an alternative measure of long-term dependence in stochastic models of hydrological series. For the case study presented in this paper this approach showed that the H4 model has a higher probability than the widely used lag-one autoregressive (AR(1)) model of run-lengths longer than 30 years. However a key difficulty with the comparison of long-dependence statistics is the short observed record.

A drought risk analysis was included to determine indicative effects of different models on long-term reservoir behaviour. The approximate size and characteristics of the proposed Tillegra Dam were used. Drought risks were found to be slightly higher for the H4 model than the AR(1) model for some annual demand scenarios, though the differences are relatively small.

This framework has the capability of utilising data other than hydrological sources to characterise climate variability on multiple time-scales. Further research will utilise full parameter uncertainty techniques. Sub-decadal variability will be incorporated into the framework by stochastic processes for the ENSO and Indian Ocean Dipole (IOD) phenomena. It is suggested that drought risk is possibly more sensitive to these mechanisms than the IPO. The connection between the IOD and winter rainfall in Eastern Australia means that the inclusion of the IOD in the framework might provide a better characterisation of the winter climate processes. The addition of palaeo-data in formal observation processes would expand the observed record, providing more certainty in the characterisation of long-term dependence.

1. INTRODUCTION

The recent drought in many parts of Australia has brought water security and climate science into the forefront of public thinking. Water authorities are presented with the significant challenge of managing water supply systems, due to the highly variable and persistent hydrology of Australia. Stochastic models play a key role in the estimation of drought risk to water supply systems. They provide stochastic simulations of hydrological inputs which are used in Monte-Carlo simulations of reservoir systems in order to estimate the water supply performance and optimise operating procedures under scenarios of interest.

It has been extensively documented that Australian hydrological data are modulated by large-scale climatic mechanisms such as the El Niño-Southern Oscillation (ENSO) phenomenon and the Interdecadal Pacific Oscillation (IPO) (Kiem & Franks, 2001, Verdon *et al.*, 2004). When sea surface temperatures (SST) in the tropical pacific warm, (IPO positive phase), the link between ENSO and Australian climate weakens, and the La Nina's frequency decreases. When SST's in the tropical pacific cool (IPO negative phase), the frequency of 'replenishing' La Nina's is increased (Verdon *et al.*, 2004). This is shown in **Figure 1** which compares the Interdecadal Pacific Oscillation (IPO) index (11-year moving average) with the annual rainfall anomaly (11-year moving average) at Stroud in the Hunter Valley, NSW.

Much work has been done to improve the ability of stochastic models to capture the observed statistics of hydrological records, including parametric (Salas, 1993) and non-parametric approaches (Sharma *et al.*, 1997). Often, these models have the somewhat contestable underlying assumption of a static long-term climate. The lag-one autoregressive (AR1) model is used widely by the water resources industry for stochastic rainfall simulations (Thyer & Kuczera, 2000). Thyer and Kuczera (2000) applied a hidden Markov model (HMM) to mimic the wet-dry climate state behaviour. Whiting *et al.*, (2003) criticised the HMM for possibly resulting in simply a mixture of two normal models, and suggested an alternative that informed a stochastic model with the Pacific Decadal Oscillation (PDO) index and the Southern Oscillation Index (SOI) using a multiple-linear regression approach.

Henley *et al.*, (2006) presented an approach that incorporated climate indices into a stochastic rainfall model via a formal Bayesian hierarchical framework. This approach aimed to explicitly incorporate climate mechanisms into the model design and actively incorporate the natural quasi-periodic variability of hydro-climatological records. The aim is to improve on the traditional approaches such as the AR(1) model which (it could be argued) have a rather passive means of inducing persistence into rainfall simulations.

Previously, Henley *et al.*, (2006) used a simplified approach to simulate the IPO index and inform rainfall simulations at the annual time step. It was suggested that the annual simulations were too coarse to properly characterise the seasonal influences of the IPO index on the observed rainfall record.

