A Stochastic Space-Time Rainfall Model for Engineering Risk Assessment

by Michael Leonard

Submitted in fulfilment of the requirements for the degree of
DOCTOR OF PHILOSOPHY

April 14, 2010

FACULTY OF ENGINEERING, COMPUTER AND MATHEMATICAL SCIENCES

School of Civil Environmental and Mining Engineering

THE UNIVERSITY OF ADELAIDE AUSTRALIA

Chapter 8

Climatic and Seasonally Partitioned Extreme Rainfall

At daily and sub-daily time scales heterogeneity can be due to different types of storms such as convective and frontal rainfall. However this difference is not always modelled explicitly in stochastic models of rainfall. For example, the DRIP model of *Heneker et al.* [2001] uses only one type of storm, but the storm parameters are sufficiently flexible to capture the broad range of rainfall observations. In contrast to this, there are numerous stochastic rainfall models having multiple types of storms [*Cowpertwait*, 1995; *Sansom*, 1999; *Pegram and Clothier*, 2001a]. The parameters for each storm type are more defined, for example, one set of parameters might generate short duration high intensity storms, while another set might generate longer duration low intensity storms. To reduce the overall number of parameters many models allow for some parameters to be common to both types of storm, for example the frequency, the internal variability, the direction or the areal coverage.

At an intra-annual scale, temporal heterogeneity appears as a seasonal process. The months of the year have widely varying weather regimes which give rise to different storm properties. Seasonality is a ubiquitous aspect of rainfall models and there are two common techniques for modelling it: (i) to use a new parameter set each month of the year and (ii) to use harmonic series fitted to monthly parameters. Harmonic series have two main advantages in that they reduce the overall number of parameters and they provide more meaningfully smooth changes in parameter values. However, because of the smoothness constraint of harmonic curves, they are less flexible than independent monthly parameter sets. *Srikanthan and McMahon* [2001] discusses seasonality and parameterisation methods for a wide range of models. For the Neymann-Scott model, *Cowpertwait* [2004] provides a detailed comparison. Irrespective of these two approaches, few models address

the correlation of monthly rainfall in time. This point is often overlooked since monthly rainfall correlation is often low (say ~ 0.1). The issue of correlated monthly rainfall is nonetheless explored in this chapter since it is significant to understanding the variance of annual rainfall totals.

Inter-annual variability is increasingly acknowledged as an important aspect of rainfall. For example, numerous authors attribute inter-annual variability as a reason for under-estimation of variability in annual rainfall totals [*Wilks and Wilby*, 1999; *Srikanthan and McMahon*, 2001; *Frost*, 2004]. These studies demonstrate links between rainfall totals and climate indices. The links imply that rainfall is auto-correlated as a result of larger scale climatic events that persist over time. The variability and correlation of annual rainfall is important to water resource planners since, for example, the short-term risk of a drought may differ significantly from the long-term risk. Models for water resources planning typically require monthly or annual aggregates, whereas models at daily and sub-daily scale have been developed with different applications in mind, such as flash flooding or urban drainage. In these models, monthly and annual processes are the aggregate result of the many daily and sub-daily events and for this reason the correlation at the annual scale is not preserved well. To remedy this, climate states are introduced into the model whereby each state has different parameters tailored either to wetter or drier years.

8.1 INTRODUCTION

An important role of hydrologic design has been to quantify and mitigate the risks of flooding that arise from the variability of extreme rainfall and streamflow values. Seasonal fluctuations are a significant source of variability in both rainfall and runoff records across many regions. However, seasonality is often overlooked when evaluating flood risk due to the use of annual maxima for defining extreme values. The phrase "1 in 100 year" flood does not inform whether a given extreme value is more likely to come from one season over another. The oversight of seasonality is also common to the peak-over-threshold method, even though this method is capable of obtaining more than one extreme value per year. One example where seasonal variation can be important to the analysis of extreme values is in the design or assessment of dual purpose basins or reservoirs to be used jointly for water harvesting for supply and flood attenuation. Specifically, a reduced flood risk within a certain season may be exploited to allow for additional storage of water. In this context, it is possible to consider the trade-off between the economic benefit of additional (seasonal) storage to the cost of increased flood-risk. The potential benefits of such a scheme are likely to be more significant for locations having marked seasonal

variability. As an example, consider a tropical location having two seasons termed 'wet' and 'dry' and that the 'annual' extremes, by exaggeration, occur only during the wet season. Consequently, during the wet season the structure functions as a flood mitigation structure and during the dry season (when water is a more highly valued commodity) it could be used for seasonal water storage without impinging on the flood security.

Several locations in Australia, having contrasting climates, are used to illustrate the characteristics of seasonal extreme rainfall values. The sites include Adelaide and Perth (Mediterranean climates), Sydney and Melbourne (temperate climates) and Cairns (topical climate). Figure 8.1 shows extreme rainfall values corresponding to the 5 locations about Australia. Each figure shows distributions of maximum rainfall intensities occurring in a 72 hour time window, plotted against log-normal probabilities. Similar results are obtained for other time aggregates. The three series in each plot correspond to annual maximum values from the December to November year (solid circles) and the two calendar seasons of summer (Dec-Feb; hollow circles) and winter (Jun-Aug; crosses) respectively. Log-normal distributions were fitted to the seasonal maxima and confidence intervals were evaluated at the 90% level (summer - solid lines, winter - dashed lines) using FLIKE software [*Kuczera*, 1999]. The emphasis of Figure 8.1 is as a motivating example, providing comparisons between seasonal distributions of summer and winter for varying climatic regions. Similar observations can be made for the autumn and spring seasons (not shown, but for the cases presented lie predominatly between the summer and winter distributions). A comparison of the seasonal distributions to the distribution of annual maxima is also made in Figure 8.1, but is restricted to a qualitative discussion due to the omission of spring and autumn seasons. The methodology for formally comparing multiple seasons and seasons of various length (e.g. 6-month) is presented later in this chapter.

Figure 8.1 (a), Adelaide, presents a case where the median seasonal extreme is similar for the two seasons, yet the summer season is more variable than the winter. Thus, the less frequent summer extremes have a higher intensity. By way of contrast, Figure 8.1 (b), Cairns, presents a case where the seasonal extremes follow distinctly different distributions and where there is a large difference in the magnitude of intensity between the two seasons - attributed to the tropical climate in this region. Figure 8.1 (c), Melbourne, presents another case where the summer distribution is more variable than the winter distribution, with similar magnitudes occurring only in lower tails for the respective seasons. Figure 8.1 (d), Perth, also has summer as the more variable season, yet it differs from the other sites in that winter has the higher extreme values (excepting the

upper tail). Figure 8.1 (e), Sydney, provides a case where the extreme value distributions are indistinguishable with respect to season.

The annual maximum in a given year is the maximum value of the four observed seasonal maxima. Figure 8.1 (a) shows that the annual maximum in any given year can be contributed by either the summer or winter distributions (or autumn or spring, though not shown). For example, in some years the summer maximum will come from the upper tail of the summer extreme distribution and be larger than the extremes of the other seasons, but in other years it is possible that a winter maximum (from the upper tail of the winter extreme distribution) is larger than the extremes of the other seasons. Figure 8.1 (c) and Figure 8.1 (e) provide similar cases where the annual maximum is likely to come from any of the seasons, Figure 8.1 (b) provides an example where the annual maxima mostly come from the summer season, and Figure 8.1 (d) provides an example where the annual maxima are mostly attributed to the winter distribution. It is evident that seasonality contributes an important source of variability to the annual extreme values and that there is a wide variety of relationships between seasons across varying locations.

Whereas seasonal variation occurs on an intra-annual timescale, climatic variation in hydrologic variables has been demonstrated at inter-annual and inter-decadal scales. Records of sea-surface temperatures, pressure differences and other variables about the oceanic basins have been used to explore these underlying relationships. Examples include the Southern Oscillation Index (SOI) in the Pacific Ocean, the Dipole Mode Index (DMI) of the Indian Ocean and the Interdecadal Pacific Oscillation (IPO), the latter showing oscillations at a lower frequency than the SOI. Though not focusing on extreme values, numerous studies have investigated the role of inter-annual climatic variability and its links to seasonal variability present in Australian hydrologic records [*Chiew et al.*, 1998; *McBride and Nicholls*, 1983; *Simmonds and Hope*, 1997; *Verdon and Franks*, 2005]. With respect to extreme values, *Franks and Kuczera* [2002] demonstrate statistically significant differences in the distribution of extremes by investigating a step-change in the year 1945 (associated with a shift in the IPO) for 41 streamflow records across New South Wales. The results of this stratification suggest that annual maxima are not independent and identically distributed as is assumed by standard flood-frequency analyses. *Kiem et al.* [2003] extended this observation by investigating different climate states of the El Niño Southern Oscillation (ENSO), namely El Niño and La Niña, using a regional flood index based on 40 streamflows across for New South Wales. They demonstrate not only that flood-risk levels are strongly associated with ENSO but also that the IPO, acting at a longer timescale, modulates both the amplitude and frequency of flood events in the various
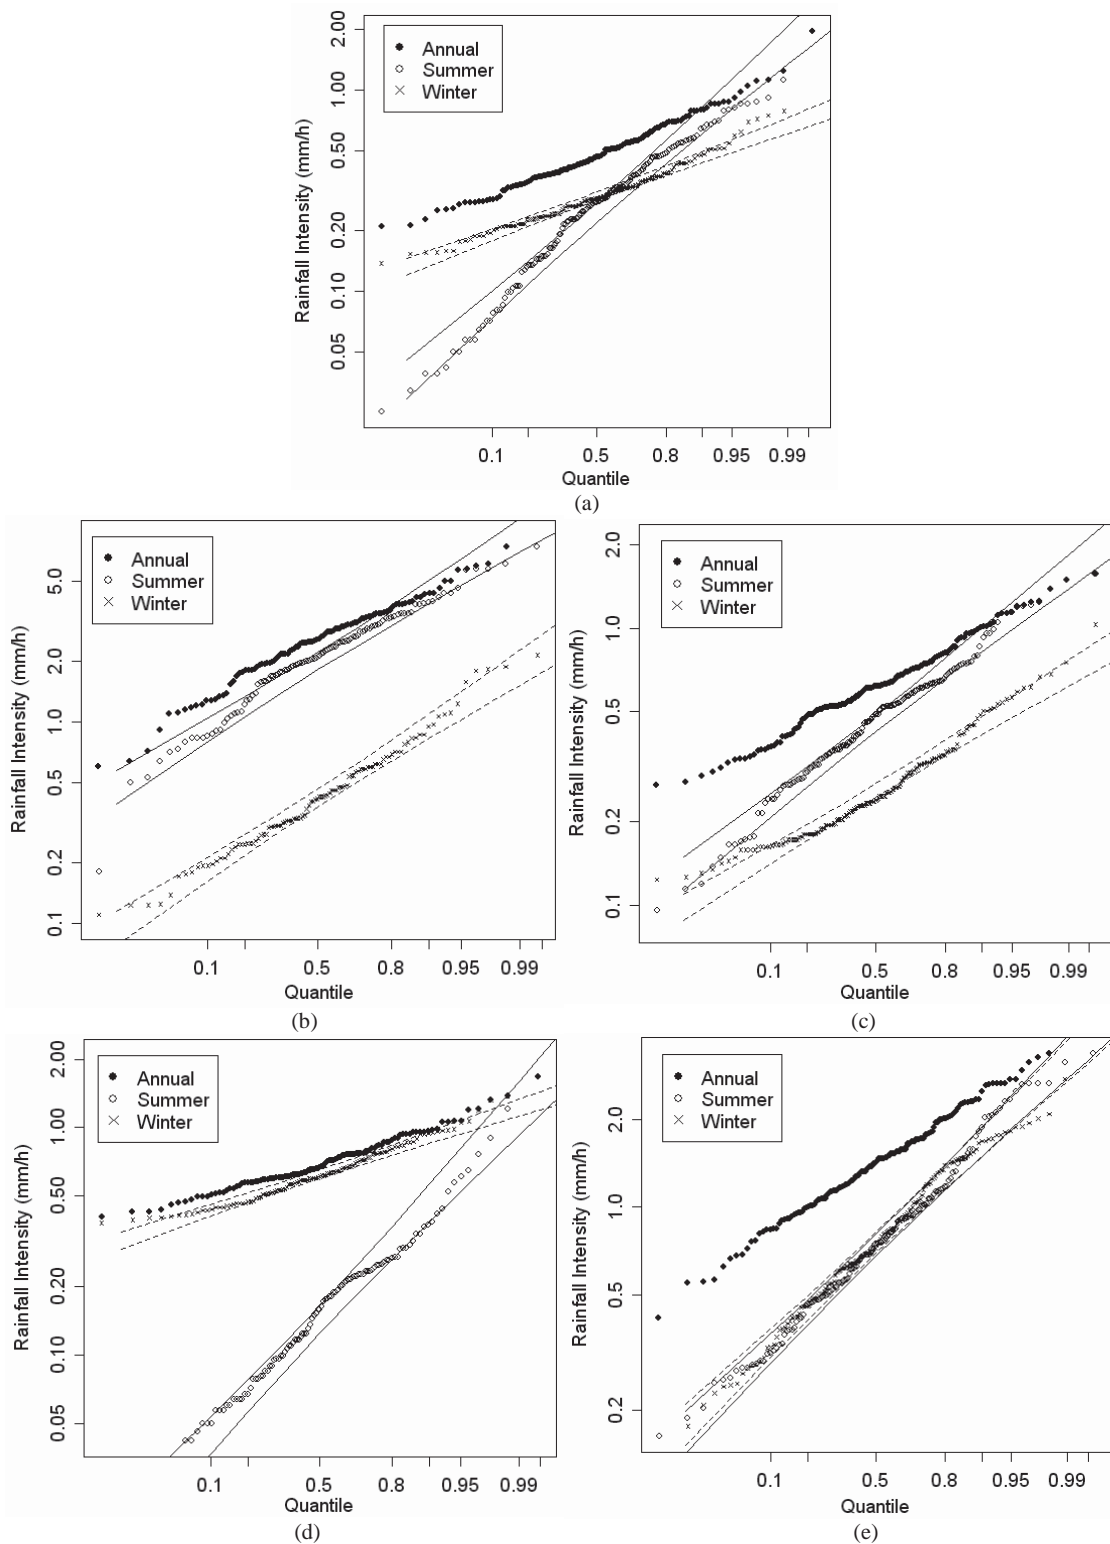
Figure 8.1 Comparison of distributions of summer and winter seasonal maxima of 72-hour rainfall for selected sites: (a) Adelaide, (b) Cairns, (c) Melbourne, (d) Perth, (e) Sydney. Lines represent 90% confidence interval from fitted log-normal distributions.

