

Chapter 9 Extensions to the standard HMM formulation

The results presented thus far in this thesis have been produced through the implementation of two-state and three-state HMMs to describe hydroclimatic persistence. Although the standard HMM formulation is parsimonious, it can also be developed in a variety of ways to model accurately different characteristics of hydrologic data. The previous chapter showed non-parametric HMMs to be a valuable extension to the standard HMM formulation. The model developments that are described in this chapter relax various basic assumptions of HMMs, producing time series models that may improve the descriptions of certain persistent time series.

The first aspect of the HMM framework analysed here is the efficacy of the Markovian assumption for climatic state series. Secondly, the issue of conditional independence of rainfall data is discussed. The third investigation focuses upon identifying the correct scale at which persistence should be analysed, and the development of a novel hierarchical HMM is presented. The time series models produced through these developments are then calibrated to the deseasonalised monthly rainfall for Sydney as an initial example of their practicality.

9.1 Incorporating explicit state durations

9.1.1 Hidden semi-Markov model (HSMM) structure

A major assumption of conventional HMMs fitted to discrete-timed data is that state durations follow a geometric distribution, which is inherent in the Markovian assumption as discussed in Section 3.3.1. For many physical series this geometric density may be inappropriate (Rabiner, 1989), and an explicit duration density such as the negative binomial may be preferred. The negative binomial distribution is a two parameter distribution that decays at a slower rate than the geometric. The longer tail of such a distribution allows longer durations in each state, which produces a more persistent process. If random independent trials result in a success (state transition) with probability p (the HMM transition probability), the distribution of X , which is the number of the trial on which the first success occurs, constitutes the geometric distribution. The number of trials however to produce k successes (where $k > 1$) is provided by the negative binomial distribution, given as

$$p(X) = \binom{X-1}{k-1} p^k (1-p)^{X-k} \quad (9.1)$$

with k taking integer values. The negative binomial distribution is a useful density in this context, as it includes the geometric distribution as a special case when $k = 1$. For the case of

$k = 2$, the probability mass function of the negative binomial distribution is compared to the geometric distribution in Figure 9.1, with each being generated using a transition probability of $p = 0.3$. This figure shows higher values of X , analogous to longer state durations, to have a higher probability of occurring with the negative binomial distribution than with the geometric.

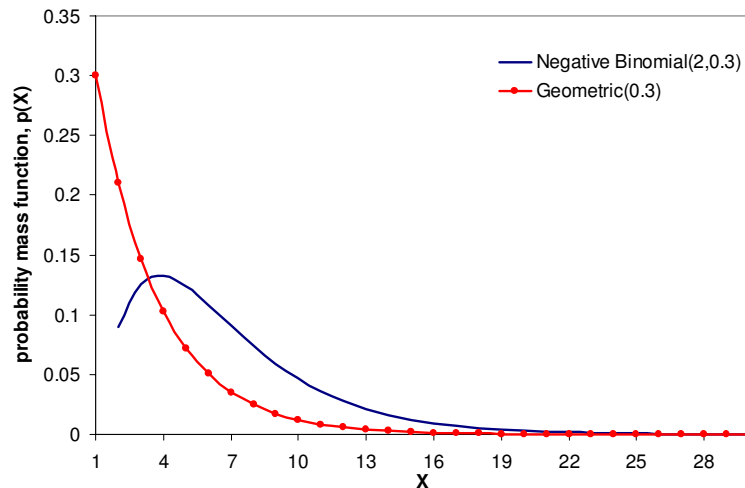


Figure 9.1 Probability mass function for negative binomial and geometric distributions

The inclusion of a specific duration density transforms a conventional HMM to a hidden semi-Markov model (HSMM), also termed a variable duration HMM. When adopted to describe hydrological persistence, state conditional distributions defining observations in conventional HMMs are retained in this extended framework. The number of discrete model states and the form of state conditional distributions are separate modelling assumptions, such that two-state HSMMs are the simplest implementation. When an explicit duration density is incorporated into a HMM, self-transition probabilities take a value of zero, and transitions between states only occur after specific numbers of observations defined by duration densities. It is clear then that by setting the explicit duration density to a geometric, the HSMM is equivalent to a conventional HMM.

Hidden semi-Markov models have been described for various meteorological applications, mainly daily rainfall modelling (eg Sansom, 1998; 1999), yet have received little attention for modelling persistence at monthly or annual scales. Ferguson (1980) presented the original investigation into variable duration models for discrete time series, mostly in speech recognition, although the estimation algorithm presented by Ferguson is computationally expensive. Levinson (1986) presented an efficient method to calculate the joint probability distribution of a sequence of observations, evaluating the HSMM likelihood by adapting the forward-backward algorithm of conventional HMMs. Yu and Kobayashi (2003a; 2003b) further refined this estimation method by proposing a new forward-backward algorithm that dramatically reduces the calculation complexity of the HSMM likelihood. This algorithm is

used later in this work. The SCE algorithm is used to evaluate MLEs for the HSMM parameters, with parameter uncertainty evaluated with the Adaptive Metropolis algorithm as with conventional HMMs. Models with lognormal and gamma state conditional distributions are calibrated to the Sydney deseasonalised monthly rainfall, using negative binomial duration densities.

9.1.2 Calibration of a two-state lognormal HSMM to Sydney monthly rainfall

Two-state lognormal HMMs provide suitable descriptions of persistence in the series of monthly rainfall anomalies for Sydney; however the generalisation of state duration probabilities for this series has not been addressed up to this point. In order to investigate the efficacy of this model to describe hydrological persistence in Australian rainfall, two-state HSMMs with negative binomial duration densities are now fitted to Sydney data. With state conditional distributions of the HMM framework more closely approximating random draws from lognormal distributions than gamma distributions within the Sydney series, lognormal distributions are also assumed to describe monthly observations in the HSMM. The eight parameters for this model are estimated with the SCE algorithm, and 60,000 Metropolis samples are generated to estimate parameter uncertainty. The posterior distributions for the parameters of the lognormal conditional distributions of the HSMM are summarised in Table 9.1.

Table 9.1 Comparison of posteriors for parameters of conditional distributions for two-state lognormal HSMM and two-state lognormal HMM, with median and 90% credibility interval

	Wet state location	Dry state location	Wet state scale	Dry state scale
Two-state lognormal HMM	4.814 (4.67, 4.97)	3.945 (3.78, 4.15)	0.555 (0.50, 0.62)	0.575 (0.51, 0.66)
Two-state lognormal HSMM	4.840 (4.63, 4.98)	4.253 (4.15, 4.32)	0.561 (0.47, 0.66)	0.698 (0.67, 0.73)

The posterior median estimates of the HSMM wet state parameters are very similar to those from calibrating the HMM, whereas the dry state parameters are both higher than the HMM parameters, with narrower credibility intervals. Figure 9.2 shows the wet and dry state conditional distributions for this model, using the posterior medians as parameter estimates. The higher dry state parameters values for the HSMM produce a dry state distribution with higher mean, variance and skew than the HMM dry state.




Figure 9.2 Conditional distributions from the calibrations of a two-state lognormal HMM and a two-state lognormal HSMM to the deseasonalised monthly rainfall for Sydney

The calibration of the HSMM estimates the wet state negative binomial density to have order 2 and the wet state order 4. The posterior distributions for the real-valued probabilities for the two negative binomial distributions are shown in Figure 9.3. Median values of the wet state and dry state distributions are 0.631 and 0.430 respectively, with the dry state showing much tighter credibility bounds.




Figure 9.3 Posterior distributions for probabilities associated with negative binomial duration densities from the calibration of a two-state lognormal HSMM to the deseasonalised monthly rainfall for Sydney, with medians shown

With negative binomial densities of order greater than 1, it is difficult to describe persistence in a manner similar to transition probabilities in a HMM. The probability mass functions associated with the wet and dry state distributions are shown in Figure 9.4 alongside geometric distributions, with each described by the posterior medians. The minimum durations within the wet and dry states of the HSMM are 2 and 4 months respectively. Although the decay of the HSMM wet state distribution is quite rapid, the dry state demonstrates a much stronger persistence than the dry state of the HMM, with a higher probability for all durations greater than 5 months.

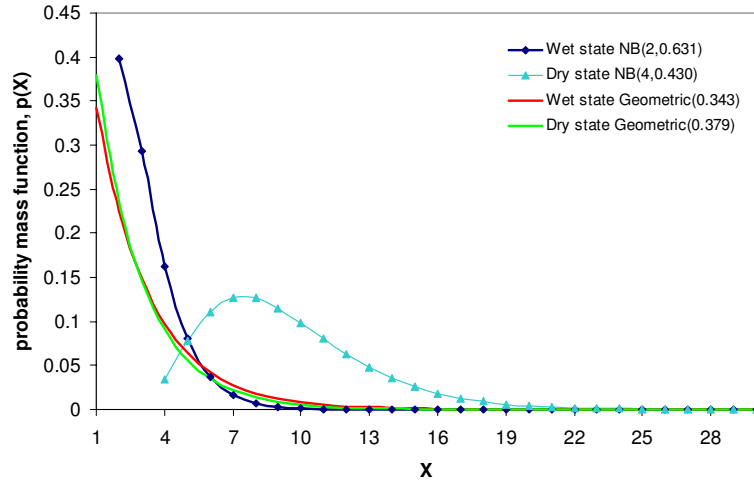


Figure 9.4 Probability mass functions for duration densities from the calibrations of a lognormal HMM and a lognormal HSMM with negative binomial duration densities to the deseasonalised monthly rainfall for Sydney

In order to calculate the marginal likelihood of the HSMM, uniform prior distributions over the interval (1, 6) were taken for the integer order of the negative binomial distributions, with uniform priors over (0, 1) used for the negative binomial probabilities. This Bayesian model selection produces an estimate of $\ln BF_{HMM,HSMM} = 3.6$, which in this instance marginally favours the simpler model.

9.2 Incorporating temporal dependence in observations

9.2.1 Autoregressive hidden Markov model (ARHMM) structure

Conventional HMMs assume that a sequence of observations are conditionally independent, being estimated as random draws from a defined parametric distribution. This assumption can also be extended to include alternative modelling approaches, by considering an observation at time t in terms of the following time series regression

$$y_t = \beta_{s_t} a_t + z_t \quad (9.2)$$

where a_t is a vector of past observations and z_t is a series of independently and identically distributed variation. The notion of underlying model states conditioning the observation series is incorporated in this model by specifying different values of the β parameters for each discrete value of the state variable x_t . A change in model state realises a change in regression model parameters. For first-order dependence in observations, the previous equation reduces to

$$y_t = \beta_{s_t^j} y_{t-1} + z_t \quad (9.3)$$

By estimating transition probabilities for movement between model states, the basic modelling assumptions of the HMM are preserved and this produces the class of autoregressive HMMs (ARHMMs), which are useful time series models accounting for changes in regime. This class of models reduce to standard autoregressive models when remaining in a constant state. Furthermore, the relationship between a first-order ARHMM and a conventional HMM is shown through representing the former model as the Bayesian net diagram in Figure 9.5. In this diagram, circles represent the unobserved state variables and squares show observed values. The thick arrows indicate Markovian transitions, with the thin red lines representing conditional relationships between variables. In a conventional HMM, the relationships designated with red lines are absent.




Figure 9.5 Bayesian net diagram for first-order ARHMM

The ARHMM class of models have received a broad application in the econometrics literature (eg Goldfeld and Quandt, 1973; Hamilton, 1989; Hamilton, 1990) as they are suitable for modelling nonstationarities that are due to abrupt changes of regime in the economy (Bengio, 1999). Like HSMs, these are also particularly suited to speech recognition (Juang and Rabiner, 1985; Rabiner, 1989), yet they have received little attention in the field of stochastic hydrology. Parameter estimation in ARHMMs adopts the standard forward-backward algorithm of conventional HMMs by replacing independent observations densities with an autoregressive density. As a consequence, the distribution of observations in a first-order ARHMM $P(y_t | s_t^j, y_{t-1})$ takes a specific form such as a Gaussian distribution whose mean is a linear function of the previous observation y_{t-1} .

9.2.2 Calibration of a two-state lognormal ARHMM to Sydney monthly rainfall

The ARHMM provides a mechanism to incorporate regime changes into an autoregressive model. With the Sydney monthly rainfall series displaying significant persistence at a monthly scale, the inclusion of temporal dependence in observed values may improve the description of monthly variability. In order to determine a suitable order of this model, it is necessary to first observe its autocorrelation. The natural logarithm of scaled monthly anomalies produces a time

series that approximates random draws from a Gaussian distribution $N(4.402, 0.713^2)$. Within the series of log values, only the first order autocorrelation is significant, having a value 0.093 with standard error 0.024. As a result, a two-state ARHMM with lognormal state conditional distributions and first-order autocorrelation in each state is calibrated to the Sydney monthly series. The posterior distributions for parameters of these conditional distributions are summarised in Table 9.2 alongside posteriors from a standard two-state lognormal HMM.

Table 9.2 Comparison of posteriors for parameters of conditional distributions for two-state lognormal ARHMM and two-state lognormal HMM, with median and 90% credibility intervals

	Wet state location	Dry state location	Wet state scale	Dry state scale
Two-state lognormal HMM	4.814 (4.67, 4.97)	3.945 (3.78, 4.15)	0.555 (0.50, 0.62)	0.575 (0.51, 0.66)
Two-state lognormal ARHMM	4.932 (4.77, 5.19)	4.136 (3.91, 4.32)	0.498 (0.36, 0.58)	0.637 (0.55, 0.70)

Posterior medians in Table 9.2 suggest that conditional distributions for this model are similar to those obtained from both the two-state HMM and two-state HSMM, with greater dry state variability that is consistent with the latter. Posterior distributions for the first-order autocorrelation parameters are shown in Figure 9.6. The dry state autocorrelation has a posterior median value of 0.080, which includes the autocorrelation of the marginal distribution within its 90% credibility interval (0.002, 0.161). Although this parameter estimate appears consistent with the observed features of the data series, the autocorrelation within the wet state has a posterior median of -0.084. The broad posterior distribution for this parameter however has a 90% credibility interval (-0.217, 0.038) that includes zero, suggesting that this parameter may not have statistical significance.