This paper investigated the seasonal impact of the IPO on the rainfall to develop a seasonal model and determine if this seasonality improves the stochastic rainfall model (section 2). The rainfall data from Stroud was chosen as the case study because of its proximity to Hunter Water Corporation's (HWC) water supply system (the industry sponsors of the project). The Stroud data is part of the Bureau of Meteorology's high quality data set (Lavery *et al.*, 1997). The model calibration is described in section 3 and the results from the model calibration and testing are given in section 4. A simple reservoir simulation was included to compare the effects of different models on simulated long-term drought risks. Reservoir characteristics similar to HWC's proposed Tillegra Dam were used. It should be noted that the results were not indicative of real-world drought risk as operating rules and the integration with the overall water supply system was not included.

2. SEASONAL ANALYSIS AND MODEL DEVELOPMENT

2.1. Identification of the Impact Season

In order to develop a seasonal hierarchical model, hereafter referred to as H4, an analysis of the monthly and seasonal influence of the IPO was undertaken for several sites located within the vicinity of the HWC water supply catchments. Henley *et al.*, (2006) only found statistical justification for two IPO phases, a positive and negative phase, hence this approach was adopted.

The monthly statistics in IPO-positive and negative epochs for Stroud are shown in **Figure 2**. These results are typical for other sites within the region. The influence of the IPO can be seen most clearly in January, February, March and June, especially in the differences in upper quartile values. This result is concurrent with the higher frequency of high-rainfall, 'replenishing' La Niñas in IPO-negative epochs.

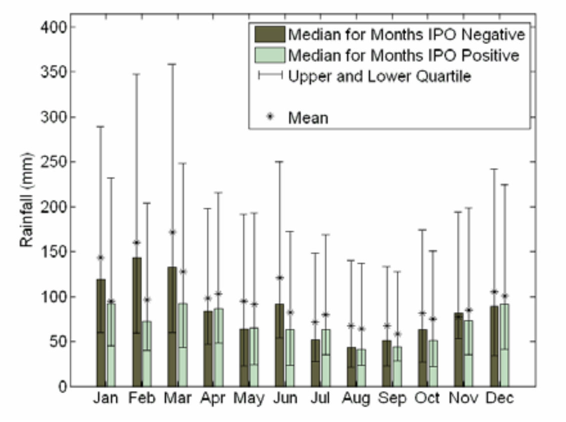

Figure 2 Monthly Rainfall Statistics for Stroud Stratified by IPO Epoch

The Wilcoxon rank-sum test is a non-parametric test to determine whether two samples are from the same distribution. This test was applied to monthly and seasonal rainfall totals at the 5% significance level. The results are summarised in **Figure 3**. All possible contiguous impact-season definitions at a monthly resolution were tested. This included durations from 1 to 12 months (inclusive of the starting month) for starting months January to December.

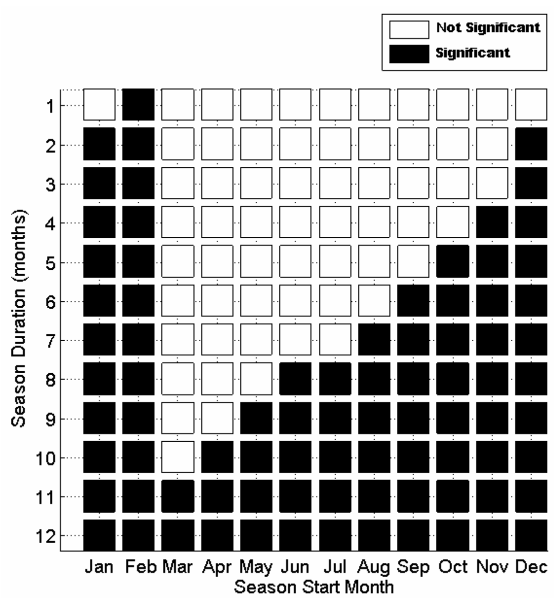

Figure 3 Significance test results for all contiguous season definitions at Stroud using the Wilcoxon rank-sum test on IPO-stratified seasonal rainfall at the 5% significance level

This demonstrates that many season definitions show a significant IPO-impact at Stroud. The resulting 'p-values' from the rank-sum tests are the probabilities that the IPO positive and negative epoch rainfall sample sets are from the same distribution. The IPO-impact season for the H4 model was chosen by taking the season with the lowest probability, indicating the maximum impact season. This step can be seen as part of the calibration process. For Stroud the resulting impact season was the three month season from January to March. Several sites near to the Stroud case study (**Figure 5** and **Figure 4**) also reveal a summer-centred maximum impact period with season lengths of two to four months. For most sites many other season definitions were also statistically significant. The comparison between season definitions with small differences in their p-values might be somewhat arbitrary, however the minimum p-value technique used here was a convenient and statistically sound method for obtaining the maximum impact period.