states. *Verdon and Franks* [2006] further emphasise the links between ENSO and IPO using reconstructed paleo-climate records over a 400 year period. *Micevski et al.* [2006] investigate the spatial relationships of streamflow maxima for a wide variety of catchments across Eastern Australia and also conclude that the IPO has a significant influence on flood-risk levels.

Building on these observations of inter-annual variability, this chapter considers the additional influence of seasonal extremes and their relationship to annual maxima. The emphasis of the chapter is not to repeat observations of climatic variability, but to present a generic framework for modelling the intra-annual and inter-annual partitions of a dataset. The methodology for this is outlined in Section 8.2 and it allows both seasonal partitions and climatic partitions to be considered in isolation should either process be non-identifiable for a given dataset. An application of the framework is provided in Section 8.3 a small, realistic design scenario involving Scott Creek, located in the Adelaide Hills, Australia. Section 8.4 demonstrates a method to partition the extremes of the SNSRP model and demonstrates this for a case study in the Bourke region.

8.2 METHODOLOGY

Let Y represent a hydrologic variable of interest such as rainfall or streamflow, where observed records are comprised of individual measurements, y_l^k, where $k = 1 \ldots K$ is a year index and $l = 1 \ldots L$ is an index for the time period within that year, e.g. daily increments. Consider then a range of climate states $c = 1 \ldots C$, where C is the number of climate states that form the basis of inter-annual partitions of a dataset. The climate states are defined to be mutually exclusive and are exhaustive of the data, so that any year, k, belongs to one of the states. For this reason, 'neutral' climate states should be included in the analysis. Also note that the year can be defined to follow intervals other than January to December. For a record having K years in total, the year k is defined as belonging to climate state c according to some function, $Climate(k)$. The years belonging to state c are then defined as, $N_c = \{k | \, Climate(k) = c, \, k = 1 \ldots K\}$.

Consider also a range of seasons $s = 1 \ldots S$, where S is the number of seasons that form the basis of intra-annual partitions of a dataset. The seasons are defined to be mutually exclusive and are exhaustive of data within a year. For this reason, the year must be a perfect multiple of the season length, for example, seasons of length 1, 2, 3, 4, 6 or 12-months satisfy this criterion. Having a generic season length allows for variation in extremes to be considered at different scales, so that 1 month seasons correspond to monthly variation, 3-month seasons can correspond to calendar seasons, 6-month seasons can give

'wet'/'dry' seasonal comparisons and 12-month seasons reproduce annual maxima. In order to relate seasonal maxima at the different scales it is important that the combined length of shorter seasons aligns with the interval of the longer season. For this reason, using Southern Hemisphere calendar seasons, summer (Dec. to Feb.), autumn (Mar. to May), winter (Jun. to Aug.) and spring (Sep. to Nov.) requires the annual 'season' to be defined for years on the interval December to November, rather than as a calendar year. The time index, l, for an element of data within a given year is defined as belonging to a season, s, according to some function, $Season(s)$. The time periods belonging to season s are then defined as, $N_s = \{l|\ Season(l) = s,\ l = 1 \ldots L\}$.

Following this definition, the hydrologic variable Y is partitioned into variables, Y_s^c, with respect to climate states and seasons according to,

$$Y = \bigcup_{c,s} Y_s^c \mid c = 1 \ldots C,\ s = 1 \ldots S. \tag{8.1}$$

The partition of observed seasonal data across all years in a given climate state, Y_s^c, is comprised of the set of elements

$$Y_s^c = \left\{ y_{s,l}^{c,k} \mid k \in N_c,\ l \in N_s \right\}, \tag{8.2}$$

where $y_{s,l}^{c,k}$ is an individual element. In a similar manner, the seasonal data in an individual year, k, corresponding to climate state c is defined as $Y_s^{c,k} = \{\ y_{s,l}^{c,k} \mid l \in N_s\}$.

For some aggregation of the data, h, let X_s^c represent the maximum values from each season s in each climate state c, X^c represent the annual maxima (across all seasons) in each climate state c and X represent the annual maxima across all seasons and states. Thus if there are 3 climate states and two 6-month seasons, there would be 6 different sets of climatic-seasonal maxima $[X_1^1,X_2^1,X_1^2,X_2^2,X_1^3,X_2^3]$, 3 sets of climatic annual maxima $[X^1,X^2,X^3]$ and the overall set of annual maxima X. The sets involving climatic partitions are shorter than others as only a portion of the years will correspond to a given state. For an individual year, k, in season s and state c, the data yield one maximum value, denoted $x_s^{c,k}$, according to some maximising function, $g_h()$, as

$$x_s^{c,k} = g_h(Y_s^c), \tag{8.3}$$

Note that different maxima may arise from varied aggregates of the observed data, for example, hourly data may give maxima corresponding to 1 hour, 6 hour, 12 hour, 24 hour and 72 hour periods. In the event of aggregated data, the maximising function, $g_h()$, may use a moving window of length, h, from which the maximum value is selected. A moving window is able to detect the maximum 'burst' of duration h, which is in contrast

to obtaining the maximum from data aggregated at duration h, due to the binning artefact of the aggregation process.

Denoting $x^{c,k}$ as the annual maximum value of year k, it is related to the seasonal maxima as the maximum of those values,

$$x^{c,k} = max\{x_1^{c,k}, x_2^{c,k}, \ldots, x_S^{c,k}\}. \tag{8.4}$$

The cumulative distribution function of the annual maxima, $F_X(X^c)$, is consequently related to the seasonal maxima as,

$$
\begin{aligned}
F_X(X^c) &= Pr(X^c < x) \\
&= Pr(X_1^c < x \cap \ldots \cap X_S^c < x).
\end{aligned}
\tag{8.5}
$$

If the seasonal maximum values in a given climate state are assumed independent from one another, the distribution of annual maxima can be related to the seasonal maximum distributions, as

$$
\begin{aligned}
F_X(X^c) &= Pr(X_1^c < x) \times \ldots \times Pr(X_S^c < x) \\
&= \prod_{i=1}^{S} F_X(X_i^c).
\end{aligned}
\tag{8.6}
$$

In the event that there is correlation between maximum values in any given pair of seasons then the assumption of independence is invalid. A convenient method for accommodating correlations between the seasonal marginal distributions, $F_X(X_i^c)$, is to use a copula function [*Zhang and Singh*, 2006], denoted as,

$$F_X(X^c) = C(F_X(X_1^c), F_X(X_2^c), \ldots, F_X(X_S^c)). \tag{8.7}$$

Copulas are a flexible representation of multivariate distributions, as each marginal distribution can have a different distributional form and there is a wide range of copulas to select from to be able to yield the correlated joint distribution.

As the set of observed annual maxima across all climatic states is the union of the annual maxima within each climate state, the overall distribution of maxima is equivalent to a mixture distribution of the annual maxima in each state. This is given as,

$$F_X(X) = \sum_{c=1}^{C} w_c F_X(X^c), \tag{8.8}$$

where the weight, w_c, is the ratio of the number of years in climate state c to the total number of years in the record. The distribution, $F_X(X)$, given in Eq. 8.8, represents the conventional concept of annual maxima, specifically, one maximum value obtained per each year in the record. Therefore, Eq. 8.6 (or Eq. 8.7 in the event of correlated maxima)

and Eq. 8.8 provide the probabilistic relationship between the conventional concept of annual maxima, $F_X(X)$, to maxima conditioned on being within a given season, climate state or both. This relationship provides the basis for considering trade-offs in seasonal or climatic levels of risk that give the same annual exceedance probability. This method can also be used to consider seasonal or climatic partitions separately, as it is possible to consider the data as belonging to only one climate state, or one (annual) season.

An alternative method for constructing distributions of extreme values, particularly in cases of short data records, is to use all of the peaks in the dataset that occur above a specified threshold value. The resulting distribution is not the same as the distribution of annual maxima given by Eq. 8.8, but a similar argument can be applied for determining the relationship it has to seasonal or climatic partitions of the data set. A key difference is that Eq. 8.4 is not valid with respect to the peak-over-threshold method, as, when a common threshold value is used for all seasons, the set of peaks within a given year is the union of the seasonal peaks. For this reason, when using the peak-over-threshold method, both seasonal and climatic partitions of the data are related as mixture distributions to the unpartitioned distribution of peaks. The peak-over-threshold method is less common with long records because of the problem of defining independence between maxima. Also, the use of seasonal and climatic partitions on short records is inappropriate as the partitioning reduces the data for fitting distributions to only those within that partition. For these reasons, the peak-over-threshold method is not considered further in this chapter, though similar results would be expected applying either method to long records of data.

8.3 URBAN DESIGN APPLICATION - SCOTT CREEK

The application of seasonal extremes to a hypothetical design situation typical to urban areas is illustrated for Scott Creek, a small catchment located in the hills $20\ km$ south of Adelaide, South Australia. The Scott Creek flow-gauge has 38 years of continuous record from 1969 to 2006. Adelaide rainfall is winter dominated, as shown in Figure 8.2, with average annual rainfall approximately $530\ mm$, and being slightly higher in the Adelaide Hills. Adelaide is not highly influenced by the ENSO phenomenon when compared to the Eastern States of Australia, thus climatic partitions were not considered for this investigation.

The objective of the design is to size the height of the downstream wall of a detention basin on this creek. The role of the seasonal extreme analysis is to determine if the seasonal influences can be exploited to allow for additional storage in the basin and the impact on the flood mitigating properties of the basin. Thus the basin may function as a
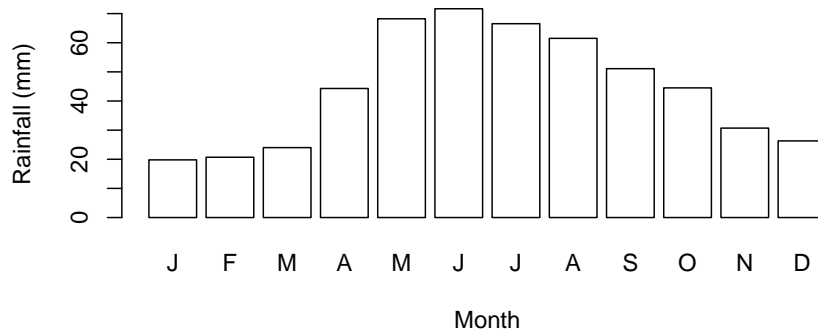
Figure 8.2 Average monthly rainfall, Adelaide, South Australia.

dual-purpose basin and the additional (seasonal) storage may be used for water harvesting, for example by pumping water down to the aquifer under an aquifer storage recovery scheme. The observed daily streamflow data were used as input to a simple basin model covering an area of $1000\ m^2$ and having two circular-culvert outlets of diameter $0.525\ m$. In this way, the observed streamflow volumes were converted to simulated water depths at the downstream wall of the basin.