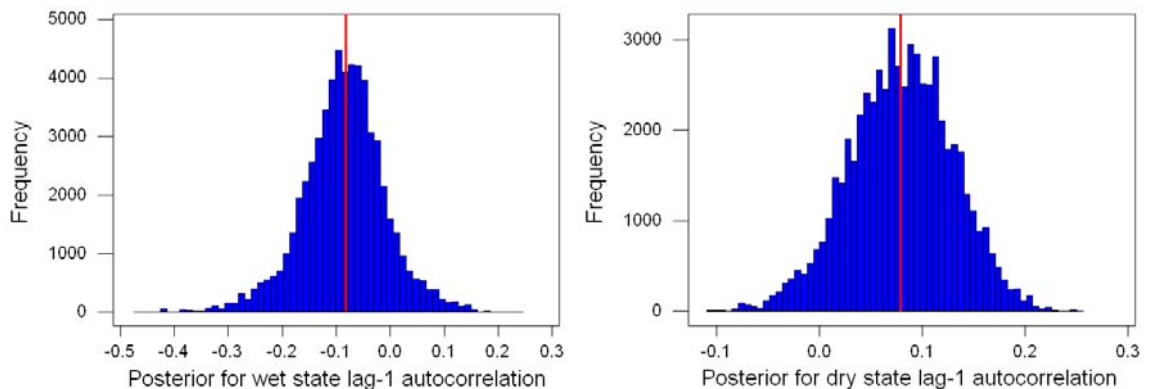


Figure 9.6 Posterior distributions for lag-1 autocorrelations estimated from the calibration of a two-state lognormal ARHMM to the deseasonalised monthly rainfall for Sydney, with medians shown

The estimation of ARHMM parameters in the Sydney monthly series shows that observed autocorrelation is modelled only within a dry state, suggesting that this model might have too many parameters. In order to evaluate the Bayes Factor between this model and the simpler two-state HMM, conjugate priors described earlier are again used with Gaussian priors $N(\mu_\phi, \sigma_\phi^2)$ described for autocorrelations, with μ_ϕ the estimated lag-1 autocorrelation and σ_ϕ the standard error of this estimate. This produces an estimate of $\ln BF_{HMM, ARHMM} = 1.8$, which marginally favours the simpler model.

The posterior distribution of $P_{WD} + P_{DW}$ for the first-order ARHMM is shown in Figure 9.7, with a median of 0.761 that is similar to the median from fitting a two-state lognormal HMM to these data. This distribution however has a wider 90% credibility interval (0.568, 1.092) than the HMM, and since this includes a value of 1, evidence against persistence is obtained at a 10% level. The inclusion of autocorrelation in the HMM therefore appears to mask the two-state persistence that is evident in the monthly rainfall series for Sydney; this is likely due to the weak autocorrelation of this time series. This modelling approach may indeed provide an improved description of persistence for hydrologic series showing strong temporal dependence.

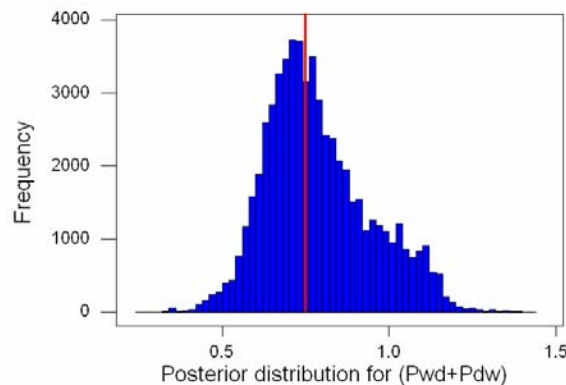


Figure 9.7 Posterior distribution for the sum of transition probabilities from the calibration of a two-state lognormal ARHMM to the deseasonalised monthly rainfall for Sydney, with median

9.3 Analysing persistence at multiple time scales simultaneously

9.3.1 Hierarchical hidden Markov model (HHMM) structure

The results presented in this section are from the calibration of various HMMs to time series of both annual and monthly totals. With results indicating that either there is insignificant persistence at an annual scale, or that there is simply insufficient data available to detect such persistence, the most suitable models have used monthly data. From a climatic perspective, the physical processes producing hydrological persistence at a monthly scale are very different from

those producing annual variability. ENSO episodes tend to persist for 15 months on average, during which the monthly rainfall totals for parts of Australia are in turn amplified and moderated. The issue of whether annual rainfall totals are also influenced by this climatic mode depends upon the periods over which annual periods are defined.

In order to describe the interaction of hidden state processes at various scales, a hierarchical HMM (HHMM) is described here to simultaneously fit climate states operating over both annual and monthly periods. In this model approach, regional climates are assumed to fluctuate between wet and dry states over annual time scales, and within these states there exists higher frequency variability with fluctuating wet and dry months. By defining the annual wet and dry states as aW and aD , and monthly states mW and mD , four model states are then estimated within the time series of monthly observations: $(mW_t | aW)$, $(mD_t | aW)$, $(mW_t | aD)$ and $(mD_t | aD)$. Persistence within the monthly states of this model is governed by two transition probabilities for months in a wet year (defined as $P_{mWD|aW}$ and $P_{mDW|aW}$) and two for months in a dry year ($P_{mWD|aD}$ and $P_{mDW|aD}$). Annual persistence is controlled by the values of two more transition probabilities, P_{aWD} and P_{aDW} , producing six unknown probabilities that need to be estimated in model calibration. It is clear that by remaining in a constant annual state, this model can degenerate to the simple two-state monthly HMM. The relationship between these various model states is shown in Figure 9.8.

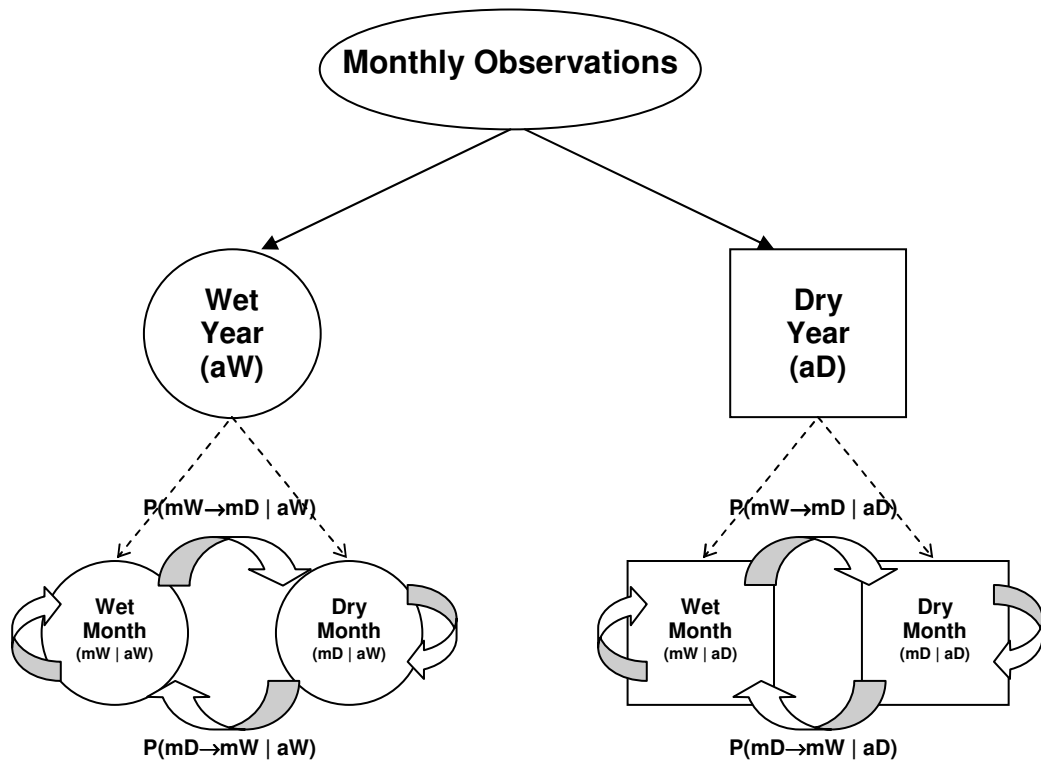


Figure 9.8 Framework for the hierarchical HMM

Although hidden states are defined at both monthly and annual scales, state conditional distributions are only defined with monthly values. One set of monthly parameters are associated with an annual wet state (θ_w) with one set associated with an annual dry state (θ_d). In calibrating this model, monthly totals (y_t) are assumed to approximate random draws from either lognormal or gamma distributions. The probability of rainfall in month t being generated in a wet monthly state of a wet year is expressed as $p(y_t | mW_t, aW)$ and likewise for other combinations. In order to calibrate the HHMM, the standard HMM forward-backward likelihood function is evaluated using the monthly data series with parameters for an annual wet state and then separately with the monthly parameters associated with an annual dry state.

Since the climate is assumed in this model to fluctuate between two states at an annual level, a method to adapt the monthly series is developed. The application of the forward-backward likelihood function at an annual scale requires the probability of total rainfall in year i , A_i , to be evaluated in each state, $p(A_i | aW_i)$ and $p(A_i | aD_i)$. Although these probabilities are not defined explicitly, monthly pdfs are used to facilitate their calculation, adopting the relationship between an annual total A_i and the sequence of 12 monthly totals within the same period. An annual pdf is therefore replaced with a probability of observing a sequence of twelve monthly totals as:

$$p(A_i | aW_i) = p(y_{(i-1) \times 12 + 1}, y_{(i-1) \times 12 + 2}, \dots, y_{(i-1) \times 12 + 12} | aW_i) = p(Y_i | aW_i) \quad (9.4)$$

which is also the likelihood for a 12-month sequence given an annual climate state. The monthly forwards variables from the existing implementation of the HMM likelihood function are multiplied together over 12-month sequences to obtain this monthly likelihood. Annual forwards and backwards variables are calculated from these annual pdfs and the HMM likelihood function can proceed. The posterior probabilities for the climate being in each annual state are evaluated, and these then weight the two monthly posterior state series to produce an overall monthly state series. For models that contain two annual states and with monthly observations following lognormal, Gaussian or gamma distributions in each of two discrete states, a total of 14 parameters require estimation.

9.3.2 Calibration of a two-state lognormal hierarchical HMM to Sydney monthly rainfall

Results shown in Chapter 6 indicated that there is no significant persistence in the series of annual rainfall totals for the Sydney gauge. This observation was reinforced by the calibration of a two-state Gaussian HMM, which failed to reject the possibility of the marginal distribution being constructed from a mixture of two Gaussian distributions without temporal persistence.

This contrasts with the analysis of monthly totals, which provided evidence for two-state persistence at this scale. The HHMM framework develops the assumptions of the standard two-state HMM by allowing monthly persistence to be conditioned by annual variability. The calibration of the HHMM to the Sydney monthly rainfall series assumes monthly data in each annual state to be lognormally-distributed such that $P(\ln(y_t) | s_t^j) \approx N(\mu_j, \sigma_j^2)$. The posterior distributions for monthly parameters of the HHMM are summarised in Table 9.3.

Table 9.3 Summary of posterior distributions for parameters of wet and dry years from a two-state lognormal HHMM, showing medians and 90% credibility intervals

	P_{WD}	P_{DW}	Wet state location	Dry state location	Wet state scale	Dry state scale
Wet	0.334 (0.21, 0.46)	0.398 (0.24, 0.51)	4.806 (4.67, 4.94)	3.914 (3.76, 4.11)	0.556 (0.50, 0.62)	0.559 (0.49, 0.65)
Dry	0.400 (0.02, 0.92)	0.520 (0.07, 0.95)	4.895 (2.84, 10.99)	3.432 (1.28, 7.06)	0.645 (0.20, 1.36)	0.887 (0.21, 1.44)

A main feature of these results is that the credibility intervals around posterior medians in dry years are much wider than for monthly parameters from a wet year. Furthermore the posterior medians for wet year parameters closely match estimates from simple two-state lognormal HMMs. It is straightforward to show that the HHMM can degenerate to the simpler model by remaining within a single annual state, such that there is not annual persistence. Annual transition probability estimates for the HHMM indeed support this result. The annual P_{WD} has a posterior median of 0.010, with 90% credibility interval (0.002, 0.152), which contrasts dramatically with the posterior distribution for P_{DW} , which has median 0.547 and 90% credibility interval (0.080, 0.951). These estimates are vastly different from the posterior medians obtained from the calibration of a Gaussian HMM to the time series of annual totals, and strongly indicate that the most likely model structure is predominantly within a single wet state. The posterior distribution of the monthly $P_{WD} + P_{DW}$ in a wet annual state has a median of 0.732 with 90% credibility interval (0.601, 0.834), similar to the interval size obtained from the standard two-state HMM. The posterior distribution for monthly $P_{WD} + P_{DW}$ in a dry annual state however shows a median of 0.926 with a 90% credibility interval (0.323, 1.608) that indicates that monthly persistence is not observed in dry years.

These results demonstrate that the absence of significant annual persistence causes the more complex HHMM to relax to the simpler model. After calculating the marginal likelihoods for both the lognormal HHMM and the lognormal HMM, the Bayes Factor between the models is estimated as $\ln BF_{HMM,HHMM} = 95.3$, which is strongly in favour of the two-state monthly model, and supports the inference made from parameter estimates. Although this development

of the monthly HMM was inappropriate for Sydney data, monthly hydrologic series showing the HHMM to be a superior model to the HMM are described in Chapter 10.