Figure 4 Location of several indicative rainfall sites near to the Stroud case study (image: Google Earth)

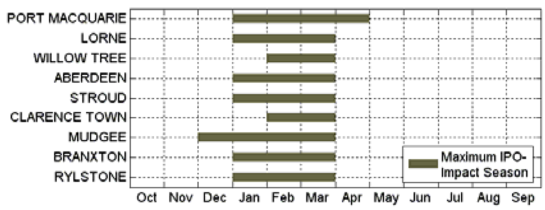

Figure 5 Maximum IPO-impact seasons based on Wilcoxon rank-sum tests on IPO-stratified seasonal rainfall

2.2. Statistical Analysis of Seasonal Data

Posterior distributions of the mean seasonal rainfall at Stroud stratified by IPO-epoch in **Figure 6** show a very clear distinction between the distributions during the maximum impact season of January to March. The threshold of zero was used for the IPO phase identification. The probability that the IPO-negative (wet) mean is less than the IPO-positive (dry) mean is less than 0.1% during the impact season. The Kolmogorov-Smirnov test indicated there was no evidence to reject the assumption of Gaussian seasonal rainfall totals at the 5% significance level.

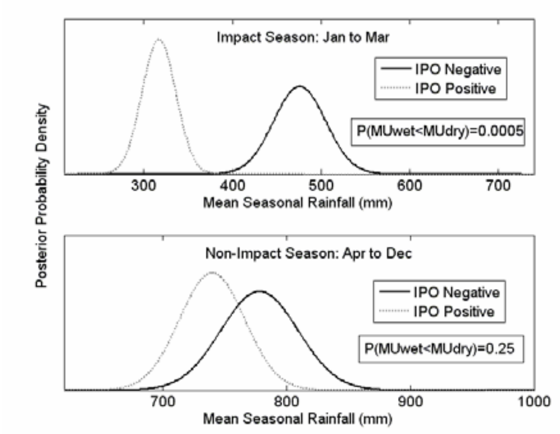

Figure 6 Posterior distributions of the mean seasonal rainfall

Cross correlations between impact and non-impact seasons and autocorrelations within the seasons were calculated for the Stroud seasonal rainfall. The lag-1 autocorrelation in the impact season was statistically significant. The cross correlations between impact and non-impact seasons included the case where impact season precedes non-impact season and vice versa. Neither option gave statistically significant cross correlation. Lag-1 autocorrelations of both the non-impact season and the annual data were not statistically significant.

2.3. Model Structure

As mentioned, the hierarchical framework used in this paper was developed previously in Henley *et al.*, (2006). Hence, a brief description of the overall framework is given here, and an outline of the changes to the previous model implemented in the H4 model. The two-level general framework is presented in the directed acyclic graph (DAG) in **Figure 7**. It is a conceptual representation of two processes: the random unobservable natural process on the left and the observation process through which the framework is informed with real world data on the right of the graph.

This framework was used to model the impact season rainfall. At level 0, the H4 model uses the IPO data y^0 to explicitly inform the long-term process, s^0. Assuming $y^0 = s^0$ means there is no observation model and no β^0 parameters; likewise for y^1 and β^1. Uppercase S^0_{t-1} represents all simulations prior to time step t. Henley *et al.*, (2006) chose a lag-two auto-regressive (AR(2)) stochastic process to simulate the long-term process s^0_t shown in equation 1. This was based on the simplest ARMA model that displayed quasi-periodic run-length behaviour. The process has parameters $\alpha^0 = \left[\mu_{IPO}, \sigma_{IPO}, \phi_{1\,IPO}, \phi_{2\,IPO}\right]$.