Six intra-annual periods of 2-month lengths were used as this was considered a convenient seasonal aggregate for the analysis. These periods are denoted by their months, Jan/Feb, Mar/Apr, May/Jun, Jul/Aug, Sep/Oct and Nov/Dec. For each year, the seasonal maxima of water depth in the basin were determined and a log-normal distribution, having 2 parameters, was fitted. The simulated maximum water depths and the fitted distributions are shown in Figure 8.3 and have good agreement. It is important that the seasonal distributions are well fitted otherwise discrepancies between the directly simulated annual maxima and the annual maxima obtained as a composite of the seasonal distributions will be evident. Noting the logarithmic scale, it can be seen that the largest depths come from the late winter and early spring months (up to $5.6\ m$), while the summer and early autumn months are very dry (less than $0.15\ m$).

To provide detail on the relationship between rainfall and runoff in the Scott Creek catchment, a seasonal analysis was conducted using observed rainfall for the same 38 year period from a neighbouring catchment to Scott Creek. There is no sufficiently long gauge in Scott Creek itself. The annual maxima (symbols) show the different seasons that contributed to the maximum value in a given year and it can be seen that some of the highest rainfall bursts occurred in the months of January through to April ('Figure 8.4').
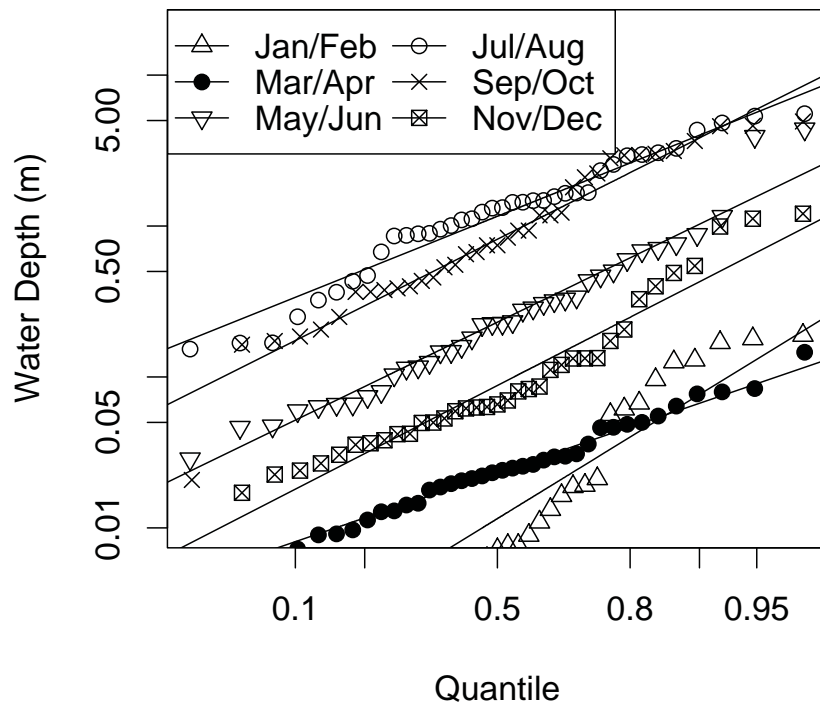
Figure 8.3 Maximum water depths in Scott Creek Basin for six two-months seasons plotted against a log-normal variate axis: simulated values (points) and fitted log-normal distributions (solid lines).

By comparison with Figure 8.3 it can be seen that these bursts did not yield corresponding high water depths. The likely explanation for this is attributed to the catchment dryness during summer months and the low water depths in the basin prior to the infrequent summer rainfall bursts. In short, the composition of extremes is not the same between rainfall and simulated water depths.

The correlation between maxima in several pairs of seasons was observed to be significantly different from zero at the 5% level. For this reason, Eq. 8.6 was not applied and it is necessary to apply Eq. 8.7 instead. Although methods exist to model correlation between each pair of seasons separately [*Zimmer and Trivedi*, 2006], a single correlation was considered between all seasons in order to reduce the overall number of parameters. The Gumbel copula (related to the Gumbel distribution by name only) was used to couple the six seasonal marginal distributions. It belongs to the Archimedian family of copulas
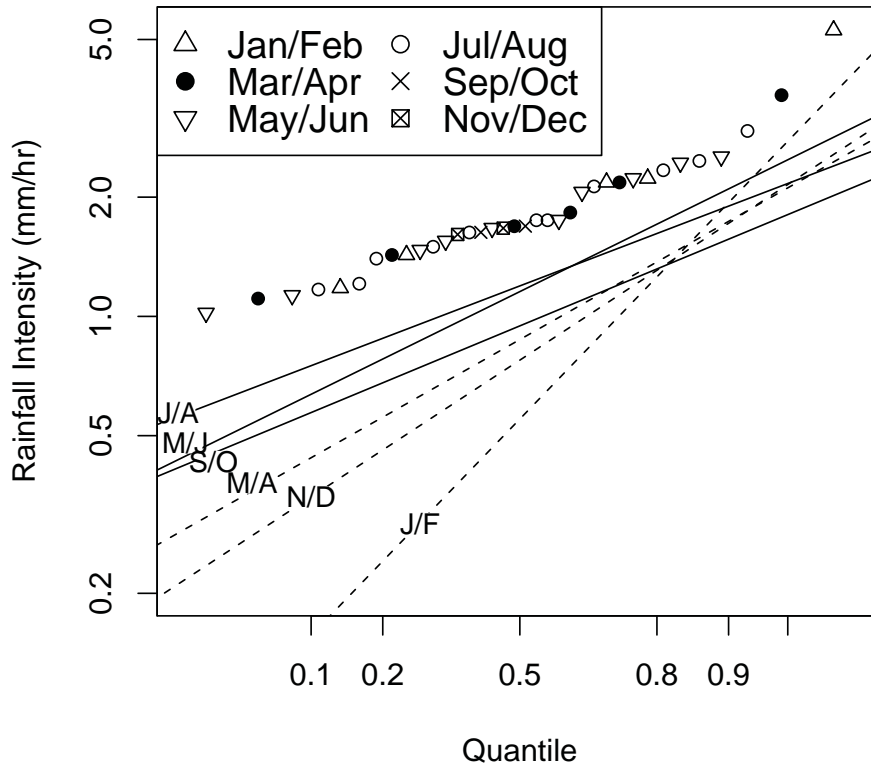
Figure 8.4   Maximum rainfall intensities in neighbouring catchment to Scott Creek for six two-months seasons plotted against a log-normal variate axis. Solid lines indicate 'wetter' months (May-Oct.), dashed lines indicate 'drier' months (Nov.-Apr.). Plotted symbols give the annual maximum distribution, where each symbol shows the season contributing to that annual maximum.

because it has the form

$$C(q, r, s, t, u, v) = \phi^{-1}(\phi(q) + \phi(r) + \phi(s) + \phi(t) + \phi(u) + \phi(v)) \qquad (8.9)$$

where ϕ is the copula function, and q, r, s, t, u and v are values between 0.0 and 1.0 that correspond to the six marginal distributions of maxima in each season: $F_X(X_1), \ldots,$ $\ldots, F_X(X_6)$. The Gumbel copula has one parameter, θ, which has the form $\phi(t) = (-\ln t)^\theta$, and was calibrated using the relation to Kendall's coefficient of correlation, τ, given as, $\tau = 1 - \theta^{-1}$. The Gumbel copula was suitable for this application because the seasons displayed only positive correlation.

Figure 8.5 shows, with good agreement, the annual maxima compared to a log-normal distribution fitted directly to those same maxima. Figure 8.5 also provides a comparison between the observed annual maxima and the composite-annual maxima, as obtained from Eq. 8.7 with log-normal seasonal distributions and the Gumbel copula. Despite

using only one parameter to model the correlation between all pairs of seasons, the copula shows a good agreement to the simulated maxima and the fitted log-normal distribution.
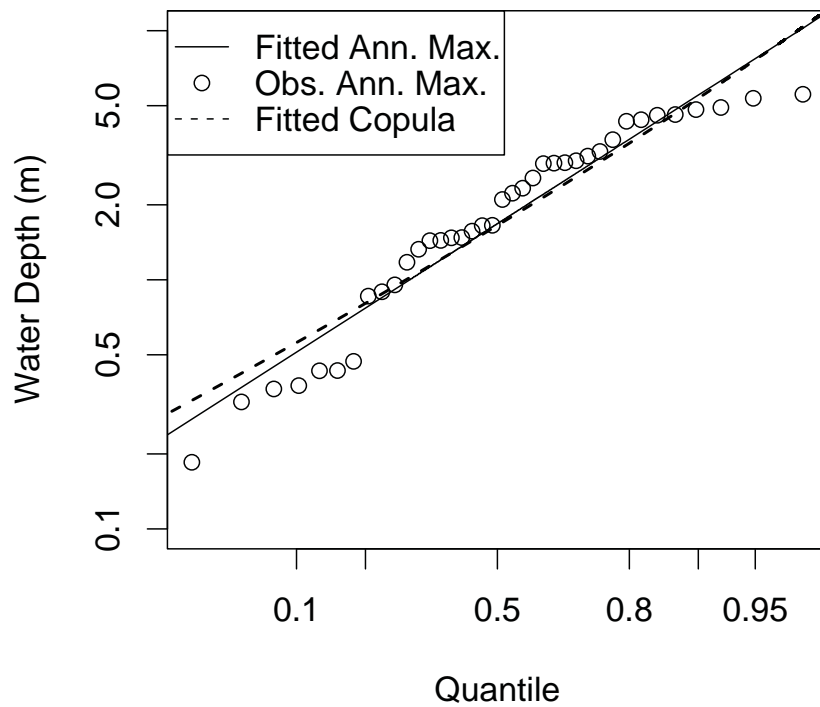

Figure 8.5   Annual maximum water depths in Scott Creek Basin: simulated values (points), fitted log-normal distributions (solid line) composite distribution from simulated seasonal maxima (dashed line).

To illustrate the utility of the methodology outlined in Section 8.2, suppose that the 5-year Average Recurrence Interval (ARI) is used to design the height of the dam and that a spillway is provided in the event of exceedances above this threshold. This frequency corresponds to the 0.8 quantile and based on Figure 8.5 requires the dam height to be $3.53\,m$. From Figure 8.3 it is evident that the months of January through to June, November and December do not experience maxima as large as July through to October. In other words, they have less impact on the annual maxima. Suppose then that a design specification is to use up to $1\,m$ of the basin for water storage during these months. This would provide room for $1\,ML$ of storage during months that coincide with the drier period of the Adelaide climate. The method for evaluating the flood-risk of the dam under this new

scenario is to shift the extreme value distributions of the seasons in question by $1\ m$. Using the shifted distributions with the previously modelled relation given by Eq. 8.7, a new distribution for the annual maxima can be obtained. Figure 8.6 shows the shifted seasonal maxima and the adjusted composite annual maximum distribution in comparison to the simulated annual maxima and the fitted annual maxima under the scenario of no storage.
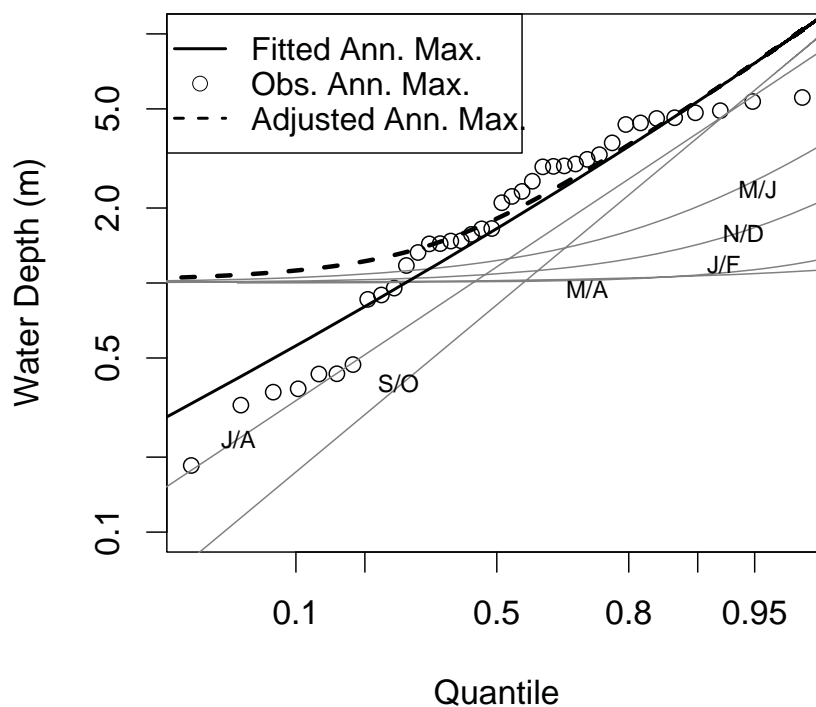

Figure 8.6 Maximum water depths in Scott Creek Basin. Bi-monthly seasonal distributions are indicated with beginning letters of each month (thin solid lines), simulated annual maxima (points), fitted log-normal annual maximum distribution (thick solid line), adjusted annual maximum composite distribution from simulated seasonal maxima (dashed line).