9.4 Summary of chapter

This chapter has introduced a range of stochastic models that are derived from the standard HMM formulation. Although hidden semi-Markov models and autoregressive HMMs are not new developments, these models have rarely been used for describing hydrological persistence at either a monthly or annual scale. The results from calibrating these models to the Sydney monthly data in this chapter indicate that these models may improve descriptions of hydrological persistence in certain data series. The ARHMM in particular provides a useful connection between the widely-used ARMA time series models and the HMM approach that better defines hydrological persistence in terms of climatic influences. The hierarchical HMM is a novel approach to modelling hydrological persistence that can account for different levels of hydroclimatic interactions. The efficacy of these models is further investigated in the following chapter, using a range of hydrologic data.

Chapter 10 Utilising NP HMMs to identify appropriate parametric models for persistent data

Previous chapters have established that a monthly time scale is more appropriate than an annual scale for the identification and explanation of hydrological persistence. In this chapter non-parametric HMMs estimate underlying probability distributions needed for the analysis of persistence in the monthly-scale hydrology of Australia. Using a range of observed data series, HMMs and variants described in Chapter 9 are employed to identify hydrological persistence, and to evaluate both its strength and its relationship to climate fluctuations. This chapter provides an overview of the benefits to stochastic modelling in hydrology provided by parametric and non-parametric HMMs, particularly as models for monthly rainfall data.

10.1 Developing HMMs to model hydrological persistence in spatial rainfall

The deseasonalised monthly rainfall for the four meteorological districts introduced in Chapter 4 (Districts 9A, 16, 27 and 71), are analysed in this section to demonstrate the method by which the non-parametric HMM formulation can augment information obtained from established parametric methodology to identify hydroclimatic persistence. As described in Table 4.1 and Figure 4.2, these districts are representative of the four main rainfall regimes across this country.

Two-state NP HMMs are calibrated to the deseasonalised monthly rainfall for each of these four districts. The posterior medians and 90% credibility intervals for transition probabilities from these calibrations are presented in Table 10.1. These results provide strong evidence for two-state persistence in each district, with 95% credibility limits for their sum of transition probabilities being well below 1. The credibility interval around an estimate for the sum of transition probabilities for District 9A does not include 1 and two-state persistence is therefore significant at a 10% level. However, the upper limit for this interval is close to 1 which indicates persistence to be much weaker than the other districts. These observations are consistent with those shown in Table 5.1, which demonstrate strong persistence in each of these monthly rainfall series using an array of runs statistics. Furthermore in this Chapter 4 analysis, the deseasonalised monthly rainfall from District 9A showed the weakest persistence of runs either side of the median value, reflecting the results from NP HMM calibration.

Table 10.1 Posterior medians and 90% credibility intervals for transition probabilities from calibrating two-state NP HMMs to deseasonalised monthly rainfall

District	P_{WD}	P_{DW}	$P_{WD} + P_{DW}$
9A	0.258 (0.078, 0.553)	0.363 (0.167, 0.668)	0.673 (0.358, 0.948)
16	0.210 (0.106, 0.330)	0.244 (0.143, 0.360)	0.452 (0.301, 0.642)
27	0.247 (0.144, 0.360)	0.197 (0.096, 0.299)	0.445 (0.254, 0.635)
71	0.141 (0.084, 0.234)	0.252 (0.178, 0.346)	0.400 (0.289, 0.532)

Following the calibration of two-state NP HMMs, state conditional distributions are estimated through random samples taken around posterior median estimates of the partition locations. Anderson-Darling goodness-of-fit statistics for these distributions are shown in Table 10.2 for Gaussian, lognormal and gamma distributions. These statistics illustrate that for five of the eight conditional distributions shown, lowest AD values are obtained for fitting lognormal distributions, suggesting that these samples most closely represent random draws from lognormals. However seven of the eight lognormal AD values are still too high to be consistent with this parametric form at a 5% statistical level.

Table 10.2 Anderson-Darling goodness-of-fit statistics for estimates of state conditional distributions from the calibration of two-state NP HMMs to deseasonalised monthly rainfall

District		Gaussian distribution	Lognormal distribution	Gamma distribution
9A	Wet	20.81	3.50	7.23
	Dry	5.73	3.14	2.19
16	Wet	40.78	1.36	6.91
	Dry	34.35	4.67	9.52
27	Wet	23.49	5.63	11.92
	Dry	23.94	19.29	18.59
71	Wet	19.63	0.41	2.35
	Dry	6.57	3.29	0.82

The calibration of NP HMMs to the deseasonalised monthly rainfall series allows the form of parametric two-state HMMs to be estimated, and this is often desirable as parametric models may provide a superior fit. Using the results of Table 10.2, two-state lognormal HMMs are now calibrated to each of the four monthly rainfall series. Transition probability estimates are then analysed in Table 10.3 to determine whether significant two-state persistence is identified through these parametric models. These results indicate that even by imposing a parametric form on state conditional distributions two-state persistence is identified within each of these four monthly rainfall series.

Table 10.3 Posterior medians and 90% credibility intervals for transition probabilities from the calibration of two-state lognormal HMMs to deseasonalised monthly rainfall

District	P_{WD}	P_{DW}	$P_{WD} + P_{DW}$
9A	0.091 (0.039, 0.226)	0.506 (0.254, 0.730)	0.607 (0.329, 0.875)
16	0.205 (0.093, 0.361)	0.175 (0.084, 0.288)	0.386 (0.252, 0.549)
27	0.154 (0.118, 0.194)	0.411 (0.334, 0.499)	0.566 (0.469, 0.670)
71	0.103 (0.063, 0.164)	0.233 (0.154, 0.329)	0.339 (0.238, 0.459)

The similarity between the calibrations of the NP HMMs and parametric HMMs is now investigated through the posterior state series of each model. Linear correlations between the median state series of the NP HMM, Lognormal HMM and Gamma HMM for each district are shown in Table 10.4. Each correlation is highly significant, $p < 0.001$ in each case, which shows that two-state parametric HMMs identified through the calibration of the NP HMMs identify similar persistence to the latter. The correlations are strongest for District 71, which suggests that a two-state lognormal HMM describes this series better than it can for any of the other three series.

Table 10.4 Linear correlations between median state series from calibrating various HMMs to deseasonalised monthly rainfall

District	NP HMM and Lognormal HMM	NP HMM and Gamma HMM	Lognormal HMM and Gamma HMM
9A	0.556	0.677	0.440
16	0.982	0.856	0.813
27	0.704	0.787	0.928
71	0.974	0.986	0.928

Another method to demonstrate relationships between the two-state persistence identified with the non-parametric and parametric HMMs is to analyse their respective associations with other measures of climate variability. After evaluating rank correlations between the monthly NINO3 index and the median state series for each of the NP, lognormal and gamma HMMs, District 27 shows the strongest correlations for the four samples analysed. For this monthly rainfall series, the correlation between the NINO3 series and the NP HMM state series ($r = -0.349$) has greater magnitude than correlations achieved for either the two-state lognormal HMM ($r = -0.274$) or two-state gamma HMM ($r = -0.295$). This suggests that ENSO variability is revealed most clearly through the calibration of the NP HMM. In Section 4.2.1, it was noted that District 27 was predicted by NINO3 the most clearly of all districts.

The relationship between the HMM calibrations and descriptors of regional climate variability can be further examined by segregating the series of months over the period (1913-2002) on the basis of ENSO phase (as defined by the 5-mem values of the NINO3 series) and most likely climate state from calibrations of the three models. The numbers of months obtained through this separation are presented in Table 10.5. These results demonstrate that for the NP HMM almost 78% of El Niño months coincided with likely dry states and 65% of La Niña months coincided with wet states. These biases reflect expected hydroclimatic interactions for the ENSO phenomenon, and demonstrate the strong influence that this source of climatic variability has upon monthly rainfall in District 27. Although the lognormal HMM identified a strong bias towards La Niña months coinciding with wet states, it was unable to distinguish clearly between wet and dry states during El Niño periods. However, with the Gamma HMM showing bias in both La Niña and El Niño periods that opposed those identified with the NP HMM, it appears that for these monthly data the lognormal HMM is a superior parametric model. The relationships between state series and NINO3 values for the other districts were much less significant than for District 27, such that their results offer little assistance to elucidate the improvements gained from calibrating parametric as opposed to non-parametric HMMs.

Table 10.5 Numbers of months in which most probable HMM states from the calibrations of two-state NP, lognormal and gamma HMMs to the deseasonalised monthly rainfall for District 27 coincide with ENSO phases

		El Niño	ENSO Neutral	La Niña
NP HMM	Wet state	71	216	174
	Dry state	249	277	93
LN HMM	Wet state	162	372	233
	Dry state	158	121	34
Gamma HMM	Wet state	201	186	71
	Dry state	119	307	196

Following the analyses of two-state persistence in these monthly rainfall series, it is pertinent to investigate whether such time series also demonstrate significant three-state climatic persistence. Three-state NP HMMs are calibrated to the four series, with Table 10.6 showing posterior medians and 90% credibility intervals for each HMM transition probability. The series of monthly anomalies from District 71 produces the lowest value for three of these probabilities, reinforcing the results in Table 10.3 that indicated this series of monthly values to be the most persistent of the four districts investigated. Furthermore the six transition probability estimates for District 9A, shown to be the least persistent series when observing two-state persistence, are higher than at least two other districts.

Table 10.6 Posterior medians and 90% credibility intervals for transition probabilities from the calibration of three-state NP HMMs to the deseasonalised monthly rainfall

	P_{WN}	P_{WD}	P_{NW}	P_{ND}	P_{DW}	P_{DN}
9A	0.231 (0.04, 0.46)	0.232 (0.07, 0.46)	0.275 (0.08, 0.54)	0.194 (0.03, 0.50)	0.246 (0.05, 0.50)	0.303 (0.09, 0.55)
16	0.194 (0.07, 0.32)	0.113 (0.02, 0.28)	0.114 (0.01, 0.36)	0.207 (0.08, 0.43)	0.148 (0.05, 0.28)	0.150 (0.02, 0.31)
27	0.146 (0.05, 0.25)	0.147 (0.04, 0.30)	0.149 (0.04, 0.28)	0.136 (0.05, 0.27)	0.459 (0.23, 0.62)	0.309 (0.13, 0.46)
71	0.241 (0.14, 0.34)	0.077 (0.01, 0.17)	0.086 (0.02, 0.21)	0.162 (0.08, 0.27)	0.196 (0.09, 0.29)	0.105 (0.02, 0.22)

The sums of self-transition probabilities from the calibrations of three-state NP HMMs to the deseasonalised monthly rainfall in each of these districts are summarised in Table 10.7. These sums show evidence of significant three-state persistence in the monthly rainfall observations of Districts 16, 27 and 71, with 90% Bayesian credibility intervals well above a value of 1. The weak evidence for monthly persistence in District 9A rainfall is reinforced in these results, with a lack of evidence at a 10% level for significant three-state persistence. Bayes Factors from calibrating two-state and three-state NP HMMs to each series show that the simpler models are more appropriate for the monthly rainfall in Districts 9A, 16 and 27. District 71 shows a low value of $\ln BF_{2NP,3NP}$ providing only weak evidence in favour of the two-state model such that the three-state NP HMM is an appropriate description of monthly persistence in District 71.

Table 10.7 Posterior medians and 90% credibility intervals for the sums of self-transition probabilities from the calibration of three-state NP HMMs, and Bayes Factors comparing the calibrations of two-state NP HMMs to the calibrations of three-state NP HMMs

	$P_{WW} + P_{NN} + P_{DD}$	$\ln BF_{2NP,3NP}$
9A	1.446 (0.984, 1.906)	5.5
16	2.028 (1.531, 2.368)	5.3
27	1.634 (1.373, 2.003)	5.4
71	2.116 (1.832, 2.338)	0.4

To describe parametric three-state HMMs for each of these series, it is again helpful to estimate the form of the state conditional distributions. Using posterior medians for partition locations to firstly estimate these distributions, Anderson-Darling goodness-of-fit statistics for fitting three continuous distributions are presented in Table 10.8. For 11 of the 12 state conditional distributions shown in this table, Gaussian distributions are less appropriate than either

lognormal or gamma. Lognormal distributions provide the lowest AD statistics for 8 of these distributions, suggesting that a three-state lognormal HMM may adequately model three-state persistence as identified with the three-state NP HMM.

Table 10.8 Anderson-Darling goodness-of-fit statistics for estimates of state conditional distributions from the calibration of three-state NP HMMs to the deseasonalised monthly rainfall

		Gaussian distribution	Lognormal distribution	Gamma distribution
9A	Wet	13.67	2.84	4.58
	Neutral	16.78	4.77	6.90
	Dry	6.71	2.09	1.75
16	Wet	21.10	1.38	2.30
	Neutral	34.03	3.26	5.71
	Dry	46.30	4.50	12.15
27	Wet	32.38	8.42	16.81
	Neutral	34.28	17.24	24.03
	Dry	9.83	10.64	8.28
71	Wet	5.39	1.89	0.47
	Neutral	8.43	1.32	0.84
	Dry	14.53	2.09	2.28

Using the information from the calibrations of three-state NP HMMs to the deseasonalised monthly rainfall series, three-state lognormal HMMs are now calibrated. Transition probabilities from these models are summarised in Table 10.9, with District 71 displaying the lowest values from these four series for five of the six variables, a result that indicates three-state persistence to be strongest in the District 71 series.