$$s^0_t = \mu_{IPO} + \phi_{1IPO}(s^0_{t-1} - \mu_{IPO}) + \phi_{2IPO}(s^0_{t-2} - \mu_{IPO}) + N(0, \sigma^2_{IPO}) \quad \textbf{(1)}$$

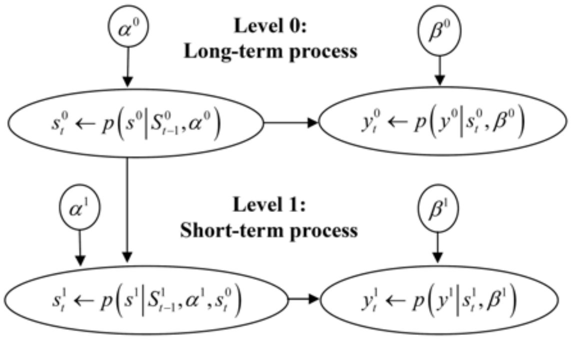

Figure 7 A two-level hierarchical framework that relates stochastic processes and observation processes at different time scales

At level 1, the short-term process s^1_t includes the rainfall in the impact period $s^1_t(1)$ and the non-impact period $s^1_t(2)$ as shown in equation 2. The impact period rainfall $s^1_t(1)$ was simulated using a dynamic AR(1) process as shown in equation 3. A deterministic relationship between the IPO and the impact season rainfall was assumed, as shown in equations 4 and 5. The non-impact season was simulated with the AR(1) process shown in equation 6.

$$s^1_t = [s^1_t(1) \quad s^1_t(2)] \quad \textbf{(2)}$$

$$s^1_t(1) = \mu_I(s^0_t) + \phi_I(s^0_t, s^0_{t-1}) \cdot (s^1_{t-1}(2) - \mu_I(s^0_t)) + N(0, \sigma^2_I(s^0_t)) \quad \textbf{(3)}$$

$$\phi_I(s^0_t, s^0_{t-1}) \begin{cases} \phi_{1_I} \ \ if \ s^0_t > \mu_{IPO} \ and \ s^0_{t-1} > \mu_{IPO} \\ \phi_{1_I} \ \ if \ s^0_t < \mu_{IPO} \ and \ s^0_{t-1} < \mu_{IPO} \\ 0 \quad otherwise \end{cases} \quad \textbf{(4)}$$

$$\mu_I(s^0_t), \sigma_I(s^0_t) \begin{cases} \mu_{w_I}, \sigma_{w_I} \ \ if \ s^0_t \le \mu_{IPO} \\ \mu_{d_I}, \sigma_{d_I} \ \ if \ s^0_t > \mu_{IPO} \end{cases} \quad \textbf{(5)}$$

$$s^1_t(2) = \mu_{NI} + \phi_{NI} \cdot (s^1_{t-1}(2) - \mu_{NI}) + N(0, \sigma^2_{NI}) \quad \textbf{(6)}$$

The parameters $\alpha^1(2)$ and $\alpha^1(3)$ are as follows:

$$\alpha^1(1) = \left[\mu_{w_I}, \sigma_{w_I}, \mu_{d_I}, \sigma_{d_I}, \phi_{1_I}\right] \quad \textbf{(7)}$$

$$\alpha^1(2) = \left[\mu_{NI}, \sigma_{NI}, \phi_{1_{NI}}\right] \quad \textbf{(8)}$$

3. MODEL CALIBRATION

A maximum likelihood approach was used to fit the H4 model to the IPO and Stroud data. It can be shown that the joint likelihood of the parameters can be simplified to equation 9. The sum of squares method was used to calculate the unconditional likelihood of the auto-regressive moving average (ARMA) process $p(Y_N^0 \mid \alpha^0)$ after Box & Jenkins, (1976). The dynamic and standard AR(1) likelihood functions were used to evaluate the rainfall likelihood $p(Y_N^1 \mid Y_N^0, \alpha^1)$. The Shuffled Complex Evolution (SCE) algorithm (Duan *et al.*, 1992) was used to maximise the joint likelihood.

$$p(Y_N^{0,1} \mid \alpha^0, \alpha^1) \propto p(Y_N^1 \mid Y_N^0, \alpha^1)\, p(Y_N^0 \mid \alpha^0) \qquad \textbf{(9)}$$

4. RESULTS

4.1. Maximum Likelihood Estimates

The maximum likelihood estimates (MLE) for the H4 model parameters calibrated to the IPO index and Stroud rainfall data are shown in **Table 1**.