Due to the additional storage requirement, there is an increase in the annual flood-risk, but this is mostly evident as an increase in the lower tail of the composite distribution. Specifically, during the two-thirds of the year when there is $1\ m$ storage in the basin, there is a significant increase in the number of times that the annual maximum reaches one to two metres. Comparing the adjusted annual maximum distribution from Figure 8.6 to the annual maxima under the condition of no additional storage, it can be seen that

there is a negligible difference at the 0.8 quantile. From this curve, if $1\ m$ of water is stored in the basin for two-thirds of the year then in order to maintain the 5-year ARI the new height of the dam needs to be $3.59\ m$ (up from $3.53\ m$). Alternatively, if the dam height is not increased then the new exceedance ARI will be 4.9 years. While the quantities involved in this scenario are negligible, the principle remains the same for larger applications and for locations having different seasonal patterns. As the methodology is general it is applicable for simulation studies, for example if the detention basin has a more detailed operation rule-set involving draw-downs or aquifer storage pumps for diverting water. Similarly, the methodology can incorporate uncertainty in the design estimate by using multiple simulated rainfall-runoff replicates. Rainfall-runoff simulation may also be applicable if there is only a short streamflow record but a long rainfall record that shows seasonal and/or climatic variability. The efficacy of any dual-purpose system would need to be evaluated based on a comparison of the benefit of additional water storage and the additional cost (increased height/flood-risk and operating costs). The operating rules for this dual-purpose basin are straight forward due to the regular period of the seasons. By way of contrast, as the climate state is quasi-periodic there is an additional level of uncertainty in determining any change in the state. Had it been used, it would require a more complex operation of the basin.

8.4 CLIMATE CONDITIONED SNSRP METHODOLOGY

Temporal heterogeneity of rainfall occurs when there is structure in the variability of the rainfall due to different physical processes. Conceptual models of rainfall, in trying to explain this variability, can choose between having a minimal set of parameters that are sufficiently flexible to explain all possible types of observation, or multiple sets of parameters that are tailored to each set of processes. Whether or not models choose to explicitly address heterogeneity is a matter of how subtle the variation is and whether the increase in the number of parameters is in some way justifiable.

To address these aspects, a simple auto-regressive model of the climate state is incorporated into the model at a monthly time-step. Parameter values are conditioned on the separate climate states. The Southern Oscillation Index (SOI) is defined in terms of a ratio of pressures between Darwin and Tahiti. The SOI has been associated with El Niño and La Niña climate states which are attributed to anomalies in rainfall over eastern Australia. Figure 8.7 shows the SOI values for the years 1876 to 2007 and that it has a quasi-periodic nature.

The auto-correlation function (acf) and partial auto-correlation function (pacf) can be
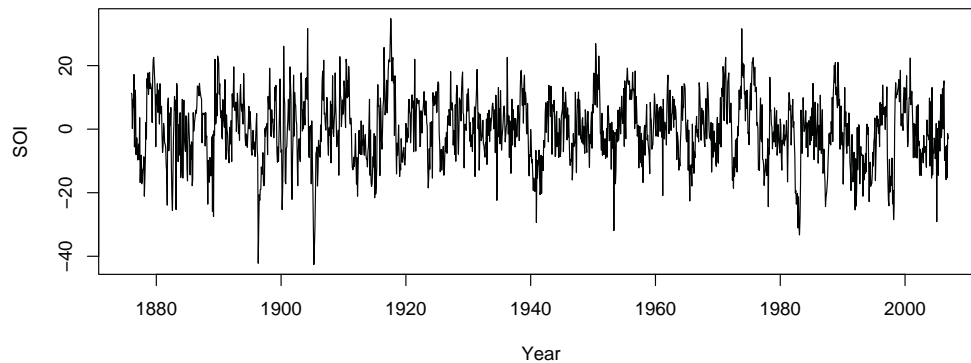
Figure 8.7   Southern Oscillation Index.

used to identify the structure in the SOI time-series. The acf in Figure 8.8 (a) shows significant positive correlations up to 10 months and slightly negative correlations from 14 to 26 months. An exponential decay in the acf suggests that an auto-regressive linear filter is an appropriate model. The pacf in Figure 8.8 (b) shows positive correlations up to 3 lags and the exponential decay in this plot suggests that a moving average linear filter is appropriate. Therefore, an Auto-regressive Moving Average (ARMA) model will be used to model the SOI time-series.
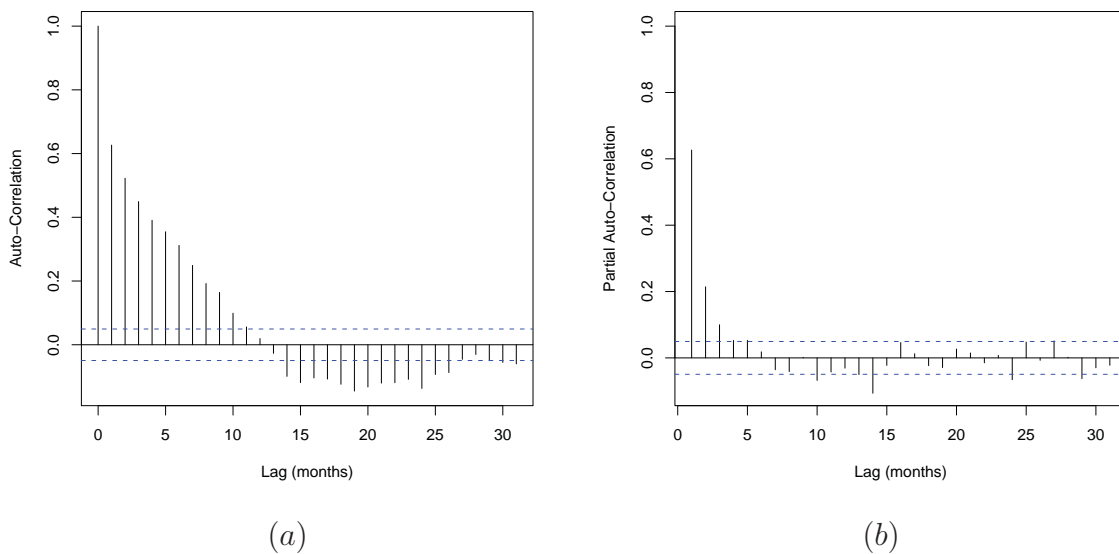

(a) (b)

Figure 8.8   Auto-correlogram and partial auto-correlogram of the SOI.

The generalised form of an ARMA(p,q) model of the SOI is given below

$$SOI_t = c + \sum_{i=1}^{p} \alpha_i SOI_{t-i} + \sum_{j=1}^{q} \beta_j \epsilon_{t-j} + \epsilon_t, \tag{8.10}$$

where t denotes the time step, c is a parameter constant, p is the number of auto-regressive terms, α_i is the auto-regressive coefficient, q is the number of moving average terms, β_j is a moving average coefficient and ϵ is a normally distributed error term. The variance of the error term is specified with an additional parameter, σ_ϵ^2. An ARMA(1,1) model was found to give the best fit to the SOI data (in terms of the Akaike Information Criterion) with the parameters shown in Table 8.4.

	α_1	β_1	c	σ_ϵ^2
Value	0.8528	-0.4010	-0.0797	62.57
Standard Error	0.0196	0.0352	0.8094	-

This model can be used to provide stochastic simulations of the SOI. A simple procedure is used to condition the rainfall model on the SOI. Two indicator variables are used for the events $SOI \geq 0$ and $SOI < 0$ which are denoted SOI^+ and SOI^- respectively. Figure 8.9 demonstrates the observed pattern of the SOI according to this definition of climate states.
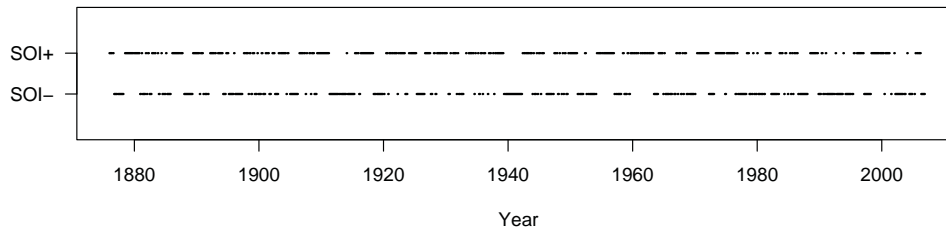

Figure 8.9   Pattern of two indicator states based on partition of SOI.

This definition of climate states can be used to partition the monthly data so that a separate set of calibration statistics is available for each month and each climate state on record. The independent calibration of the model to separate months (and separate states) proceeds as previously outlined. To simulate the climate index is first updated to indicate the current state and then the rainfall for that month is simulated using the parameters of the given state.

It is possible to specify more elaborate mechanisms for conditioning the rainfall model on the climate, but this model is sufficient for demonstrating improvement to observed statistics. One issue is that the climate state is never truly 'known' and that multiple definitions exist for how to define the state. One example is that additional states could be postulated (e.g. a neutral state for the SOI) and another example is that additional indices such as the Interdecadal Pacific Oscillation (IPO) could be used to define the state. A problem that arises with this when it is implemented for the current approach is that additional indices and states introduce more partitions to the data. The more partitions there are, the less data is in each partition and the more uncertain the statistic and parameter estimates become. An alternative approach that avoids partitioning, but is more involved than using climate states, would be to regress parameter values on the observed SOI and IPO values.

8.5 BOURKE CASE STUDY

Figure 8.10 shows the 16 daily rainfall gauges and the solitary pluviograph gauge located near the Bourke district in northern New South Wales. The details of the gauges are listed in Table 8.1. The pluviograph gauge spans 43 years, the shortest daily gauge spans 18 years and the longest gauge spans 120 years. The average span of the gauges is 60 years, but there are numerous years with missing observations at the majority of the sites. The average rainfall is approximately 350 mm per year, but it is highly variable and can range from under 50 mm to nearly 1000 mm in any given year. The amount of rainfall in each month is similar, but the storms are more intense in the summer months. As Bourke is an arid region, there is a very high portion of dry days.