Table 10.9 Posterior medians and 90% credibility intervals for transition probabilities from the calibration of three-state lognormal HMMs to the deseasonalised monthly rainfall

	P_{WN}	P_{WD}	P_{NW}	P_{ND}	P_{DW}	P_{DN}
9A	0.367 (0.08, 0.68)	0.173 (0.02, 0.53)	0.153 (0.03, 0.46)	0.161 (0.03, 0.47)	0.164 (0.02, 0.43)	0.360 (0.07, 0.72)
16	0.056 (0.01, 0.14)	0.171 (0.08, 0.29)	0.226 (0.08, 0.42)	0.264 (0.08, 0.53)	0.334 (0.14, 0.47)	0.293 (0.13, 0.57)
27	0.136 (0.02, 0.46)	0.140 (0.06, 0.22)	0.285 (0.10, 0.57)	0.239 (0.12, 0.38)	0.166 (0.04, 0.31)	0.235 (0.10, 0.38)
71	0.337 (0.18, 0.49)	0.050 (0.01, 0.15)	0.065 (0.02, 0.19)	0.091 (0.05, 0.16)	0.147 (0.07, 0.25)	0.102 (0.02, 0.22)

The appropriateness of three-state parametric HMMs, with lognormal or gamma conditional distributions, to describe persistence in each district is now examined. Table 10.10 summarises posterior distributions for the sum of self-transition probabilities from the calibration of each model. These results show that for Districts 27 and 71 in particular, the three stochastic models identify persistence of similar magnitude, with District 71 displaying the strongest three-state

persistence using all three models. This result is consistent with the estimation of two-state persistence in these monthly rainfall series. Even though the calibration of a three-state lognormal HMM to the monthly rainfall of District 9A identifies significant persistence at a 10% significance level, this is not supported through the calibration of a three-state gamma HMM. A three-state lognormal HMM is therefore a suitable model for this monthly rainfall series, as it potentially improves the NP HMM estimates of transition probabilities.

Table 10.10 Posterior medians and 90% credibility intervals for the sum of self-transition probabilities from the calibration of various three-state HMMs

	NP HMM	Lognormal HMM	Gamma HMM
9A	1.446 (0.98, 1.91)	1.482 (1.12, 1.88)	1.376 (0.80, 1.83)
16	2.028 (1.53, 2.37)	1.615 (1.29, 1.90)	1.863 (1.38, 2.23)
27	1.634 (1.37, 2.00)	1.770 (1.29, 2.07)	1.565 (1.14, 1.96)
71	2.116 (1.83, 2.34)	2.180 (1.90, 2.38)	2.069 (1.81, 2.30)

Although Table 10.10 summarises the calibrations of three-state HMMs to these monthly rainfall series, it is pertinent to investigate Bayesian model selection results from the calibrations of various parametric models. Using a two-state lognormal HMM as a reference model, the calibrations of other models are compared with Bayes Factors, which are shown in Table 10.11. Negative values for a particular model demonstrate its superiority to a two-state HMM with lognormal state conditional distributions. The best models for each monthly time series, identified by negative Bayes Factors of the highest magnitude, are shown as red. These results provide an interesting demonstration of how different stochastic models are appropriate for different monthly time series.

Table 10.11 Bayes Factors ($\ln BF_{M,LN}$) comparing the calibrations of various parametric models (M) to the calibrations of two-state lognormal HMMs to the deseasonalised monthly rainfall

Model	9A	16	27	71
2-state Gamma HMM	-38.3	-2.3	-2.4	-0.5
3-state Lognormal HMM	-44.0	-1.9	-10.6	-4.0
3-state Gamma HMM	-39.3	-2.7	-8.0	-2.9
AR(1) to natural logarithms	-31.4	3.3	55.3	10.3
AR(2) to natural logarithms	-29.9	-1.3	52.7	6.1
AR(3) to natural logarithms	-29.2	-1.6	51.7	9.8
2-state Lognormal ARHMM lag-1	45.3	47.8	8.4	-2.0
3-state Lognormal ARHMM lag-1	192.3	11.5	-5.1	5.3
2-state Lognormal HHMM	14.5	10.4	236.4	-3.3
2-state Lognormal HSMM	-37.1	-1.7	-20.3	-5.8
2-state Gamma HSMM	-29.4	1.9	5.9	8.4

It has been noted that the time series of monthly anomalies for District 71 shows strong two-state and three-state persistence. The strength of this persistence is demonstrated through Bayes Factors that provide evidence in favour of HMMs as opposed to autoregressive models. Large positive BF values for District 27 also indicate that autoregressive models fail to improve the description of these monthly anomalies, although this result is expected due to the weak temporal dependence in this series. The hierarchical HMM provides an improved description of the monthly persistence for District 71, although for the other time series the simpler two-state lognormal HMM is superior. It is suggested that the HHMM is a suitable model for time series that show strong two-state persistence.

The two-state lognormal HMM is an inadequate description of hydrological persistence in District 9A. This is supported by the TP estimates from this model, which showed a tendency to remain in a single state for a majority of values. However when this series is calibrated with a gamma HMM, estimates of the two transition probabilities are almost equal, indicating similar persistence in each of the two climate states. A three-state lognormal HMM is the most appropriate parametric model for District 9A, and this is suggested through the sum of self-transition probabilities, which has a 90% credibility interval that exceeds unity. Alternatively, the corresponding sum for the calibrations of a three-state NP HMM and a three-state gamma HMM includes a value of 1 within its 90% credibility interval.

The lognormal HSMMs that are the most appropriate models for both District 27 and 71 include negative binomial (NB) distributions that have integer values of 2 for the dry state distributions in each series and geometric distributions for the wet states. Probabilities associated with the wet state NB distributions for Districts 71 and 27 (0.077 and 0.119 respectively) are similar to posterior medians for P_{WD} from the calibrations of two-state lognormal HMMs to both series (0.103 and 0.154). This indicates that these models identify similar wet state distributions to the two-state lognormal HMMs.

The section has presented a comprehensive estimation of persistence within the monthly rainfall data for four meteorological districts that represent the main rainfall regimes of Australia. In light of the results of Chapter 5 that showed monthly streamflow data to be more persistent than rainfall data, a similar analysis is undertaken with River Murray flow data in the following section.

10.2 Developing HMMs to model hydrological persistence in streamflow

Following the application of NP HMMs in the previous section to provide unbiased estimates of hydrological persistence in various district-averaged rainfall series, the focus of this section is to use these models to investigate the persistence within the monthly flow record for the River Murray. Deseasonalised monthly streamflows for the Murray were calibrated with two-state and three-state NP HMMs in Section 8.4.3. These results indicated strong persistence in monthly flows, with wet and dry climate states persisting on average for approximately 12 and 11 months respectively with the former model. The estimated state conditional distributions, using 2000 samples around posterior median values from the calibration of a two-state NP HMM to these data (as detailed in Section 8.4.3) are shown in Figure 10.1. These plots illustrate that a dry state distribution approximates a series of Gaussian variates, with the wet state distribution being heavily skewed.

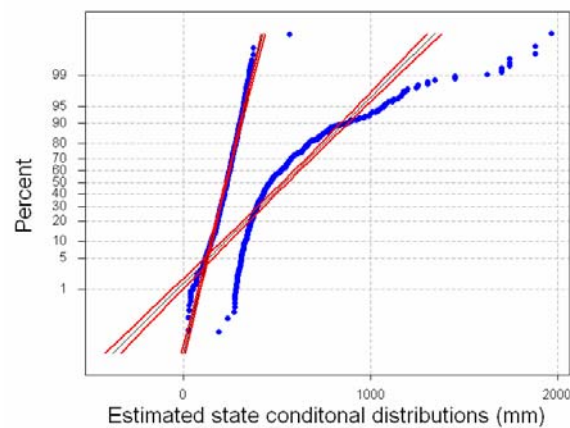


Figure 10.1 Gaussian probability plot showing estimates for state conditional distributions from the calibration of a two-state NP HMM to the deseasonalised monthly Murray flows

To analyse the parametric structure of the wet state distribution in Figure 10.1, Table 10.12 shows Anderson-Darling goodness-of-fit statistics for various parametric forms. These statistics show that conditional distributions are inconsistent with random draws from either Gaussian, lognormal or gamma distributions, with the probability of samples following these distributions rejected at 5% significance levels. These statistics demonstrate the complex nature of two-state persistence in these monthly flow data.

Table 10.12 Anderson-Darling goodness-of-fit statistics for estimates of state conditional distributions from the calibration of a two-state NP HMM to deseasonalised monthly Murray flows

State	Gaussian distribution	Lognormal distribution	Gamma distribution
Wet	55.61	15.66	30.31
Dry	3.23	33.91	13.66

The calibration of the two-state NP HMM to the deseasonalised monthly Murray flows is now analysed in terms of its association to measures of broad-scale climate variability. The time series of monthly NINO3 values were shown in Section 4.2.4 to explain more of the variability in the time series of deseasonalised monthly streamflows in the Murray than monthly rainfall series, so it is possible that the significant two-state persistence in this series is also more closely related to ENSO variability than observed with district-averaged rainfall data. Table 10.13 shows the results from segregating months in this flow record on the basis of ENSO phase and most likely climate state using median state probabilities, and indicate a tendency towards more dry state months during El Niño periods and more wet states during La Niñas. Both features are expected for the condition of ENSO being a dominant aspect of hydrological persistence in these monthly flow data.

Table 10.13 Numbers of months in which most probable HMM states from the calibration of two-state NP HMMs to the deseasonalised monthly Murray flows coincide with ENSO phases

	El Niño	ENSO Neutral	La Niña
Wet state	153	316	201
Dry state	242	242	142

Model selection results in Section 8.4.3 showed that a three-state NP HMM is superior to a two-state NP HMM for describing persistence in monthly Murray flows. To further investigate the nature of this three-state persistence, the form of state conditional distributions is examined. After taking 3000 random samples around the posterior medians for the two partitions in the three-state NP HMM, estimates of the wet, neutral and dry state distributions are obtained in the size ratio 0.24: 0.38: 0.38. Table 10.14 shows Anderson-Darling goodness-of-fit statistics from fitting Gaussian, lognormal and gamma distributions to these estimates. The neutral state distribution produces the lowest AD statistic for each parametric form, likely due to the fact that this distribution does not contain extreme values of the marginal that can complicate descriptions of parametric distributions. The AD statistics for each parametric form are clearly rejected at a 5% statistical level for both the wet state and the dry state estimates, indicating that parametric models may not be developed easily.

Table 10.14 Anderson-Darling goodness-of-fit statistics for estimates of state conditional distributions from calibrating three-state NP HMMs to deseasonalised monthly Murray flows

State	Gaussian distribution	Lognormal distribution	Gamma distribution
Wet	33.37	12.85	20.01
Neutral	8.04	5.05	5.02
Dry	12.55	53.87	27.42

The relationships between three-state persistence and ENSO variability are again used to investigate the performance of the three-state NP HMM. Table 10.15 shows the results from separating months on the basis of ENSO phase and most likely climate state. El Niño episodes are most often modelled as dry climate states, consistent with the ENSO influence upon these data. Furthermore ENSO neutral phases show a tendency to coincide with neutral climate states more often than either wet states or dry states, and La Niña episodes slightly favour wet climate states. This latter observation reinforces the fact that although ENSO is an important influence upon these monthly data, it is not the only source of climatic persistence.

Table 10.15 Numbers of months in which most probable HMM states from the calibrations of three-state NP, lognormal and gamma HMMs to the deseasonalised monthly Murray flows coincide with ENSO phases

	El Niño	ENSO Neutral	La Niña
Wet state	53	161	124
Neutral state	140	221	108
Dry state	202	196	115

As described in Section 3.2, autoregressive moving average (ARMA) time series models are regularly used to describe hydrologic time series. The most appropriate ARMA model for the monthly Murray series is determined after transforming monthly values to approximate a series of Gaussian variates such that a series of Gaussian residuals is produced. The scaled deseasonalised monthly streamflows are shown on a lognormal probability plot in Figure 10.2 as an approximate straight line, suggesting consistency with random draws from a lognormal distribution.




Figure 10.2 Lognormal probability plot showing scaled deseasonalised monthly Murray flows

Although this series has an Anderson-Darling goodness-of-fit statistic of 2.93, which is rejected at a 5% significance level, it is likely that the logarithms of these deseasonalised data are appropriate for estimating the order of the most suitable ARMA model. A useful method for

estimating which autoregressive and moving average terms are suggested in the log-transformed data is to examine both the autocorrelation function (acf) and partial autocorrelation function (pacf). The correlogram for these data, in Figure 10.3, shows large spikes at initial lags that decay slowly to zero, which is indicative of an autoregressive process. Furthermore a partial correlogram having significant spikes at only the first three lags, as shown in Figure 10.4, suggests that an AR(3) model is the most appropriate ARMA model for the log-transformed series. The correlogram for the series of residuals from fitting this model is shown in Figure 10.5. Without significant autocorrelations (or partial autocorrelations, which are not shown) at any lags, it is clear that the AR(3) model removes serial dependence in this series.




Figure 10.3 Correlogram for natural logarithms of deseasonalised monthly Murray flows




Figure 10.4 Partial correlogram for natural logarithms of deseasonalised monthly Murray flows




Figure 10.5 Correlogram for residuals after fitting an AR(3) model to the natural logarithms of deseasonalised monthly Murray flows

With an AR(3) model being appropriate for the Murray flow series, it is important to investigate the role of persistence that was identified through NP HMMs. In essence, this reduces to investigating whether modelling this series with an AR(3) removes the need to model its persistence directly? One approach is to examine whether the time series of residuals from fitting an AR(3) to the log-transformed series retains significant persistence. By using the NP HMM as an exploratory tool to analyse the structure of this time series, it is clear that the influence of temporal persistence cannot always be removed through accounting solely for autocorrelation in a time series. Importantly the calibration of NP HMMs to these residuals indeed demonstrates persistence at both two-state and three-state levels, with 90% credibility intervals around the sums of self-transition probabilities being greater than unity for each model. Model selection results demonstrate weak evidence in favour of a simpler two-state model as the most appropriate description of these residuals ($BF_{2NP,3NP} = 1.92$).

These results suggest that a model that combines both time series regression with hidden state persistence, such as an ARHMM may indeed be the most appropriate model for this monthly flow series. These models are compared to the various other models introduced throughout this section for their description of monthly streamflows in the Murray through Bayesian model selection. Table 10.16 summarises the Bayes Factors for a range of stochastic models, comparing these to two-state lognormal HMMs.