Table 1 MLE Parameters for the H4 Model Calibrated to the IPO and Stroud rainfall

Parameter	Calibrated MLE
μ_{IPO}	-0.23
σ_{IPO}	0.07
$\phi 1_{IPO}$	1.91
$\phi 2_{IPO}$	-0.93
μ_{w_I}	516 mm
σ_{w_I}	204 mm
μ_{d_I}	319 mm
σ_{d_I}	149 mm
$\phi 1_I$	0.14
μ_{NI}	756 mm
σ_{NI}	213 mm
$\phi 1_{NI}$	0.13

Comparisons of the seasonal H4 model are made to the standard AR(1) model calibrated to annual data. Parameter estimates for the annual AR(1) model were 1152 mm, 297 mm and 0.17 for the mean, standard deviation of the random perturbations and the lag-1 autocorrelation respectively.

4.2. Evaluation using distributional and temporal statistics

The annual rainfall distribution probability plot for the simulated H4 rainfall and observed record is shown in **Figure 8**. The observed data is within the 90% probability limits; similar results were obtained for the seasonal rainfall distributions and the annual AR(1) model. Spearman lag-1 autocorrelation values for the observed annual data and the annual AR(1) and H4 models are 0.16, 0.21 and 0.16 respectively.

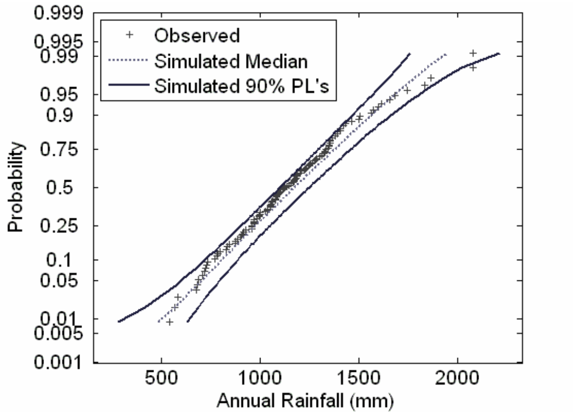

Figure 8 Probability plot of simulated and observed rainfall distributions

The Hurst coefficient has been used extensively in the hydrological and economic literature to detect long-term persistence with h>½ being an indicator of a process with long-term dependence. Salas (1993) warns against using the Hurst coefficient to select between stochastic models due to the fact that some models, such as ARMA processes, can have long-term dependence structure but will result in asymptotic Hurst coefficients of h = 0.5.

An alternative method was used to detect long-range dependence which was to examine the cumulative distribution of run-lengths of the cumulative sum of the rainfall anomalies from a large number of replicates each with the same sample size as the observed data. This and related statistics have received considerable attention in the statistical literature (Luceno & Puig-pey, 2000). The distributions are compared in **Figure 9**. The distributions for cumulative probabilities less than 0.6 were alike for the simulated and observed so the plot concentrates on the differences in the distributions at the upper end of the runlength values. From the distributions of run-lengths of cumulative sums, there appears to be some difference between the AR(1) and H4 models beyond run-lengths of approximately 30 years. For example for a run-length of 60 years, there is a probability of around 3.1% that the H4 model has run-lengths of the cumulative sum of greater than 60 years. The AR(1) model has less than 0.01% chance of a run-length of greater than 60 years.