The data for the Bourke region were pooled and a Neyman-Scott model was fitted using only one storm type. As the spatial process was not of primary interest in this chapter, typical values for the cell and storm radius were used. The skewness, autocorrelation and coefficient of variation were used in the calibration, except the 6 hour skewness was omitted. This was due to the general variability of this statistic and the small amount of data used to calculate it. The hourly skewness values are equally variable, but it is important to the calibration that some information about the sub-daily skewness is retained. The months were partitioned according to whether they belong to a SOI+ phase or a SOI- phase. Separate models were fitted to (a) the entire record, (b) the SOI- portions of the record and (c) the SOI+ portions of the record. Figure 8.11 summarises the goodness of fit of the calibration to the auto-correlation statistic. Figure 8.12 summarises the goodness of fit of the calibration to the coefficient of variation, and Figure 8.13 summarises
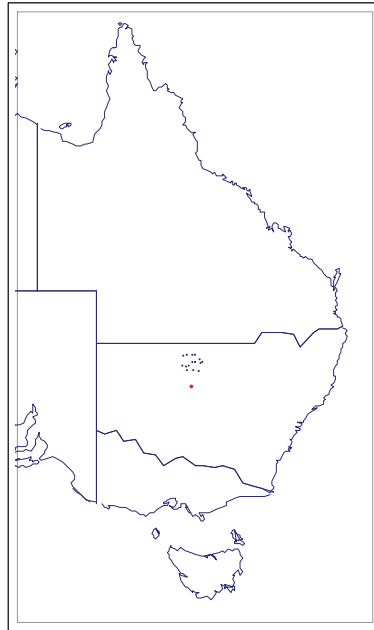
Figure 8.10 Locations of gauges used for Bourke case-study, red symbol is a pluviograph site and blue symbols are daily gauges.

the goodness of fit of the calibration to the skewness statistic.

The parameters corresponding to each of the partitions are summarised in Table 8.2 (entire record), Table 8.3 (SOI- partition) and 8.4 (SOI+ partition). The parameters from Table 8.2 were used as an initial starting point for the calibration for the two separate climatic partitions. Generally speaking the different sets of parameters should not be compared and contrasted since the conceptual framework can lead to parameter combinations of similar quality but radically different interpretation.

The two sets of parameters for the separate SOI partitions are combined with a synthetic simulation of the SOI. The parameters calibrated to the entire record are implemented separately and provide a comparison of the influence of the climatic model.

8.6 OBSERVED PARTITIONED EXTREMES

Figure 8.14 shows the annual extreme daily rainfall totals partitioned by the phase of the SOI. For each distribution a Gumbel distirbution was fitted and the 95% confidence limits are based on the parameter uncertainty for this distribution. This figure demonstrates that the annual extreme rainfall amounts are higher in the $SOI+$ state than in the $SOI-$ state. Figure 8.15 provides the same comparison, but is inspecting the maximums that occur in 'summer' months (Oct-Mar) and 'winter' months (Apr-Sep). This figure shows that there

Table 8.1   List of gauges used for Bourke case-study.

Type	Station ID	Latitude	Longitude	Description
Pluvio.	048027	-31.484	145.829	COBAR MO
Daily	048013	-30.092	145.936	BOURKE POST OFFI
Daily	048081	-29.933	146.233	WARRAWEENA
Daily	048049	-30.348	145.576	BOURKE (YANDA)
Daily	048076	-30.283	145.383	EAST TOORALE (TO
Daily	048042	-29.753	145.427	FORDS BRIDGE
Daily	048137	-30.1	146	BOURKE NORTH
Daily	048155	-30.083	145.867	FORT BOURKE
Daily	048128	-30.3	145.7	YANDA
Daily	048183	-30.167	146.317	MOUNT OXLEY
Daily	048117	-29.7	145.9	THE LAKE TANK
Daily	048026	-29.7	146	WARRELLA PLAINS
Daily	048067	-30.1	146.4	OAKLEIGH 2
Daily	048127	-30.583	145.933	YAHGUNYAH
Daily	048055	-29.7	145.6	LILA LOWER
Daily	048161	-30.6	146.2	KENILWORTH
Daily	048005	-30.558	145.595	LOUTH (BELAH)

Table 8.2   Bourke regional parameter estimates, entire record.

| k	$\hat{\lambda}$ h^{-1}	$\hat{\beta}$ h^{-1}	$\hat{\eta}$ h^{-1}	$\hat{\mu}_C$ -	$\hat{\alpha}$ -	$\hat{\phi}_c$ km^{-1}	$\hat{\phi}_s$ km^{-1}
1	0.00058	0.0745	4.66	23.4	0.766	0.01	0.001
2	0.00072	0.0337	1.48	10.7	0.677	0.01	0.001
3	0.00063	0.1257	2.49	22.5	0.865	0.01	0.001
4	0.00054	0.1316	4.35	30.6	1.278	0.01	0.001
5	0.00091	0.2113	3.06	48.5	0.586	0.01	0.001
6	0.00112	0.0984	1.04	11.7	0.781	0.01	0.001
7	0.00035	0.0033	0.5	33.9	0.526	0.01	0.001
8	0.00019	0.0021	0.37	33.9	0.67	0.01	0.001
9	0.00108	0.1444	1.84	15.1	0.663	0.01	0.001
10	0.0021	0.1147	1.36	8.1	0.555	0.01	0.001
11	0.00144	0.0974	1.46	6.9	0.632	0.01	0.001
12	0.00098	0.0734	1.73	9.5	0.681	0.01	0.001
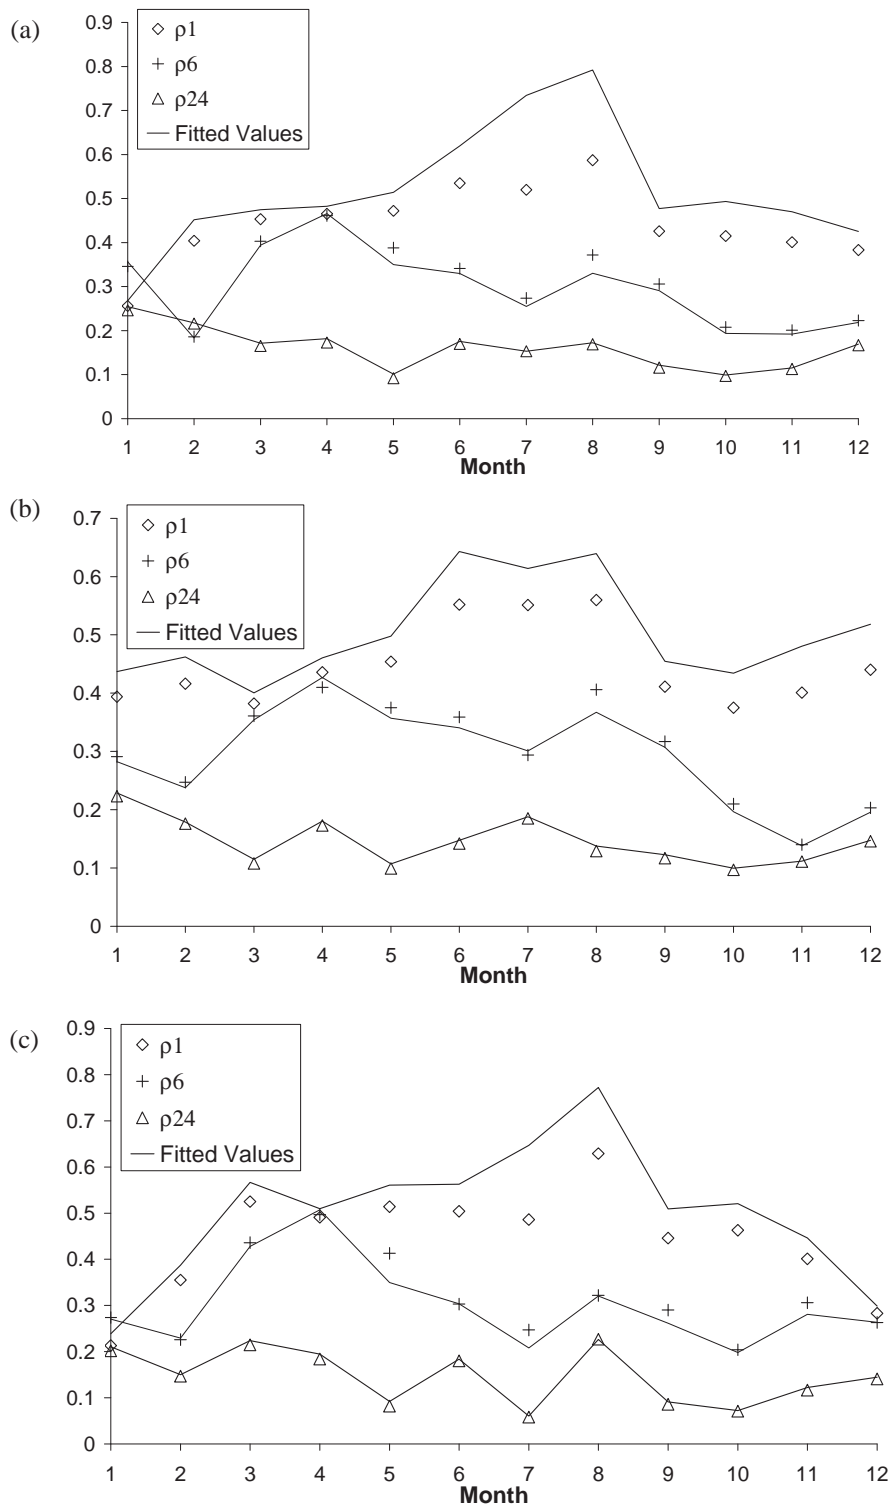
Figure 8.11 Fitted correlation statistics for Bourke region, (a) Entire Record, (b) SOI- and (c) SOI+.
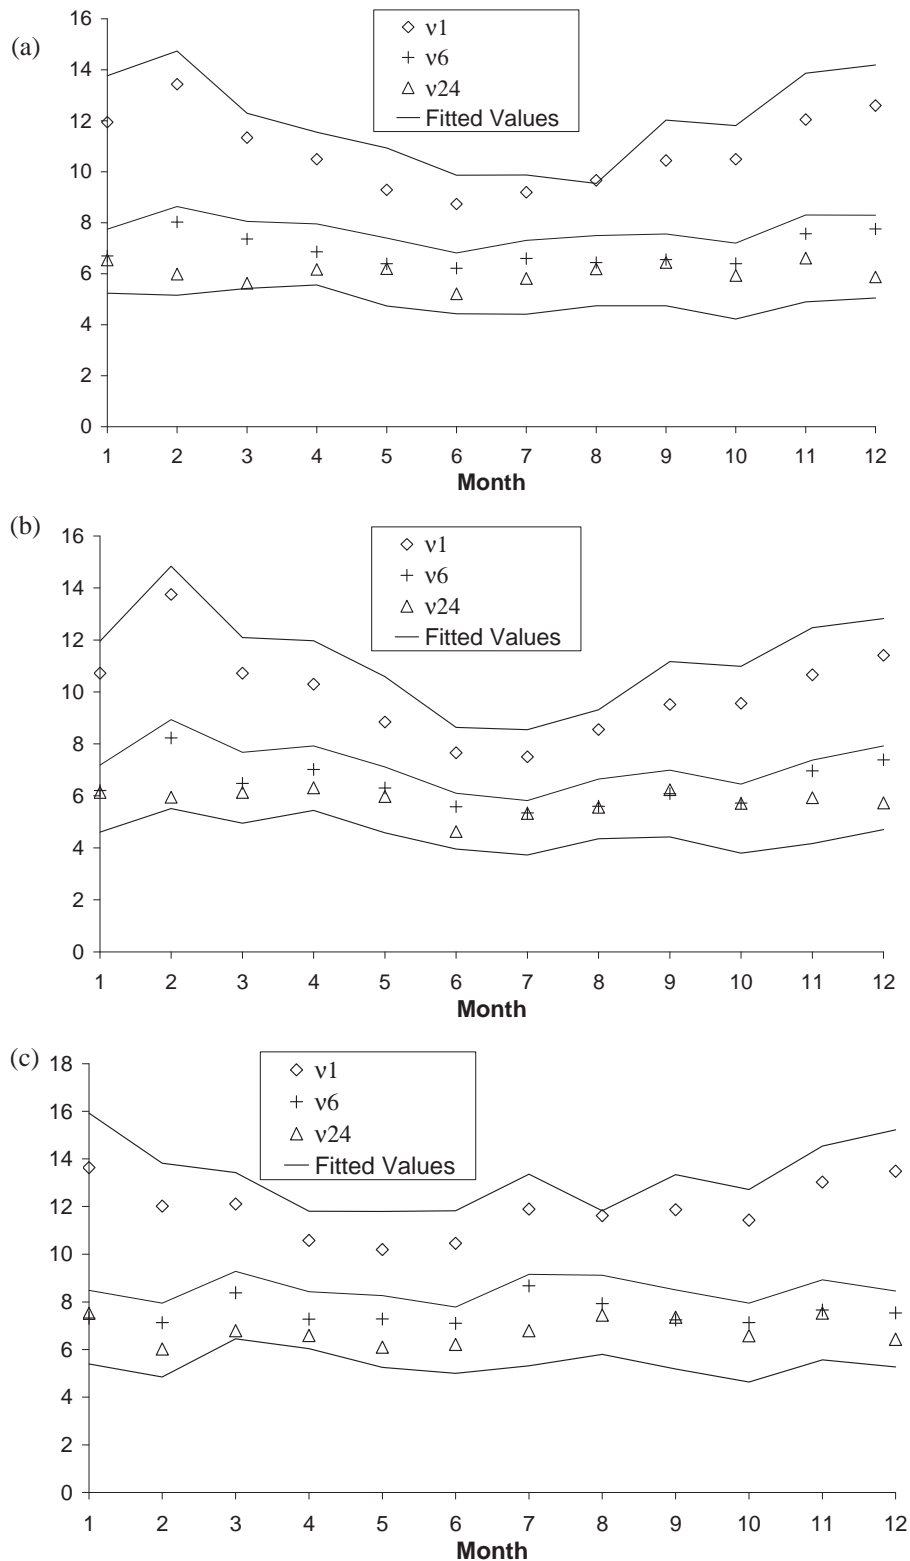
Figure 8.12   Fitted coefficient of variation for Bourke region, (a) Entire Record, (b) SOI-
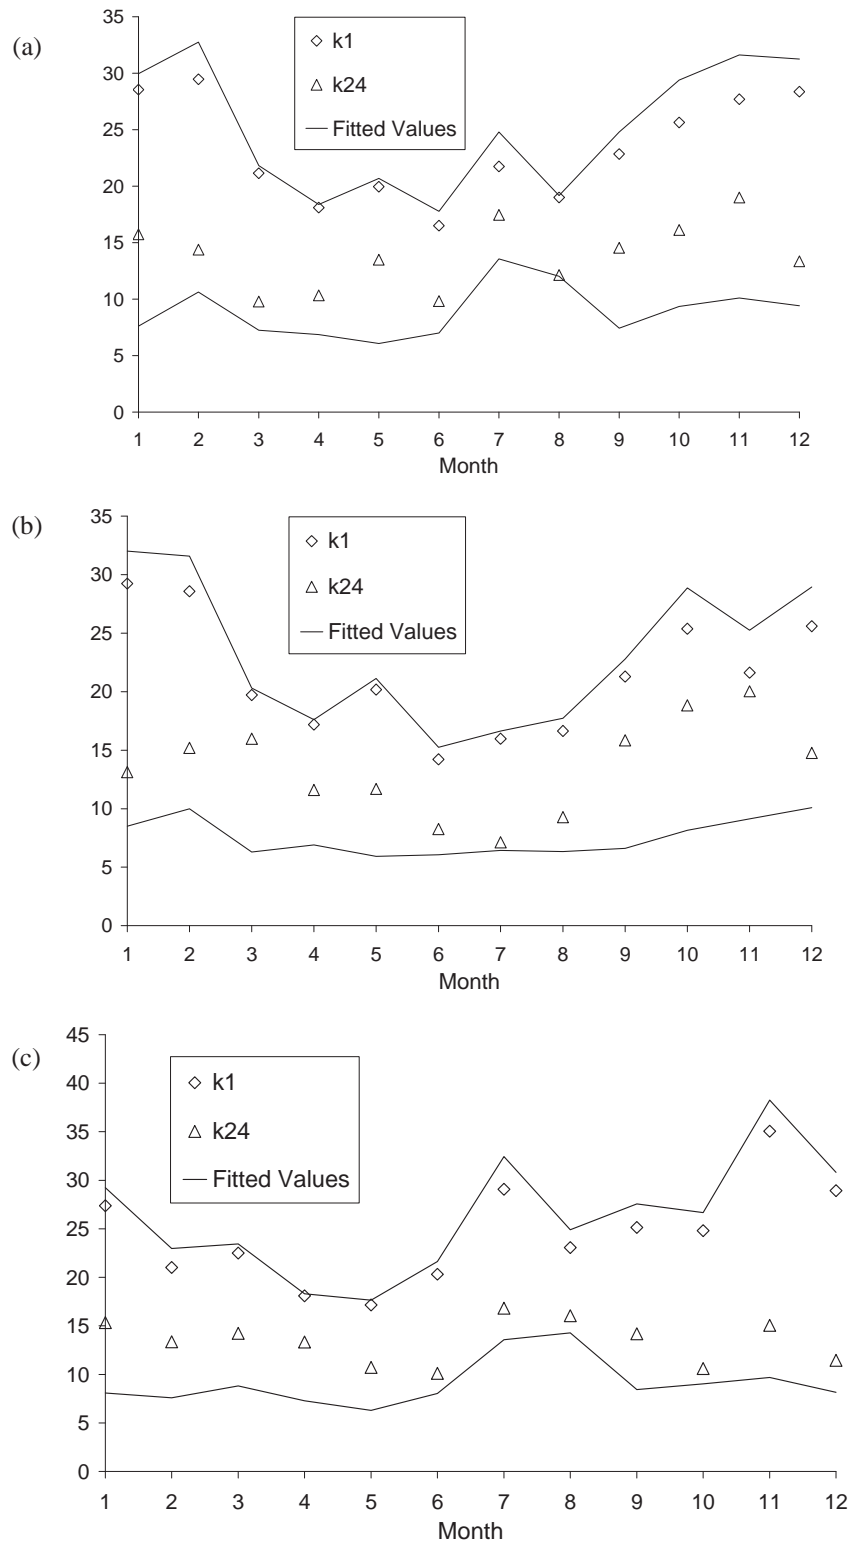and (c) SOI+.