Table 10.16 provides some interesting model selection results, with negative values for $\ln BF$ indicating that different models are superior to two-state lognormal HMMs for describing the time series. Once again the uniform-transformed data used within the calibration of NP HMMs cannot be included in this comparison, as the likelihood is dependent upon the values of the marginal distribution. The comparison of parametric models in Table 10.16 shows that multiple-state HMMs and the HHMM are improved models for this time series, indicating that persistence is strong. Furthermore both AR(1) and AR(3) models are superior to the two-state HMM, illustrating that temporal dependence is a significant characteristic of the monthly streamflow series. Importantly, the AR(3) is shown to be superior to the AR(1), reinforcing observed characteristics described in Figure 10.4. However the various ARHMMs shown in Table 10.16 are the best descriptors of the monthly persistence. The most appropriate model shown in bold in Table 10.16 combines third-order autocorrelation as suggested previously with the benefit of three model states. This clearly demonstrates that the monthly flow series for the Murray has both strong persistence and temporal dependence, a combination requiring a complex modelling approach. This section has provided further evidence for how unbiased estimations from non-parametric HMMs may inform the description of parametric time series models that develop the standard HMM framework.

Table 10.16 Model selection results from calibrations to the deseasonalised monthly Murray flows

	Maximum likelihood	Number of parameters	$\ln BF$
2-state Lognormal HMM	-8252.3	6	
2-state Gamma HMM	-8234.3	6	-19.7
2-state Gaussian HMM	-8293.5	6	56.9
3-state Lognormal HMM	-8069.5	12	-157.9
3-state Gamma HMM	-8053.1	12	-177.6
4-state Lognormal HMM	-7864.0	20	-324.3
4-state Gamma HMM	-7865.6	20	-338.2
2-state Lognormal HSMM	-8248.0	8	-14.5
2-state Gamma HSMM	-8238.7	8	-10.8
Lognormal HHMM	-8057.3	14	-170.4
AR(1) to logarithms	-7621.5	3	-638.5
AR(3) to logarithms	-7560.4	5	-689.6
2-state LN ARHMM lag-1	-7499.3	8	-747.5
3-state LN ARHMM lag-1	-7437.8	15	-773.1
2-state LN ARHMM lag-3	-7435.3	12	-762.6
3-state LN ARHMM lag-3	-7357.7	21	-834.1
3-state LN ARHMM lag-4	-7358.0	24	-827.7
4-state LN ARHMM lag-3	-7333.3	32	-813.6
5-state LN ARHMM lag-3	-7325.1	45	-777.4

10.3 Summary of chapter

This chapter has focused upon the use of non-parametric HMMs that were introduced in Chapter 8 and variants of the standard HMM formulation described in Chapter 9 alongside two-state and three-state parametric HMMs, in order to provide improved descriptions of hydrological persistence. The NP HMM estimates underlying probability distributions in the spatially-averaged monthly rainfall for four meteorological districts of Australia. Calibrations of parametric models then showed that extensions to the standard HMM formulation presented in Chapter 9 produced models that were superior descriptors of persistence in these data. The second case study analysed the strong persistence in the monthly streamflow record of the River Murray. Although the NP HMM estimated that underlying probability distributions in these data approximated lognormals, two-state and three-state parametric HMMs were inferior to more complex models such as a HSMM and a hierarchical HMM. The most suitable model for this monthly series combined three-state persistence with temporal dependence. These results show the potential for ARHMMs to provide accurate descriptions of hydrologic data that demonstrate strong hydrological persistence. Furthermore the NP HMM approach has been shown to inform the accurate calibration of parametric HMMs as a model for hydrological persistence.

Chapter 11 Simulating persistent hydrologic data with HMMs: Three case studies

The final investigation into the validity of HMMs as hydrological time series models focuses upon their application to simulate hydrologic records. Simulation has an important role in stochastic hydrology, supporting decision-making throughout water resource management. Autoregressive moving average (ARMA) models are widely-used for the simulation of hydrologic time series such as monthly and annual rainfall series; however HMMs have been rarely applied in this context. Hydrologic systems have components that act as random variables (McCuen and Snyder, 1986), and simulation provides a means for examining aspects of such systems. For example the management of a reservoir to provide optimal release patterns may be compromised through deficiencies in the length of input rainfall data, yet this can be improved through the simulation of such data. HMMs are an expedient method for describing time series that have significant persistence, and their application as a device for simulating such series is investigated in this chapter. Three case studies are investigated, with HMMs calibrated to three monthly time series that have very different statistical characteristics.

11.1 Background to hydrologic simulation

Simulation may be viewed as a mechanism for investigating the characteristics of a system (such as the hydroclimatic cycle) through a model rather than through the system itself. Historical rainfall data are sample series originating from statistical populations of unknown structure. In order to make decisions about these hydrologic data, such as conducting risk analyses for water supply systems, characteristics of the input population need to be determined, and simulation can provide this requirement.

The accuracy of decisions made from simulations depend upon the accuracy of both the model description (for an analysis of HMMs this will be related to the strength of persistence in the observed time series), and the probability description of its random elements. Simulations depend strongly upon the generation of random numbers that follow correct probabilistic distributions. The method used here closely follows the calibration approach previously described, thus engaging the Metropolis output to incorporate uncertainty into the values of model parameters. In the simulation of a time series described by a two-state Gaussian HMM (such as a time series of annual rainfall totals), the model is first calibrated and posterior distributions for the mean and standard deviation of variates in each state and the two transition

probabilities are evaluated. Multiple simulations (1000 in the examples described here) having identical length to the original time series are then generated using different parameter values for each series, sampling these values from the posteriors. This provides a method for incorporating random variation not only in the Gaussian-distributed variates but also in the first and second moments of such variates.

This section investigates the simulation of three hydrologic time series recorded at monthly scales. The HMM family are analysed here for their efficacy in providing accurate simulations of monthly hydrologic data. The simulations provide information about the distribution of the original series, and HMMs are compared to other stochastic models through the uncertainty around estimated distributions. Each simulated series is placed into ranked order, such that for each of the ranked values of the original time series, 1000 estimates are produced. The original ranked data are then placed on a probability plot, and a median estimate together with a 90% confidence interval is produced from the 1000 estimates of each datum. The accuracy of simulations is assessed through the size of the uncertainty bounds around the marginal distribution.

Although these simulations are focused on a monthly scale, it is beneficial that the characteristics of totals aggregated from original monthly values are also accurately simulated. To test this, each simulated time series is aggregated over six-month, annual, two-year and five-year periods. Statistical features for each aggregation are then compared with the aggregation of the original monthly series. The incorrect estimation of rainfall variability at a range of longer time scales is a significant criticism of models designed for higher-frequency rainfall simulation. Frost (2003) postulated that through the failure of daily rainfall models to account explicitly for long-term climatic persistence, a significant underestimate of the variability of annual rainfall results. A similar issue is faced when developing stochastic models for the simulation of monthly time series, which act as input to a variety of hydrological systems. It is suggested here that useful models of persistence will provide a superior approach for the simulation of monthly data. With each of the three monthly records in this section demonstrating significant two-state or three-state persistence, the capability of each stochastic model to simulate adequately the persistence characteristics of the original series is also evaluated. A number of the runs statistics described in Section 5.1 are therefore calculated for each monthly simulation, being compared with observed values.

11.2 Simulating monthly rainfall for District 71

The first time series simulated is the deseasonalised monthly rainfall series from District 71. This series demonstrated the strongest persistence of the four district-averaged series analysed in Section 10.1, both in terms of two-state and three-state persistence. Using Bayesian model

selection, the two-state lognormal HSMM was shown in Table 10.11 to be the best model for this series, and is subsequently used to simulate the monthly series. Comparisons are then made to simulations from an AR(2) model that is calibrated to the natural logarithms of the deseasonalised monthly series. Bayes Factors showed the AR(2) model to be more appropriate than either an AR(1) or an AR(3) model for the series of natural logarithms. Furthermore the series of autocorrelations and partial autocorrelations for these data, the latter shown in Figure 11.1, have a pattern that is consistent with an AR(2) model.




Figure 11.1 Partial correlogram for the time series of natural logarithms of deseasonalised monthly rainfall for District 71

The residuals from the calibration of an AR(2) to the series of natural logarithms are shown as an approximate straight line on a Gaussian probability plot in Figure 11.2. The Anderson-Darling goodness-of-fit statistic for these residuals 0.94 is consistent with a series of Gaussian variates, such that the modelling assumptions of an AR(2) model are justified.

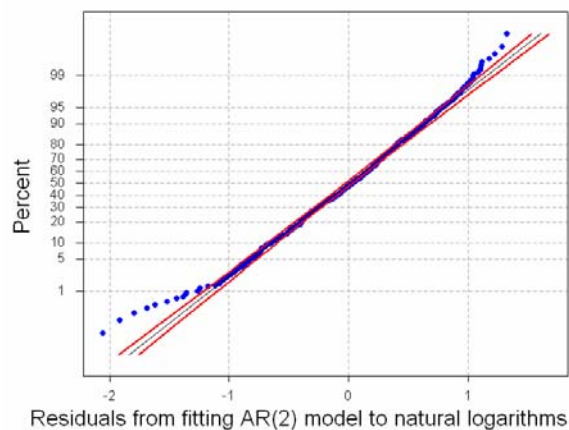


Figure 11.2 Gaussian probability plot showing residuals from fitting an AR(2) model to the natural logarithms of the deseasonalised monthly rainfall for District 71

The mean, standard deviation and skew for the deseasonalised monthly series, along with the various aggregations of this series from using a one-month moving window are summarised in Table 11.1. These statistics are then calculated for each of the 1000 simulations of length 1080 from the two stochastic models, and for the various aggregations of such simulations. The

median of these 1000 values and the interval that includes 90% of the values are also shown for each aggregation level. These results show that these two stochastic models produce accurate simulations of the monthly rainfall series, although both approaches over-estimate skew at a monthly scale. The monthly means and standard deviations are higher for simulations from the AR(2) model, and this leads to higher values of these statistics for the various aggregations. From the HSMM, twelve of the fifteen sample moments shown in Table 11.1 lie within 90% confidence intervals, which is one fewer than the number of sample moments accurately simulated with the AR(2).

Table 11.1 Statistics for the time series of deseasonalised monthly totals from District 71 over various aggregations calculated with one-month moving windows, along with the median and 90% confidence interval of statistics from 1000 simulations using a two-state lognormal HSMM and an AR(2) model with lognormal residuals

		Mean	Standard deviation	Skew
1m	Sample	77.06	37.74	0.974
	HSMM	76.32 (72.8, 79.4)	39.02 (36.2, 42.2)	1.286 (0.99, 1.81)
	AR(2)	77.54 (73.5, 81.3)	42.56 (38.6, 47.9)	1.715 (1.34, 2.70)
6m	Sample	462.39	116.82	0.564
	HSMM	458.04 (436.8, 476.6)	113.29 (102.6, 125.6)	0.270 (0.57, 0.02)
	AR(2)	465.23 (441.1, 487.6)	130.83 (115.4, 152.9)	0.818 (0.50, 1.38)
1yr	Sample	925.17	174.42	0.474
	HSMM	916.26 (874.2, 954.1)	164.54 (145.2, 187.2)	0.183 (-0.13, 0.55)
	AR(2)	930.08 (882.3, 974.6)	191.12 (164.6, 228.0)	0.584 (0.25, 1.12)
2yr	Sample	1850.05	268.91	0.510
	HSMM	1832.62 (1749.5, 1907.3)	233.82 (197.7, 275.8)	0.142 (-0.25, 0.55)
	AR(2)	1859.03 (1762.7, 1950.9)	275.31 (228.0, 334.0)	0.414 (0.01, 0.89)
5yr	Sample	4611.45	490.79	0.191
	HSMM	4579.44 (4374.3, 4767.8)	358.36 (277.5, 464.2)	0.071 (-0.47, 0.64)
	AR(2)	4650.87 (4412.0, 4876.4)	426.87 (334.3, 555.8)	0.219 (-0.26, 0.84)

Simulations from these two stochastic models are compared to the original data through Gaussian probability plots in Figure 11.3. More accurate simulations are characterised by distributions that match closely the shape of observed values, with narrow envelopes within which a majority of the simulated values lie. The deseasonalised monthly series is shown together with the distribution of 12-monthly totals (aggregated from the entire series using a 12-month moving window) and the distribution of consecutive five-year totals. After ranking both the observed series and the series of aggregated values, intervals containing 90% of simulations are obtained for each value, and across the entire distribution these intervals are shown as red dashed lines. The series of median values from these 1000 simulations is shown in these plots as a blue line. Figure 11.3 indicates that the lognormal HSMM provides an improved simulation of the marginal series, with the 90% confidence interval including more of the observed values than is achieved from AR(2) simulations. The distributions of annual and five-yearly sums of

the deseasonalised monthly series are simulated adequately with each modelling approach, with observed data lying predominantly within 90% confidence intervals.