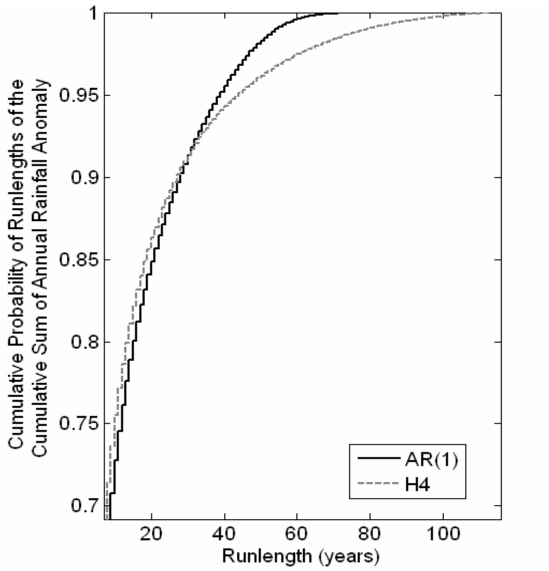

Figure 9 Cumulative distributions of run-lengths of cumulative sum of rainfall anomaly for simulated data

4.3. Drought Risk Analysis

The drought risk analysis was undertaken using an annual reservoir simulation, with characteristics similar to Tillegra Dam. Simulations of annual stochastic data for Stroud rainfall were transformed using a regression between annual Stroud rainfall and annual Tillegra runoff (adjusted $R^2 = 0.73$). **Figure 10** compares the probabilities of encountering storage levels less than 5% and 45% for a range of annual demands for the annual AR(1), H4 and an annual independent model. The results show that both annual AR(1) and H4 provide a similar drought risk which is higher then the independent model. For some of the annual demands the H4 model results in slightly higher drought risks than the annual AR(1) model.

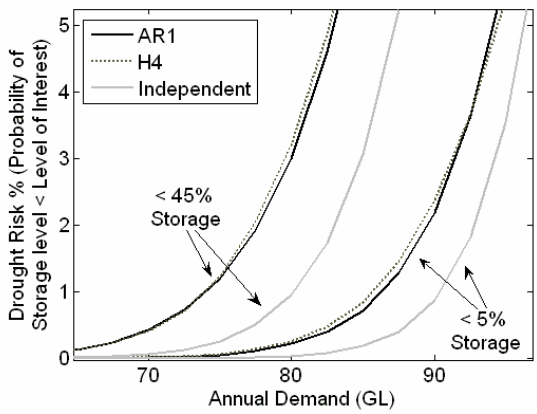

Figure 10 Simulated long-term drought risk for a reservoir with a capacity of 477.63 GL

5. DISCUSSION

The H4 model introduced in this paper utilises stochastic processes to explicitly incorporate climatic indices which characterise long-term climate processes by informing a seasonal stochastic rainfall model. The aim of this work is to improve the representation of long-term variability in Australian hydrological data.

Analysis of the H4 model showed that it was able to reproduce key observed statistics that were not used in calibration, such as the annual rainfall distribution, however at this stage there is no clear improvement on the characterisation of the temporal statistics.

The selection of a definitive impact season still remains elusive. The technique used in this paper to choose the impact season with the maximum p-value ignores other statistically significant impact seasons. Further research is required to identify the most appropriate choice of the impact season.

The distributions of run-length of the cumulative sum of the rainfall anomaly were compared for the annual AR(1) and H4 models. The H4 model demonstrated a higher probability of longer runs for run-lengths greater than 30 years however a key difficulty with the comparison of long-dependence statistics is the short observed record.

In terms of long-term drought risk, the H4 model gave similar results to the annual AR(1) model, although at some demands the H4 model showed slightly higher drought risks.

The long-term statistics and drought risk were evaluated at the annual time scale. For the H4 model, the impact season was only 35% of annual total. Climatologically, the current implementation of the hierarchical model could be improved by the addition of key climate mechanisms such as ENSO and the IOD which typically have sub-decadal cycles. It is possible that drought risk is more sensitive to these mechanisms than the multi-decadal effect of the IPO. The ignorance of climate mechanisms during the non-impact period is a shortcoming of the H4 model.