Figure 8.13   Fitted skewness for Bourke region, (a) Entire Record, (b) SOI- and (c) SOI+.

Table 8.3   Bourke regional parameter estimates, SOI-.

k	$\hat{\lambda}$ h^{-1}	$\hat{\beta}$ h^{-1}	$\hat{\eta}$ h^{-1}	$\hat{\mu_C}$ -	$\hat{\alpha}$ -	$\hat{\phi_c}$ km^{-1}	$\hat{\phi_s}$ km^{-1}
1	0.00088	0.0656	1.8	28.1	0.509	0.01	0.001
2	0.00078	0.0734	1.56	10.2	0.693	0.01	0.001
3	0.00082	0.1843	4.92	18.1	1.155	0.01	0.001
4	0.00059	0.1254	3.18	14.5	2.876	0.01	0.001
5	0.00096	0.2004	3.21	55.5	0.546	0.01	0.001
6	0.00146	0.1243	1.02	12.3	0.777	0.01	0.001
7	0.00162	0.0799	0.98	11.4	0.702	0.01	0.001
8	0.00112	0.1477	1.21	26.5	0.602	0.01	0.001
9	0.00117	0.1505	2.16	17.7	0.671	0.01	0.001
10	0.00239	0.1236	1.7	10.4	0.532	0.01	0.001
11	0.00248	0.0514	1.29	3.2	0.773	0.01	0.001
12	0.00139	0.0657	1.21	7.4	0.631	0.01	0.001

Table 8.4   Bourke regional parameter estimates, SOI+.

k	$\hat{\lambda}$ h^{-1}	$\hat{\beta}$ h^{-1}	$\hat{\eta}$ h^{-1}	$\hat{\mu_C}$ -	$\hat{\alpha}$ -	$\hat{\phi_c}$ km^{-1}	$\hat{\phi_s}$ km^{-1}
1	0.00058	0.077	1.5	32.8	0.809	0.13	0.01
2	0.00117	0.21	2.96	50.2	0.549	0.01	0.01
3	0.00188	0.183	3.22	20	0.575	0.3	0.01
4	0.00251	0.106	2.01	20.2	0.628	0.17	0.017
5	0.00523	0.121	2.51	36	0.568	0.15	0.01
6	0.00732	0.12	1.62	41	0.488	0.15	0.014
7	0.00823	0.11	2.17	73	0.433	0.2	0.01
8	0.00731	0.182	2.31	105.1	0.389	0.15	0.015
9	0.0058	0.165	3.66	44	0.493	0.17	0.02
10	0.00345	0.194	2.04	32	0.495	0.17	0.018
11	0.00293	0.124	3.4	28	0.478	0.16	0.013
12	0.00163	0.187	3.4	24	0.521	0.17	0.01

is a noticeable difference in rainfall associated with the $SOI-$ when inspecting summer extremes, but the difference in rainfall associated with the SOI is not as significant when inspecting winter extremes.
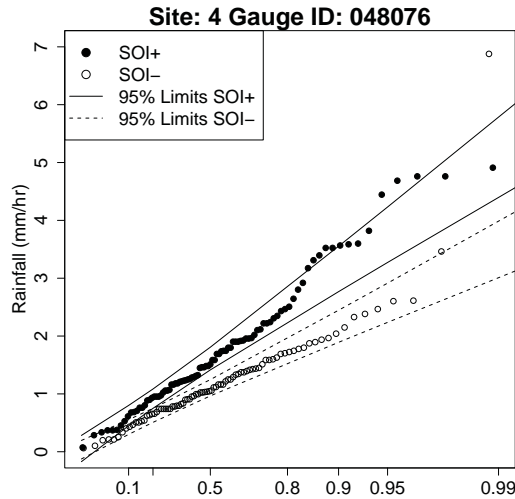

Figure 8.14 Annual extremes of daily rainfall at Bourke partitioned by +ve/-ve phases of the SOI. Probabilities shown using a Gumbel axis. 95% Confidence intervals obtained from the distribution of estimates of the Gumbel parameters.
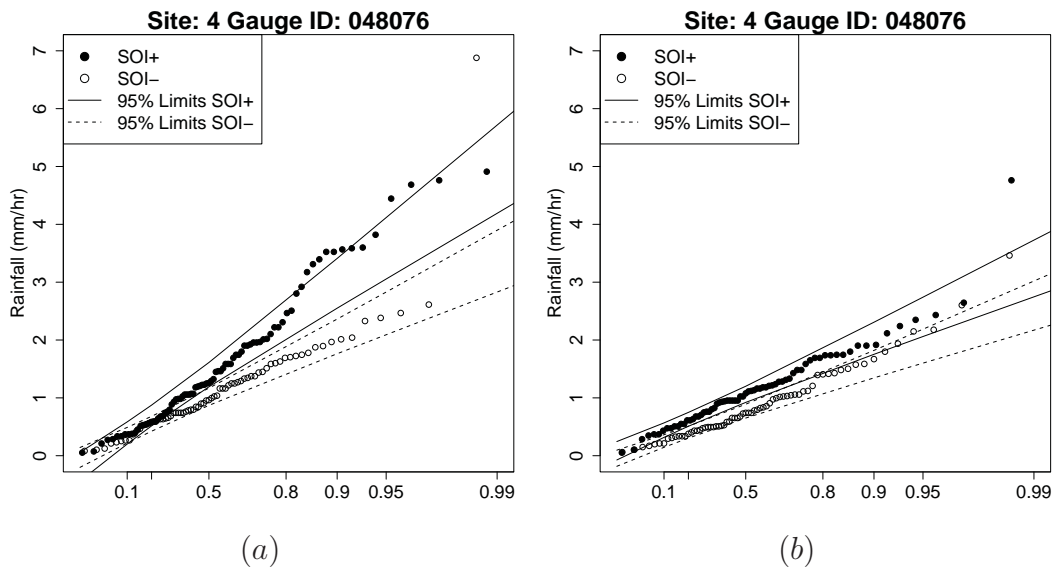

Figure 8.15 Extremes of daily rainfall at Bourke partitioned by +ve/-ve phases of the SOI (a) Summer extremes, (b) Winter extremes. Probabilities shown using a Gumbel axis. 95% Confidence intervals obtained from the distribution of estimates of the Gumbel parameters. Site 4.