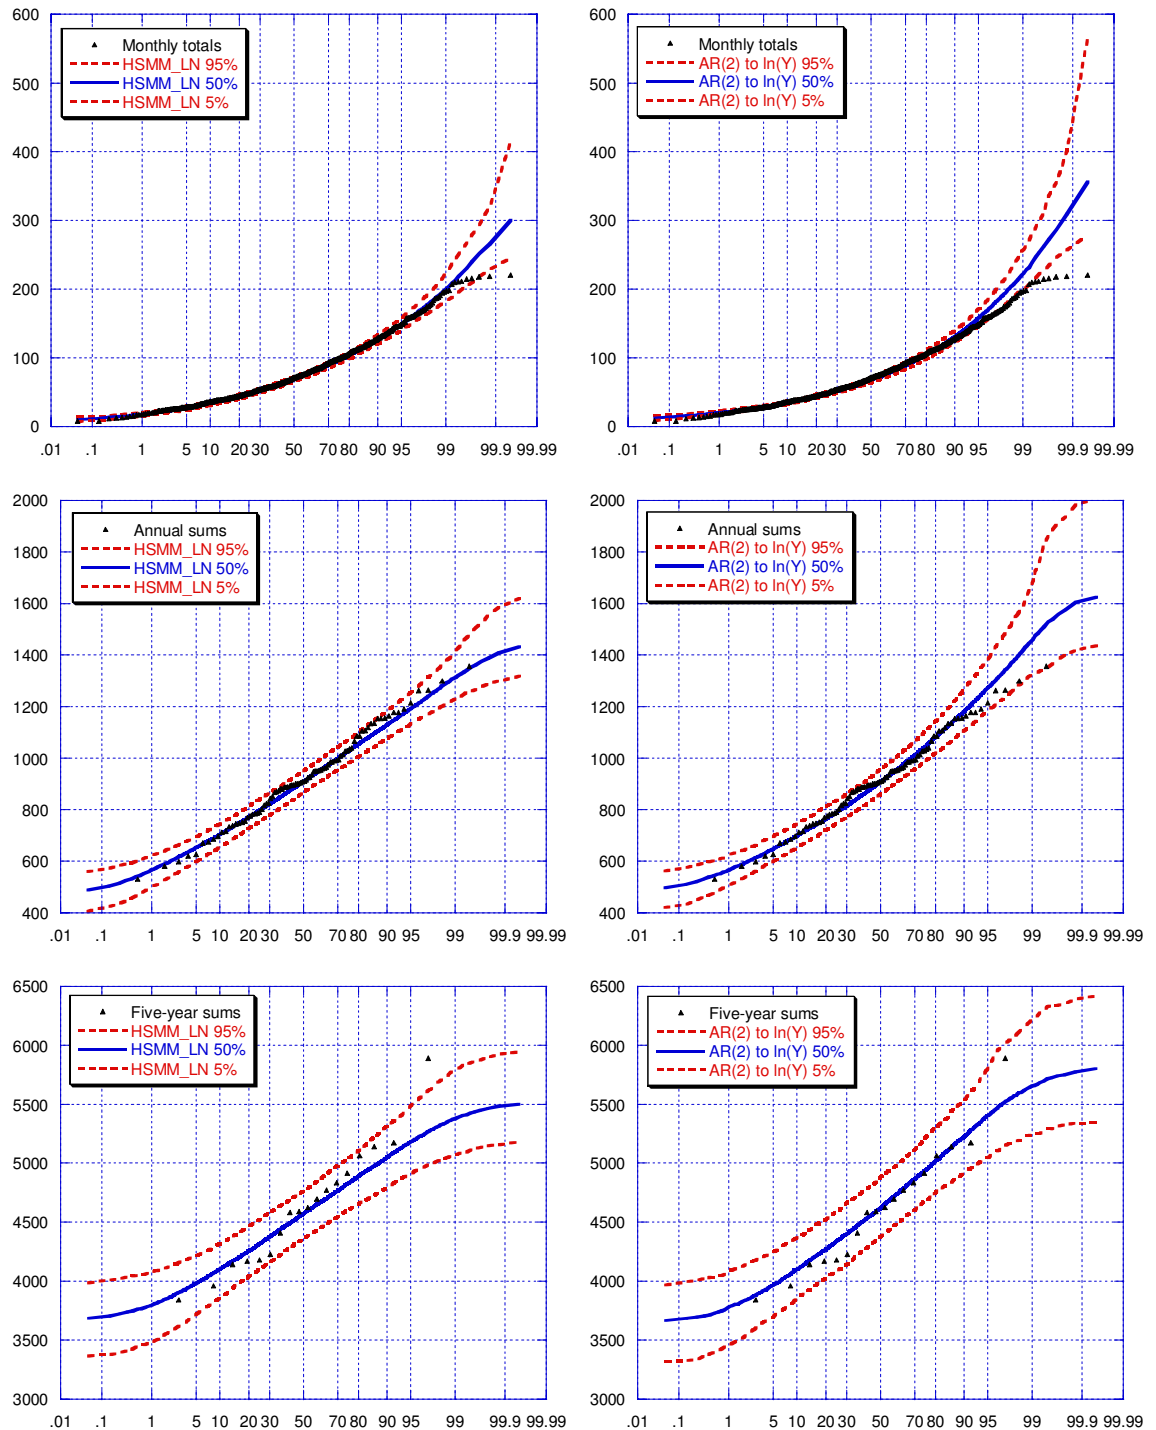


Figure 11.3 Gaussian probability plots showing deseasonalised monthly rainfall for District 71, together with annual and five-year aggregates of this series, alongside median values and intervals that contain 90% of values from 1000 simulations using a two-state lognormal HSMN (left column) and an AR(2) model with lognormal residuals (on right)

The accuracy of simulations using the lognormal HSMM and AR(2) models are further analysed in Table 11.2. These results show the capability of each model to replicate persistence in the monthly time series, with observed values within 90% envelopes. Furthermore for the lag-1 autocorrelation, LORT probability and run skew, the median values from these multiple simulations are very similar when comparing the two models. Simulations from the AR(2) model show a tendency towards longer runs than the HSMM simulations, with the longest run from 90% of these simulations exceeding the longest run in the monthly samples.

Table 11.2 Run statistics for the time series of deseasonalised monthly totals for District 71, together with the results of multiple simulations of this series using a two-state lognormal HSMM and an AR(2) model with lognormal residuals, showing medians and 90% confidence intervals

	Lag-1 autocorrelation	LORT probability	Maximum run length	Run skew
Monthly totals	0.218	1.135×10^{-32}	16	25.78
HSMM_LN (2s)	0.200 (0.130, 0.265)	2.551×10^{-16} (9.6×10^{-240} , 5.0×10^{-4})	13 (10, 19)	23.10 (17.07, 33.04)
AR(2)_LN	0.205 (0.132, 0.277)	3.937×10^{-16} (2.1×10^{-124} , 2.7×10^{-4})	23 (17, 30)	23.07 (17.91, 30.91)

The results in this section demonstrate that a two-state lognormal HSMM can produce accurate simulations of the deseasonalised monthly rainfall series from District 71. This model was chosen through Bayesian selection methods as the most appropriate description of these data. After incorporating parameter uncertainty following the calibration of this model, multiple simulations of identical length to the original series contain marginal values within 90% confidence intervals. Importantly the persistence of the monthly series is preserved.

11.3 Simulating monthly streamflows for the Burdekin River

The second time series that is simulated in this section is the series of monthly streamflows for the Burdekin River, Queensland. This series demonstrates significant two-state and three-state persistence and it is therefore a suitable choice for simulation with HMMs. The deseasonalised monthly rainfall series from District 71 was shown in the previous section to approximate a series of random draws from a lognormal distribution. In contrast to this, the Burdekin flow data shows a highly skewed distribution, even after removing seasonal nonstationarity, as demonstrated in Figure 11.4. The long tail of this marginal distribution contributes to a difficulty in transforming the data to approximate a series of Gaussian variates.

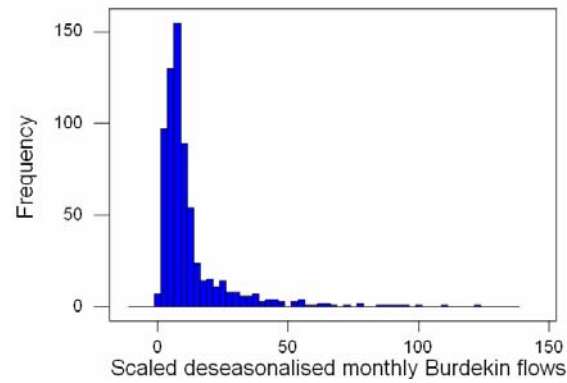


Figure 11.4 Histogram of scaled deseasonalised monthly Burdekin flows

The scaled deseasonalised monthly flow series is shown in Figure 11.5 on a lognormal probability plot, and it is clear that this distribution deviates significantly from a straight line. An Anderson-Darling goodness-of-fit statistic of 6.58 demonstrates that this series is inconsistent with a series of lognormal variates.

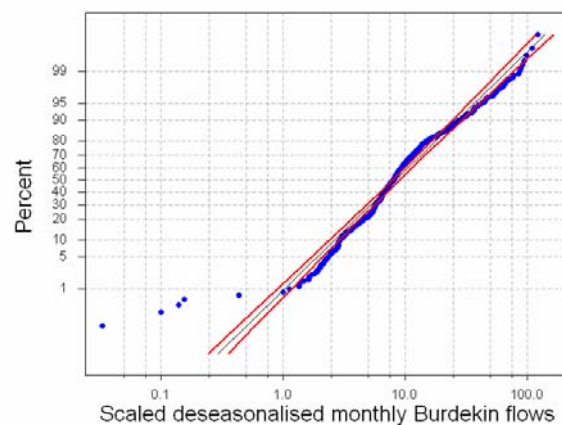


Figure 11.5 Lognormal probability plot showing scaled deseasonalised monthly Burdekin flows

With the series of natural logarithms failing to show consistency with a series of random draws from a Gaussian distribution, it is possible that the assumptions of ARMA models may then be compromised when calibrated to this series. Furthermore, the failure of the marginal distribution to approximate a known parametric form provides a useful opportunity for showing how the unbiased NP HMMs can produce accurate simulations.

In Section 8.4.3, NP HMMs identified significant three-state persistence in the time series of deseasonalised monthly flows for the Burdekin, with the three-state NP HMM also shown as a superior model to the two-state NP HMM through Bayesian model selection. Consequently, the former model is adopted in this example for the simulation of the scaled deseasonalised monthly flows. As a comparison to this model, an AR(3) model calibrated to the natural logarithms of the scaled variates is also used to obtain multiple simulations. A third order AR model is chosen

through an inspection of both autocorrelations and partial autocorrelations, in a manner similar to that used in the previous example.

Following calibration of these two models, 1000 simulations of identical length to the original series are produced using parameter values that are chosen from posterior distributions. After aggregating these monthly simulations over five different periods, the median values and 90% confidence intervals for various statistics are shown in Table 11.3. These results demonstrate that the mean, standard deviation and skewness of the deseasonalised monthly series, at each scale of aggregation, is simulated accurately by both stochastic models, with each statistic within 90% confidence intervals. Simulations from the three-state NP HMM have wider confidence intervals around the mean values of each aggregation than the AR(3) simulations, although the latter model produces wider bounds for the skew of the monthly series. This is a likely artefact of inaccurate assumptions concerning the AR(3) model fit. Both stochastic models are suitable for simulating the first three moments of the monthly flow series.

Table 11.3 Statistics for the time series of deseasonalised monthly flows from the Burdekin River over various aggregations calculated with one-month moving windows, along with the median and 90% confidence intervals of statistics from 1000 simulations using a three-state NP HMM and an AR(3) model with lognormal residuals

		Mean	Standard deviation	Skew
1m	Sample	12.72	15.12	3.378
	NPHMM	12.23 (9.9, 15.6)	14.79 (11.2, 18.1)	3.175 (2.44, 4.09)
	AR(3)	12.35 (10.2, 14.7)	13.71 (10.3, 19.9)	3.562 (2.43, 7.38)
6m	Sample	76.60	58.07	1.985
	NPHMM	73.23 (58.7, 93.0)	53.69 (39.7, 67.9)	1.647 (1.06, 2.35)
	AR(3)	74.04 (61.3, 88.3)	49.78 (36.1, 73.4)	1.840 (1.11, 3.68)
1yr	Sample	153.96	95.11	1.371
	NPHMM	146.32 (117.1, 186.1)	85.65 (62.6, 111.6)	1.279 (0.75, 1.99)
	AR(3)	148.26 (122.3, 176.3)	77.61 (54.3, 117.9)	1.352 (0.71, 2.84)
2yr	Sample	310.72	150.45	0.788
	NPHMM	291.93 (234.8, 371.9)	129.71 (89.8, 177.8)	0.885 (0.34, 1.62)
	AR(3)	296.72 (244.5, 352.7)	115.48 (77.8, 177.0)	0.953 (0.30, 2.18)
5yr	Sample	775.36	249.64	0.181
	NPHMM	730.46 (582.8, 932.1)	203.32 (133.9, 313.5)	0.439 (-0.19, 1.28)
	AR(3)	746.02 (613.8, 887.8)	180.41 (106.9, 289.0)	0.470 (-0.17, 1.38)

Simulations of the Burdekin flow data from both the three-state NP HMM and AR(3) models are compared further with the Gaussian probability plots shown in Figure 11.6. In a similar fashion to Figure 11.3, these probability plots show the distributions of the scaled deseasonalised monthly flows with median and 90% confidence intervals from simulations at aggregation scales of one year and five years. The plots indicate that the median values for each ranked value in the marginal distribution from the 1000 NP HMM simulations closely match original values, whereas the inaccurate simulation of monthly skew by the AR(3) model is

clearly demonstrated by the wide confidence interval around larger values in the marginal. Distributions at each aggregation are contained within 90% confidence intervals for both models, indicating that these two approaches are appropriate for simulating this flow series.