6. CONCLUSION

The ability of the Bayesian hierarchical framework to incorporate climate data to inform a seasonal stochastic rainfall model was demonstrated. The model was calibrated to the Stroud rainfall, located in close proximity to Hunter Water Corporation's water supply catchments. Impact-seasons of the climate mechanism (the IPO) were detected by statistical analysis, allowing the data to dictate the model structure. Model calibration was undertaken using maximum likelihood techniques. Evaluation of the calibrated model showed it was able to reproduce observed statistics not used in calibration, such as the annual distribution. This evaluation included a comparison with the commonly used annual AR(1) model.

The cumulative distribution of run-lengths of the cumulative sum of the rainfall anomaly is presented as an alternative measure of long-term dependence in stochastic models. For the case study presented in this paper this approach showed that the H4 model has a higher probability of run-lengths longer than 30 years than the AR(1) model. Drought risks were found to be slightly higher for the H4 model than the AR(1) model for some annual demand scenarios, though the differences are relatively small. When comparing the H4 model to the AR(1) model no distinct improvement in capturing the observed statistics could be found for the case study used here.

This framework has the capability of utilising data other than hydrological sources to characterise climate variability on multiple time-scales. The augmentation of observed data is the major advance on techniques that focus on simulating low frequency components of rainfall time series. Further research will utilise full parameter uncertainty techniques. Sub-decadal variability will be incorporated into the framework by stochastic processes for the El Niño Southern Oscillation and Indian Ocean Dipole phenomena. The inclusion of palaeo-data in formal observation processes will further augment the observed record. The availability of palaeo reconstructions of the IPO and PDO makes a hierarchical framework of climate mechanisms an attractive simulation alternative.

7. ACKNOWLEDGEMENTS

The authors would like to acknowledge Assoc. Prof. Stewart Franks for his guidance and ongoing assistance with this project, the Australian Research Council, Hunter Water Corporation (the industry sponsors), Dr. S. Power and UKMO for providing the IPO data, the BOM Australia for the rainfall data and the reviewers of this paper for their helpful insights and welcomed corrections.

8. REFERENCES

Box, G.E.P., & Jenkins, G.M. (1976), *Time series analysis: Forecasting and control*. Holden-Day.

Duan, Q. Y., Sorooshian, S., & Gupta, V. (1992), Effective and efficient global optimization for conceptual rainfall-runoff models. *Water Resources Research*, **28**(4), 1015–1031.

Henley, B.J., Thyer, M.A., Kuczera, G., & Franks, S. (2006), Incorporating long-term climate information into stochastic models of annual hydrological data: A Bayesian Hierarchical approach. *In: 30th Hydrology and Water Resources Symposium, Launceston, Tas.*

Kiem, A. S., & Franks, S. W. (2001), On the identification of ENSO-induced rainfall and runoff variability: a comparison of methods and indices. *Hydrological Sciences Journal - Journal des Sciences Hydrologiques*, **46**(5), 715–727.

Lavery, B. Joung, G. & Nicholls, N., (1997), An extended high-quality historical rainfall dataset for Australia. *Australian Meteorological Magazine*, **46**, 27-38

Luceno, A., & Puig-pey, J. (2000), Evaluation of the run-length probability distribution for cusum charts: Assessing chart performance. *Technometrics*, **42**(4), 411–416.

Salas, J.D. (1993), *Handbook of Hydrology*. McGraw-Hill. Chap. 19 "Analysis and modeling of hydrologic time series", pages 19.1–19.72.

Sharma, A., Tarboton, D. G. & Lall, U. (1997), Streamflow simulation: A Nonparametric Approach. *Water Resources Research*, **33**(2), 291–308.

Thyer, M. & Kuczera, G. (2000), Modeling long-term persistence in hydroclimatic time series using a hidden state markov model. *Water Resources Research*, **36**(11), 3301–3310.

Verdon, D.C., Wyatt, A.M., Kiem, A. & Franks, S. (2004), Multidecadal variability of rainfall and streamflow: Eastern Australia. *Water Resources Research*, **40**.

Whiting, J.P., Lambert, M.F., & Metcalfe, A.V. (2003), Modelling persistence in annual Australian point rainfall. *Hydrology and Earth System Sciences*, **7**(2), 197–211.