Figure 8.16 shows that summer extremes are higher than winter extremes and that the annual distribution closely follows the distribution of summer extremes. Comparing Figure 8.16 (a) and Figure 8.16 (b), there is a stronger distinction between the two seasons during the $SOI+$ phase than during the $SOI-$ phase, as indicated by the overlap of the confidence limits.
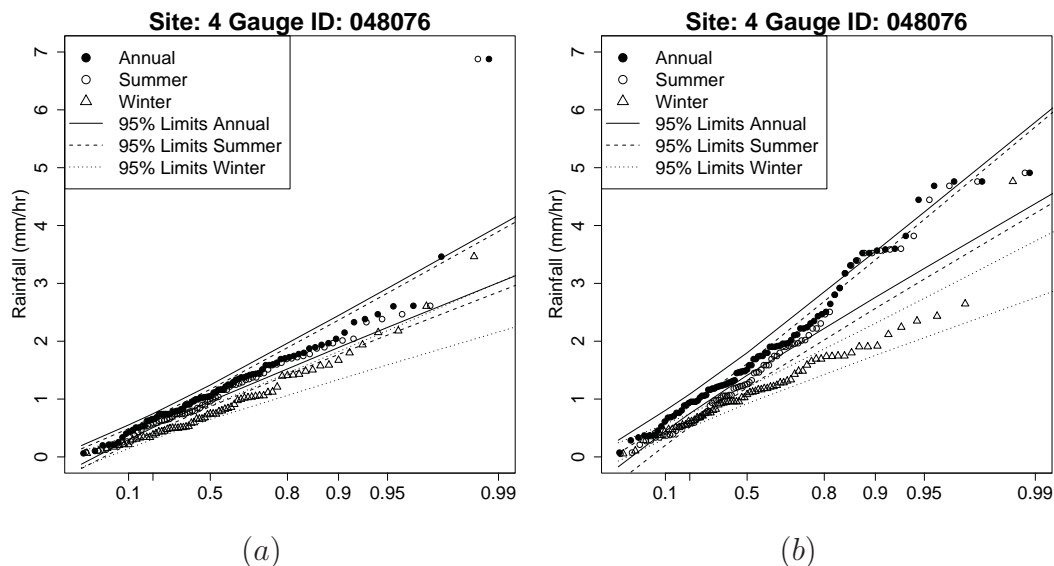

Figure 8.16 Annual and 6-month seasonal extremes of daily rainfall at Bourke (a) for SOI- and (b) for SOI+. Probabilities shown using a Gumbel axis. 95% Confidence intervals obtained from the distribution of estimates of the Gumbel parameters. Site 4.

8.7 COMPARISON OF OBSERVED AND SIMULATED EXTREMES

Simulations were conducted using 1000 replicates of length 100 years. The results for site 4 are presented in the body of the chapter for illustrative purposes. Equivalent plots for all other daily rainfall sites are available in Appendix B. Site 4 has a record length of 97 years so the observed sampling variability is similar to the simulated sampling variability for this gauge. Figure 8.17 (a) shows the comparison of the simulated annual maxima to the observed annual maxima. It is clear that the model is undersimulating the extremes at this site. Figure 8.17 (b) provides insight into the performance of the model with respect to each season. This figure shows that the winter extremes are consistently under-simulated and that the summer extremes match at the median, but the variance of the distribution is undersimulated. As a result the upper tail of the summer extremes is undersimulated but the lower tail is oversimulated. This oversimulation does not show in the annual extremes

since the annual extremes are comprised of the upper tails of both the summer and winter extremes. In other words, a simulated winter maximum is likely to be higher than a summer maximum simulated from the lower tail.
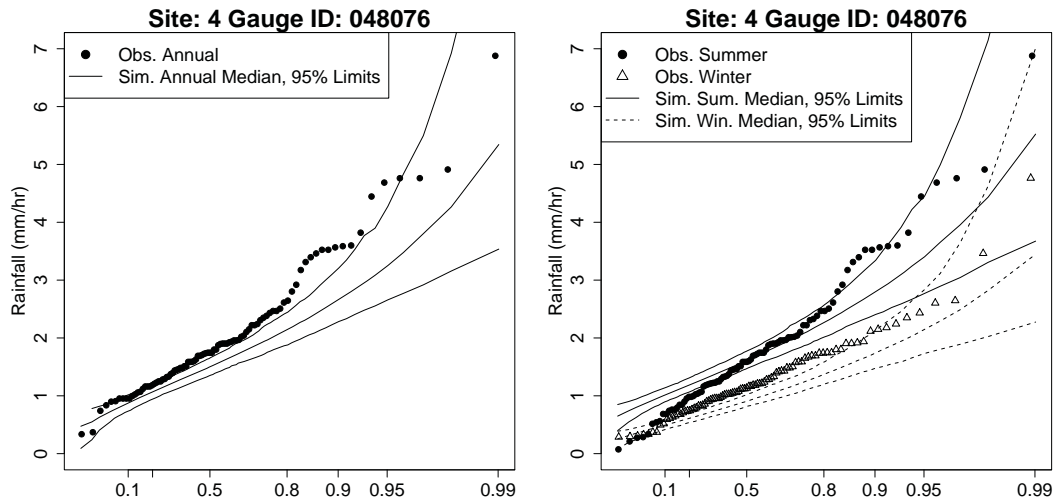

Figure 8.17 L.H.S. compares observed annual extremes with the median and 95% limits of extremes from simulated daily records. R.H.S. figures compare observed 6-month seasonal extremes with simulated confidence limits. Simulated values from 1000 replicates of 100-year records. Probabilities shown using a Gumbel axis. Site 4.

8.8 CONCLUSIONS

Previous studies have highlighted the fact that annual extremes are not independent and identically distributed due to the presence of climatic epochs [*Franks and Kuczera*, 2002]. This chapter has extended previous work by including intra-annual (seasonal) partitions of a dataset in addition to inter-annual (climatic) partition. In particular, the estimation of the distribution of annual maxima is considered, which is of practical significance given the role that intensity-duration-frequency curves of extreme variables have in the design of water-related infrastructure.

There is a broad range of possible applications for incorporating seasonal and climatic partitions into the design concepts for water-infrastructure. Examples include, (i) investigating that simulation models correctly reproduce the seasonal composition of extremes, (ii) verifying whether seasonal rainfall extremes give streamflow extremes of similar recurrence interval, (iii) inferring seasonal/climatic conditional levels of risk and necessary draw-down levels ,(iv) design of dual-purpose water basins for water retention and flood

mitigation, (v) allowing for in-stream processes in a design such as placing aquifer-storage pumps in a flood detention basin, (vi) design of structures or channels where seasonal flooding is acceptable and (vii) assessing the influence of climate change scenarios on extreme values under changing seasonal trends.

A simple illustration was provided to demonstrate the utility of seasonal partitions for a hypothetical design involving a dual-purpose basin. It highlighted that extreme summer rainfall did not correspond to extreme water depths in the basin and that as a result, a significant level of seasonal storage could be provided with only a marginal increase in the height of the downstream wall.

A second case study was provided to demonstrate the application of seasonal partitioning in the verification of the SNSRP model. The rainfall record was partitioned according to a climatic index and separate models were fitted to each partition. This procedure is not specific to the $SNSRP$ model and whatever the model, it requires the duplication of parameters. It would be possible to reduce the overall number of parameters by making some common across climatic states, but that was not the main focus of this investigation. The main aim of this chapter was to demonstrate that climatic and seasonal partitions are an important part of extreme rainfall. Partitioning an observed extreme allows for a better diagnosis of areas where the model may be performing poorly. This can either be with respect to climatic partitions or seasonal partitions.

Chapter 9
Spatially Inhomogeneous Neyman-Scott Model

9.1 INTRODUCTION

The term homogeneous is used to suggest that underlying rainfall statistics are the same across the domain which a model is applied to. Whether applied to the temporal or spatial domain, this assumption is often justified on the grounds of pragmatism, but it is rarely met in a strict sense. In the time domain, heterogeneity of the rainfall process has been well accounted for. For example, multiple storm types have been used [*Sansom*, 1999], multiple parameter sets and harmonic functions have been used to account for monthly variation [*Heneker et al.*, 2001; *Cowpertwait*, 2004] and long-term persistence models have been used to allow for climatic modulations [*Thyer and Kuczera*, 2003]. These advances have mainly been implemented for point rainfall models which are more mature than space-time models. While point rainfall models do well to reproduce temporal variability, they are intrinsically poor at capturing spatial heterogeneity, since without evidence to the contrary from additional gauges, the point of interest is assumed to be representative of the surrounding region. Even though space-time models have the capacity to improve the spatial representation, they too are often assumed to be homogeneous over some region. This is irrespective of whether the rainfall observations are made via radar or rain-gauge networks. The reasoning for this is that the application of rainfall models has typically been at smaller spatial scales common to urban environments, and examples include studies of flood mitigation, stormwater re-use and sewer overflows [*Cowpertwait et al.*, 1996a; *Sempere-Torres et al.*, 1999; *Willems and Berlamont*, 1999; *Hardy et al.*, 2005]. While some models provide scaling factors to account for variation in the mean rainfall, they rarely consider changes in other statistics, for example the probability of a dry period. This is largely due to the significant increase in complexity needed to account

for spatial heterogeneity above a homogeneous or stationary formulation.

As a region becomes larger the assumption of homogeneity becomes harder to justify since geographic effects make differences in the rainfall statistics more pronounced. Arising from this, some models allow for regionalisation, where the model parameters (usually from a point rainfall model) are scaled according to variables such as elevation and distance from the coast or such that key rainfall statistics are preserved. Despite this, regionalisation does not necessarily address the issue of heterogeneity within a region, only that variation occurs between regions. This is because, having translated the parameters across a large distance, they are then assumed to be representative of rainfall for the catchment of interest regardless of how large it is.

Jothityangkoon et al. [2000] provide a good example of rainfall heterogeneity using the Avon River basin, which is a $400\ km \times\ 400\ km$ region in south-west Western Australia, shown below in Figure 9.1. The region has very little relief, with the Darling escarpment surrounding Perth at a distance of approximately 50 km from the coast and the inland region forms a low plateau with an average elevation of 300 m. The region has a strong rainfall gradient associated with the distance inland from the coastline. The region has a dry interior and higher totals on the western coastline with respect to the southern coastline due to the prevalence of storm fronts from Indian Ocean.

NOTE:
This figure is included on page 138
of the print copy of the thesis held in
the University of Adelaide Library.

Figure 9.1   Study region of *Jothityangkoon et al.* [2000] encompassing Avon Basin.

Jothityangkoon et al. [2000] used a fractal-cascade model to spatially disaggregate lumped totals of daily rainfall across the region. While most fractal models are stationary over a given region, *Jothityangkoon et al.* [2000] proposed a weighted cascade such that the rainfall is biased toward the coastline in preference to the interior. The complexity of the observed data is reflected in the high number of weights required to obtain accurate disaggregations down to a $12.5\ km$ resolution: 1364 weights per month. This case study will be considered in more detail in this chapter for the development of an inhomogeneous version of the Neyman-Scott model.

9.2 MODEL DEVELOPMENT

9.2.1 Spatial process

An inhomogeneous poisson random field model is specified on a two dimensional grid where each pixel has a specified spatial rate. The spatial rate determines the mean number of storm cells from a Poisson distribution that land in that pixel. The model is inhomogeneous because the spatial rate varies in each cell, thus some regions will have more cells than others. The coordinates of each cell are located within the pixel with uniform probability. The cells are assumed to occupy a circular area with a random variable for the radius and a random intensity giving the constant amount of rainfall contributed at all points that the cell covers. The cell radius is assumed to follow an exponential distribution with parameter, ϕ_c. For numerical reasons this parameter is assumed to be the same for all pixels. Figure 9.2 shows a schematic of the model where one of the pixels is shown to have five rain cells each having different coordinates and radii. Where the cells overlap at a point the total intensity is the sum of the intensities from each overlapping cell. The total rainfall for a pixel is the sum over all points within that pixel.

9.3 INTEGRAL APPROXIMATION

Consider two pixels D_{ij} and D_{kl} where the subscripts refer to the row and column and each pixel has an area L^2. The pixels are separated by a distances ΔX and ΔY in the respective dimensions shown in Figure 9.3. It is necessary to determine the probability that an event landing somewhere in D_{ij} will cover a point landing somewhere in D_{kl}.