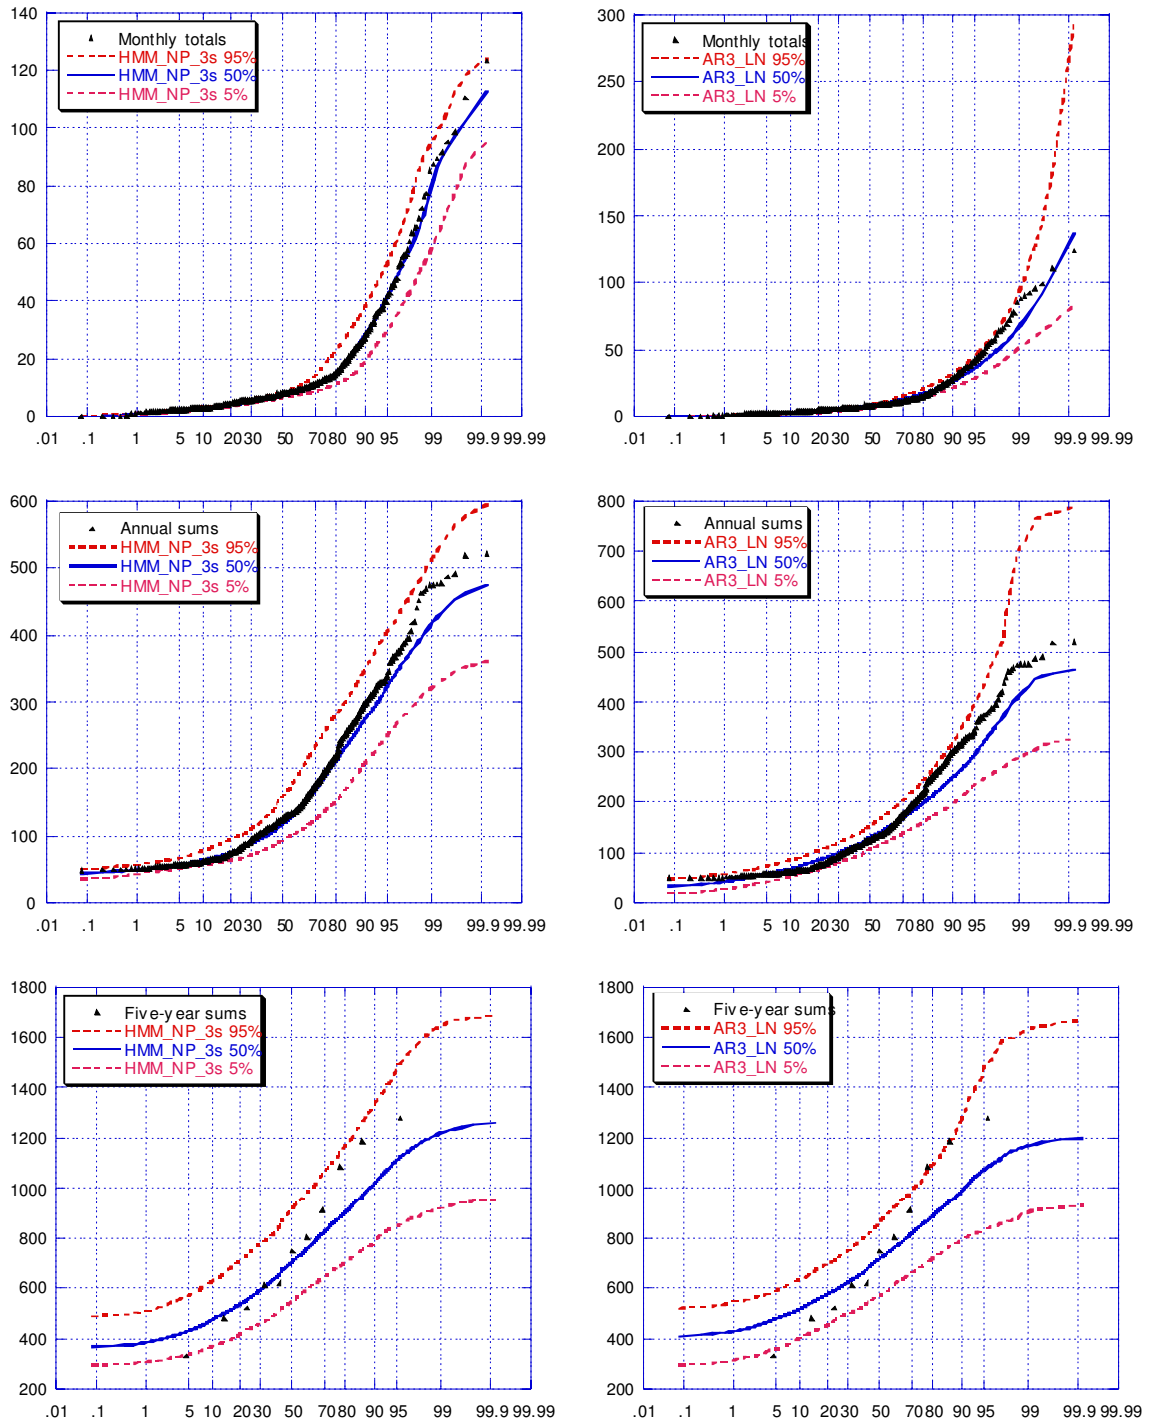


Figure 11.6 Gaussian probability plots showing deseasonalised monthly flows for the Burdekin River, together with annual and five-year aggregates of this series, alongside median values and intervals that contains 90% of values from 1000 simulations using a three-state NP HMM (left column) and an AR(3) model with lognormal residuals (on right)

The simulation of various runs statistics by these two models are summarised in Table 11.4, although this differs from Table 11.2 in the previous example through its exclusion of simulation results for the LORT probability due to this statistic having low values that are beyond computational bounds. These results show that the NP HMM underestimates the lag-1 autocorrelation in the deseasonalised monthly flow record. Although the observed lag-1 autocorrelation is captured within 90% confidence intervals of the AR(3) simulations, this model is unable to simulate accurately the maximum run length within monthly simulations having length of 684 values. The run skew statistic is simulated accurately by both models. These statistics display the suitability of the three-state NP HMM for simulating a persistent monthly hydrologic series.

Table 11.4 Run statistics for the time series of deseasonalised monthly flows from the Burdekin River, together with the results of multiple simulations of this series using a two-state lognormal HSMM and an AR(3) model with lognormal residuals, showing medians and 90% confidence intervals

	Lag-1 autocorrelation	Maximum run length	Run skew
Monthly totals	0.455	21	63.29
NP HMM (3s)	0.343 (0.253, 0.439)	20 (14, 31)	56.20 (33.34, 111.58)
AR(3)_LN	0.357 (0.229, 0.474)	51 (31, 94)	51.65 (31.8, 94.7)

11.4 Simulating catchment-scale rainfall for the Warragamba Reservoir

The third simulation example in this chapter focuses upon a deseasonalised monthly rainfall series that represents rainfall across the Warragamba Reservoir catchment in New South Wales, Australia. The Warragamba Reservoir is located to the east of Sydney, and is the major water storage in the Hawkesbury-Nepean Basin. Its catchment is the main source of water supply to Sydney, supplying up to 80% of its annual demand. Shown in Figure 11.7 alongside four adjacent smaller catchments, the Warragamba catchment covers an area of approximately 9050 km². As previously mentioned, simulation plays an important role in the risk analysis of reservoirs. One method to simulate inflows to a reservoir is to first simulate catchment-scale rainfall, and using rainfall-runoff transformations to capture the basic physical processes of catchment hydrology, to then generate long-term catchment runoff data. In this example, the effectiveness of both HMMs and ARMA models to provide accurate simulations of catchment-scale rainfall are considered together with the influence of monthly persistence upon risks to reservoir supply of such simulations.

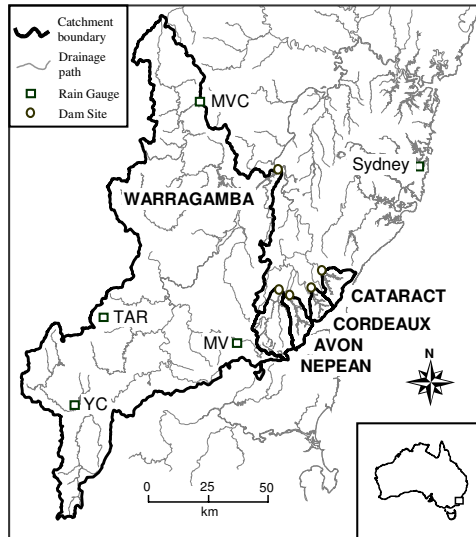


Figure 11.7 Locations of the Warragamba Reservoir catchment and adjacent smaller catchments, together with rain gauges that provide long-term rainfall data (after Thyer and Kuczera, 1999b)

Monthly rainfall totals are available from over 30 gauges within the Warragamba catchment, although these records vary in both length and quality. In order to represent the spatial distribution of rainfall across this catchment, monthly data are used from four individual gauges from which continuous records over a common 111-year period (1883-1993, 1332 months) are available. The quality of these data was verified, along with some in-filling of missing values, by Thyer and Kuczera (1999b). The locations of these four gauges: Moss Vale, Taralga and the Yarra-Goulburn composite and Mt. Victoria composite series, are shown in Figure 11.7 by their respective abbreviations MV, TAR, YC and MVC. Some statistics of these series are presented in Table 11.5, together with their Bureau of Meteorology (BOM) identification number. These statistics demonstrate the inherent variability in rainfall across the large catchment area. In order to analyse monthly persistence within rainfall across this area, a catchment composite rainfall series is constructed from these individual gauge records using the Thiessen polygon method. The ratio of the entire catchment area for which each of the four records are representative is also shown. The catchment composite rainfall series (y_t) is constructed from the four individual records in the following manner

$$y_t = \sum_{i=1}^4 f_i R_{t,i} \quad (11.1)$$

where $R_{t,i}$ is the observed rainfall from gauge i in month t and f_i is the ratio of the catchment area associated with gauge i to the whole catchment area. Statistics for the composite rainfall series are also shown in Table 11.5.

Table 11.5 Statistics for monthly rainfall from the four rain gauges in the Warragamba catchment, together with ratios of catchment area represented by each rain gauge, and the catchment composite rainfall series

Rain gauge	BOM number	Mean (mm)	Standard Deviation (mm)	Skew	f_i
Mt. Victoria	63056	89.24	78.97	1.961	0.3026
Moss Vale	68045	82.82	75.58	2.178	0.2836
Taralga	70080	68.02	51.72	1.853	0.2136
Yarra-Goulburn	70088	55.86	42.36	1.503	0.2002
Catchment composite		76.20	58.42	1.739	1

Prior to analysis, the seasonal nonstationarity in the composite series is removed in a similar manner as described previously, producing a series having zero mean and unit variance. After scaling these deseasonalised monthly variates to produce a series having mean (94.16) and standard deviation (58.06) equal to those of observed January values, the resulting marginal is shown on a lognormal probability plot in Figure 11.8. This figure indicates that the scaled deseasonalised composite record is consistent with random draws from a lognormal distribution, supported by an Anderson-Darling goodness-of-fit statistic of 0.84.

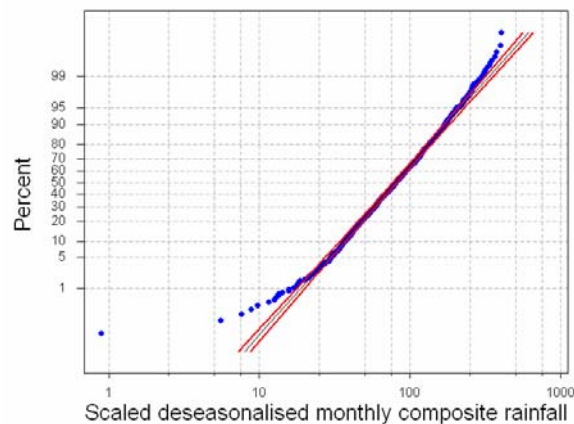


Figure 11.8 Lognormal probability plot showing scaled deseasonalised monthly values for the Warragamba composite rainfall series

The calibration of NP HMMs to the scaled deseasonalised composite series identifies significant two-state and three-state persistence; the sum of transition probabilities from the former has a 95th percentile of 0.608, with 90% of the posterior distribution for the sum of self-transition probabilities for the latter lying within the interval (1.307, 2.647). Furthermore, estimates of conditional distributions from the two-state NP HMM calibration closely approximate lognormal variates (AD statistics 1.79 and 2.20 for the wet and dry state distributions respectively). Therefore lognormal HMMs are likely to describe these data accurately. From the calibration of lognormal HMMs to this series, Bayesian model selection identifies that a three-state lognormal HMM is superior to a two-state lognormal HMM with $\ln BF_{3LN,2LN} = 3.0$.

Consequently, a three-state HMM is implemented to produce simulations of the scaled deseasonalised composite series. As a comparison to this model, an AR(3) model calibrated to the series of natural logarithms for this series is also used to generate simulations. Model selection results indicate that a three-state lognormal HMM is a superior for this series when compared to an AR(3) model with lognormal residuals, with $\ln BF_{3LN,AR(3)} = 12.4$.

The results from generating 1000 simulations of the Warragamba series are shown in Table 11.6. The median values for statistics from the HMM simulations closely match observed values, at each aggregation level, with each statistic shown contained within 90% confidence intervals. Conversely, simulations from the AR(3) model tend to over-estimate the variability in the monthly series, failing to simulate accurately the standard deviation or skewness at this scale. This inaccuracy is further demonstrated in the probability plots of Figure 11.9. At monthly, annual and five-yearly scales, the median values from HMM simulations tend to more closely match observed distributions than AR(3) simulations. By closely reproducing the general shape of the marginal at a range of aggregations, including the extreme values, the three-state lognormal HMM is an appropriate model for simulations of the composite series.

Table 11.6 Statistics for the time series of deseasonalised monthly totals from the Warragamba composite series over various aggregations calculated with one-month moving windows, along with the median and 90% confidence intervals of statistics from 1000 simulations using a three-state lognormal HMM and an AR(3) model with lognormal residuals

		Mean	Standard deviation	Skew
1m	Sample	94.16	57.83	1.526
	HMM	93.74 (88.3, 98.5)	57.84 (52.9, 63.4)	1.771 (1.40, 2.55)
	AR(3)	95.30 (90.3, 100.1)	65.05 (58.2, 74.2)	2.217 (1.75, 3.59)
6m	Sample	564.69	164.18	0.600
	HMM	562.59 (529.4, 591.2)	160.34 (144.5, 179.1)	0.664 (0.40, 1.03)
	AR(3)	571.06 (542.3, 601.0)	187.33 (165.1, 217.0)	0.987 (0.62, 1.66)
1yr	Sample	1129.78	241.97	0.437
	HMM	1125.05 (1058.7, 1183.2)	235.62 (208.3, 269.7)	0.447 (0.17, 0.81)
	AR(3)	1142.95 (1084.9, 1202.7)	274.36 (236.3, 324.9)	0.712 (0.35, 1.28)
2yr	Sample	2262.32	380.40	0.384
	HMM	2249.27 (2116.6, 2365.5)	338.23 (288.1, 402.0)	0.293 (-0.06, 0.72)
	AR(3)	2285.13 (2168.4, 2406.6)	391.44 (326.8, 473.7)	0.494 (0.11, 1.05)
5yr	Sample	5662.12	695.22	0.406
	HMM	5624.90 (5295.0, 5911.4)	528.77 (414.3, 666.1)	0.115 (-0.37, 0.67)
	AR(3)	5712.66 (5413.1, 6019.8)	614.39 (476.9, 791.2)	0.243 (-0.22, 0.92)

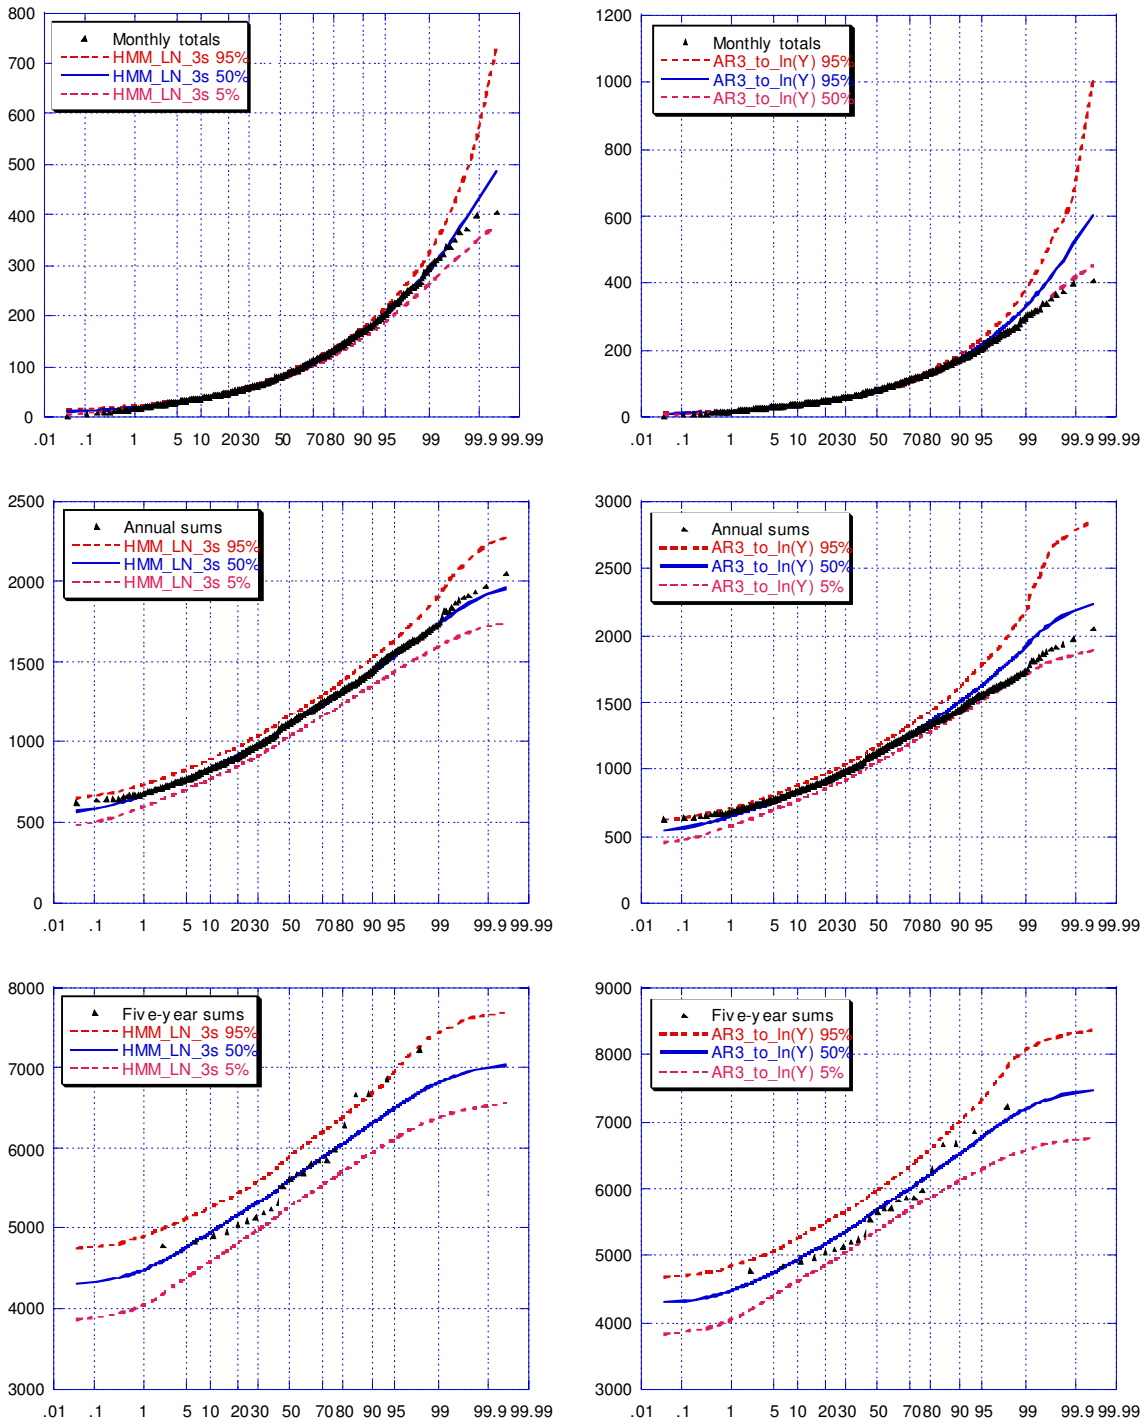


Figure 11.9 Gaussian probability plots showing deseasonalised monthly rainfall for the Warragamba composite rainfall series, together with annual and five-year aggregates of this series, alongside median values and intervals that contains 90% of values from 1000 simulations using a three-state lognormal HMM (left column) and an AR(3) model with lognormal residuals (on right)

In order to demonstrate the effectiveness of the three-state lognormal HMM to replicate the persistence of the monthly composite series, the simulation results for various runs statistics are summarised in Table 11.7. These result indicate that the HMM and AR(3) models replicate characteristics of persistence in a similar manner, although the latter simulates the lag-1

autocorrelation of the composite series more accurately. This outcome is somewhat expected as ARMA models are designed to replicate autocorrelation explicitly. Overall the three-state lognormal HMM is shown here to be a suitable model for the simulation of the monthly composite rainfall series for Warragamba. This further demonstrates that models designed to replicate explicitly persistence in monthly hydrologic records can improve simulations from alternative linear time series models.

Table 11.7 Run statistics for the time series of deseasonalised monthly totals from the Warragamba composite series, together with the results of multiple simulations of this series using a three-state lognormal HMM and an AR(3) model with lognormal residuals, showing medians and 90% confidence intervals

	Lag-1 autocorrelation	LORT probability	Maximum run length	Run skew
Monthly totals	0.113	1.689x10⁻¹⁰	14	20.01
HMM_LN (3s)	0.094 (0.033, 0.154)	1.450x10⁻⁸ (7.5x10 ⁻⁸⁹ , 6.7x10 ⁻²)	19 (14, 25)	19.14 (15.00, 25.86)
AR(3)_LN	0.128 (0.054, 0.199)	5.718x10⁻⁹ (8.4x10 ⁻⁵³ , 2.6x10 ⁻²)	19 (15, 25)	19.47 (15.24, 25.32)

Simulations from the two stochastic models presented in this section are further compared by analysing their influence upon reservoir storage volumes. In order to conduct this comparison, a relationship between rainfall across the catchment and runoff must be used. Using observed values of catchment runoff for a period of 31 years, the relationship with annual values of the catchment composite rainfall series is shown in Figure 11.10. A piecewise linear regression, with breakpoints at annual rainfall values of 500mm and 982mm, is used to relate annual rainfall to annual runoff totals.

Let $\{r_t, t = 1, 2, \dots, T\}$ be the time series of annual runoff corresponding to the time series of annual rainfall $\{y_t, t = 1, 2, \dots, T\}$. The piecewise linear regression line is described by zero r_t values for $y_t < 500$, together with

$$r_t = 0.18y_t - 90, (500 < y_t < 982) \quad (11.2)$$

and also

$$r_t = 0.75y_t - 650, (y_t > 982) \quad (11.3)$$

For the 33 pairs of rainfall and runoff data shown on this plot, this regression explains 83% of the variability in the rainfall data, which is an improvement over a simple linear regression that explains 77% of variability.




Figure 11.10 Relationship between the annual composite rainfall data for Warragamba and annual runoff (blue points) together with a piecewise linear regression over the period (1961-1993)

The piecewise linear regression is a suitable description of the relationship between total annual rainfall in the Warragamba catchment and streamflow responses. For this simulation investigation, the annual rainfall-runoff relationship is adapted to produce a method to calculate runoff totals corresponding to simulated monthly rainfall.

From the sequences of 1332 simulated monthly values, series of 111 simulated annual totals are obtained by aggregating values 1-12, 13-24, etc. As seasonal variability within the composite monthly series was removed prior to simulation, monthly simulations are scaled, prior to aggregation, to produce monthly means and standard deviations equal to the observed values in each calendar month. These annual totals are then used to obtain annual runoff values, using the relationships outlined previously. Each annual runoff total is then disaggregated to a series of 12 monthly runoff values, using the proportions of the total annual rainfall that are simulated in each month. Simulated monthly runoff values from the catchment (measured in millimetres) are then transformed into reservoir inflow volumes after multiplying by the approximate catchment area (9056 km²). One result of this simplified method is that if simulated annual rainfall totals are below 500mm, the annual runoff will be zero, as will the runoff corresponding to each month within that specific annual period. This technique is adopted for monthly simulations from both a three-state lognormal HMM and an AR(3) model calibrated to the time series of natural logarithms of the deseasonalised monthly composite rainfall.

The time series of monthly reservoir inflows obtained with simulations from the two stochastic models are adopted in order to perform simplified reservoir water balance investigations, as a means to compare the accuracy of these simulation approaches. The water balance is conducted on a time series of reservoir storage volumes for each month (t), calculated as

$$Storage_t = Storage_{t-1} + Inflow_t - Demand_t \quad (11.4)$$

The capacity storage for the Warragamba reservoir is assumed to be 1,886 GL, above which the reservoir spills. The main function of this analysis is not to incorporate complex rainfall-runoff relationships that define reservoir inflows, rather to compare the accuracy of monthly rainfall simulations from two models, so a number of generalisations are included. It is assumed that average annual demand from the Warragamba remains constant at 750 GL across the simulation period of 111 years. The monthly demand from the system is one-twelfth of this annual demand (therefore 62,250 ML), except in years of low storage volumes. In these periods monthly demand is restricted, such that monthly demand levels are multiplied by the restriction factors defined in Table 11.8.

Table 11.8 Restriction factors for monthly demand volumes from Warragamba reservoir corresponding to simulation storage volumes

Monthly storage (% of capacity)	Monthly storage (GL)	Demand restriction level	Demand restriction factor
55% – 100%	1037 – 1886	0	1.00
45% – 55%	849 – 1037	1	0.93
40% – 45%	754 – 849	2	0.88
35% – 40%	660 – 754	3	0.80
25% – 35%	472 – 660	4	0.70
5% – 25%	94 – 472	5	0.50
0% – 5%	0 – 94	6	0.50

The time series of monthly storage volumes that correspond to each simulated series are then analysed, with various statistics calculated. Firstly the tendency for demand restriction for each simulation is estimated as the percentage of months in which storage volumes fall below 50% of capacity. Secondly the reliability of monthly simulations is calculated as the percentage of months in which restriction levels are greater than 0 and finally a measure of the security of reservoir storages is obtained as the percentage of months which the restriction levels remain below 6. A restriction level of 6 corresponds to a storage volume below 5% of the total capacity. By calculating these statistics for each of the 1000 simulations, average values for the three-state lognormal HMM are then compared to average values from the AR(3) model. These results are presented in Table 11.9.

Table 11.9 Average water balance statistics from 1000 simulations of monthly composite rainfall series for Warragamba

	3-state lognormal HMM	AR(3) with lognormal residuals
Reservoir restriction	24.0%	25.9%
Reservoir reliability	98.7%	97.4%
Reservoir security	74.4%	54.8%

These results demonstrate that a three-state lognormal HMM produces monthly simulations of the catchment composite rainfall that have lower drought risk on the supplies in Warragamba reservoir than simulations from an AR(3) model. This is consistent with previous results that showed the HMM to provide improved simulations of the monthly rainfall across the Warragamba catchment. The ability of the HMM to describe the persistence within the monthly rainfall of the Warragamba catchment leads to improved descriptions of reservoir hydrology. This is an important result in terms of possible applications of HMMs in water resources management.

11.5 Summary of chapter

This chapter has demonstrated that HMMs provide a useful method to simulate persistent hydrologic data. Previous studies have not compared the usefulness of these models with the widely-used ARMA family of models as tools to simulate monthly or annual rainfall data. This chapter showed the simulation of three monthly time series, each of which are indicative of the possible applications of such an approach. From the range of models in the HMM family that have been introduced throughout this thesis, a different model formulation was shown to be the most appropriate for each of these series.

For each these simulation examples, the accurate simulation of statistics and also of the shape of marginal distributions, each at a range of temporal aggregations, were compared for a range of stochastic models. In each example, the HMM model used was at least as accurate as the appropriate AR model. In the simulation of catchment-scale runoff for the Warragamba, the three-state lognormal HMM provided reservoir inflows of lower drought risk than provided from an AR(3) model. These results suggest that HMMs not only provide accurate descriptions of hydrological persistence, but also provide accurate simulations of persistent data. This latter point provides an important benefit for the application of these models in stochastic hydrology.