The probability of overlap is evaluated as the expected value from integrating over all points within the two respective pixels. Therefore, consider the points (u, v) and (x, y) which are the dummy coordinates used when integrating over the pixels. For an event that occurs at the coordinate (u, v) and has a random radius R, the probability that it covers
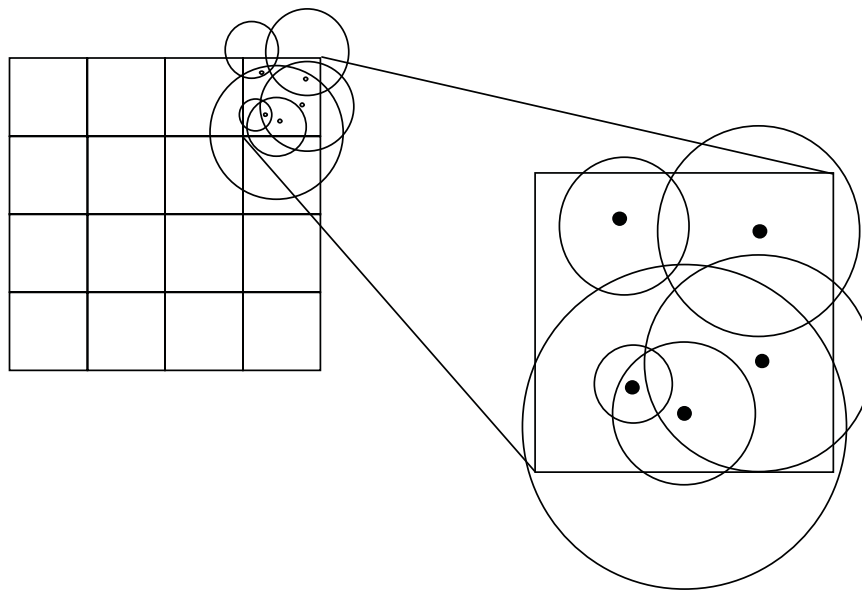
Figure 9.2 Schematic of rainfall process, discs represent raincells of a given storm from one pixel only.
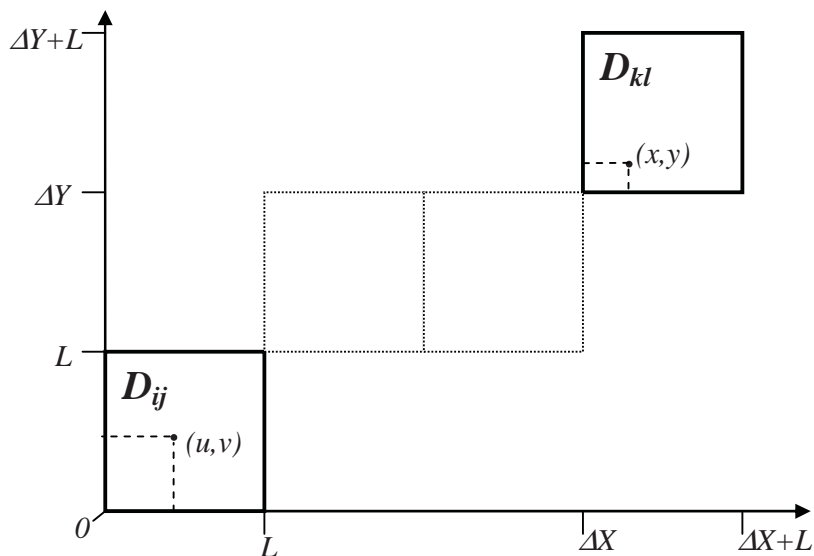

Figure 9.3 Schematic of two pixels from gridded rainfall data. Points *(u,v)* and *(x,y)* represent dummy coordinates used when integrating over the respective pixels.

the point (x, y) is,

$$P(R \geq \|(u, v) - (x, y)\|); (u, v) \in D_{i,j}, (x, y) \in D_{k,l}$$

$$= 1/L^4 \int_{u=0}^{L} \int_{v=0}^{L} \int_{x=\Delta X}^{\Delta X+L} \int_{y=\Delta Y}^{\Delta Y+L} e^{-\phi_c \sqrt{(x-u)^2 - (y-v)^2}} \, dx \, dy \, dv \, du$$

(9.1)

Where $|| \; ||$ indicates the distance which is evaluated as $\sqrt{(x-u)^2 + (y-v)^2}$. The exponential term occurs due to the exponential distribution for cell radii, but it is not possible to derive an analytic solution for the probability of overlap due to the nonlinearity of the term inside the exponential. It is possible to employ numerical procedures to solve Eq. 9.1, however these are not desirable owing to the large number of pixel combinations at which the overlap probability is to be evaluated. A further reason is that the parameter ϕ_c needs to be estimated and thus it is likely to repeat the procedure numerous times. It is possible however to find a linear approximation that can be analytically integrated. Figure 9.4 depicts the hypotenuse which is the distance between the two points (u, v) and (x, y). From this figure it is evident that the linear sum of the right angled sides will always provide an upper estimate of the length of the hypotenuse. Using just one of the triangle's sides as an estimate of the hypotenuse will always give a shorter distance, and if the longer of the two sides is selected then this will give the better estimate of the two.
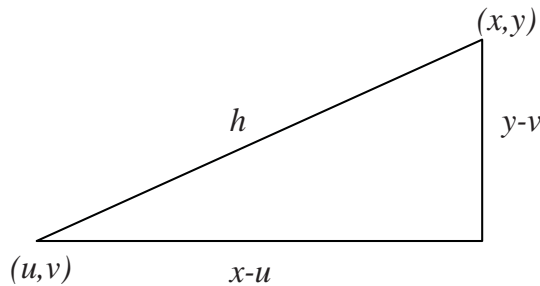

Figure 9.4   Schematic of distance between two points *(u,v)* and *(x,y)*.

If the longer distance is denoted h^+ and the shorter h^-, then a reasonable approximation of the hypotenuse is given as the average of the two, as shown in Eq. 9.2. Triangles with even sides will give poorer approximations than when one side is considerably longer than the other.

$$
\begin{aligned}
h^- &= max(|x-u|, |y-v|) \leq h \\
h^+ &= |x-u| + |y-v| \geq h \\
h &\approx \tfrac{1}{2}(h^+ + h^-)
\end{aligned}
\tag{9.2}
$$

Eq. 9.2 is a linear approximation for the distance. Due to the linearity, the estimated probability of overlap can be approximated as the average of probability of overlapping the longer distance and the probability of overlapping the shorter distance, as shown in

Eq. 9.3.

$$P(R \geq \|(u,v) - (x,y)\|); (u,v) \in D_{i,j}, (x,y) \in D_{k,l}$$
$$\approx P(R \geq \tfrac{1}{2}(h^+ + h^-)) \tag{9.3}$$
$$\approx \tfrac{1}{2}P(R \geq h^+) + \tfrac{1}{2}P(R \geq h^-)$$

Substituting Eq. 9.2 into Eq. 9.1, the analytic expression for the probability of overlap can be evaluated.

Evaluating the term $P(R \geq h^+)$ is straightforward since the exponential term in Eq. 9.1 becomes $e^{|x-u|+|y-v|}$ which is a linear sum. There are four separate cases needed to evaluate this integral to ensure that the distance terms are non-negative: (i) $x \geq u, y \geq v$ (ii) $x \geq u, y < v$, (iii) $x < u, y \geq v$, (iv) $x < u, y < v$.
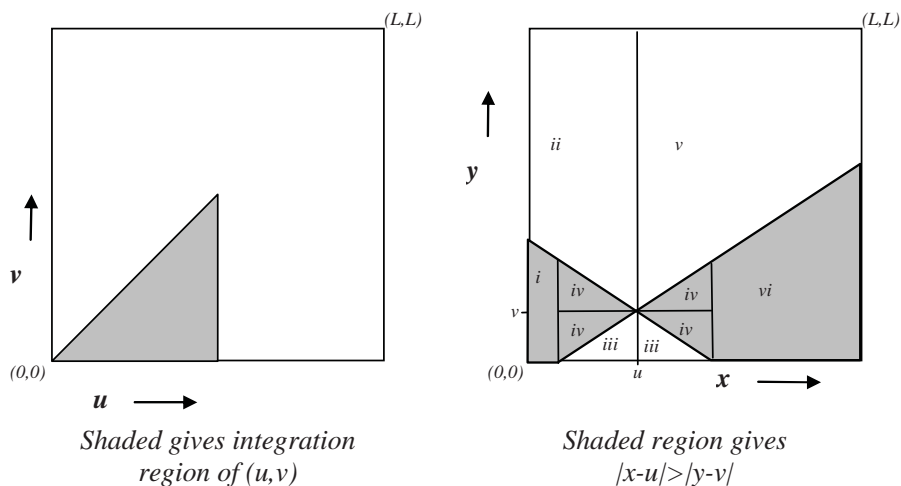

Shaded gives integration region of (u,v)

Shaded region gives |x-u|>|y-v|

Figure 9.5 Integration regions used for h^- case with $i = k; j = l$ (same cell). Roman numerals assigned in order of appearance in Eq. 9.4.

Evaluating the term $P(R \geq h^-)$ is more complicated than $P(R \geq h^+)$ since it is necessary to determine which side of the triangle in Figure 9.4 is longer. This constraint places bounds on the integrals given in Eq. 9.1. The following section outlines the method for evaluating this integral and there are different scenarios depending on whether the two pixels are (i) the same, (ii) in either the same row or column or (iii) on the diagonal to each other. Only the cases of the same pixel and diagonal pixels are outlined here as the same row (column) scenarios can be obtained by applying similar principles. Figure 9.5 shows the two domains of integration $D_{i,j}$ and $D_{k,l}$ which represent the same pixel (i.e. the pixel that a raincell originates within does not necessarily cover all points in that pixel). The shaded region inside the $u - v$ plane shows the integration region for u and v. All other

portions of this plane by symmetry will give the same total, thus a multiple of 8 is used in the resulting equation .

For a given point (u, v), the $x - y$ plane in Figure 9.5 shows 6 different regions which correspond to the integration terms in Eq. 9.4. The shaded regions show those where $|x - u| > |y - v|$, that is, the boundary within which one side of the triangle is longer than the other side. The conditions $x > u$ and $y > v$ also control the bounds of integration since the distance must always be positive, which is the reason for multiple regions. The numbering for some of the regions is repeated owing to arguments of symmetry.
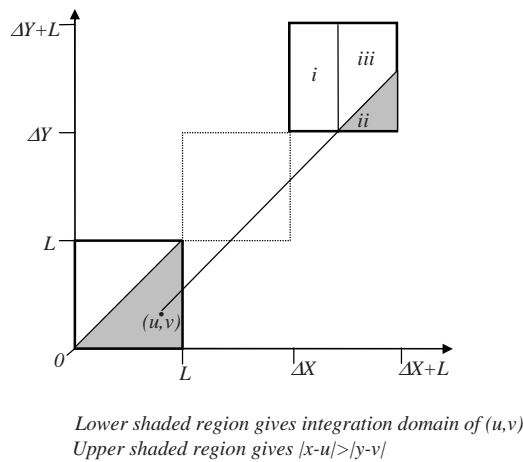

Lower shaded region gives integration domain of (u,v)
Upper shaded region gives |x-u|>|y-v|

Figure 9.6 Integration regions used for h^- case with $|i - k| = |j = l| \geq 0$ (strictly diagonal pixels). Roman numerals assigned in order of appearance in Eq. 9.4.

Figure 9.6 shows the orientation of two pixels that are on the same diagonal, separated by some distance $\Delta X = \Delta Y$. The shaded region of the $u - v$ plane gives the integration bounds used for du and dv in Eq. 9.4. The remaining area gives the same total due to symmetry and the integral is therefore multiplied by two. The $x - y$ domain is split into three regions which appear in order of number in Eq. 9.4. Region (i) is bounded on $(0, L)$ up to the point at which the line from (u, v) intersects the pixel. Beyond this point there are two regions which have different distances in the exponent of the integrand because of the condition $|x - u| > |y - v|$.

The resulting analytic expressions are shown in Eq. 9.5. While this equation looks 'large' it is an efficient means to estimate the probability. An observation from these equations is that the probability that a cell overlaps itself is not strictly equal to one. This is because the point (u, v) which the cell extends from does not cover all other points (x, y) within that same cell since there is a chance the cell radius is smaller than the pixel.

$$P(R \geq h^+) = \begin{cases} \dfrac{1}{L^4} \displaystyle\int_0^L \int_0^L \left(\begin{array}{l} \displaystyle\int_0^u \int_0^v e^{-\phi_c(u-x+v-y)} dy\, dx + \int_0^u \int_v^L e^{-\phi_c(u-x+y-v)} dy\, dx + \\[2mm] \displaystyle\int_u^L \int_0^v e^{-\phi_c(x-u+v-y)} dy\, dx + \int_u^L \int_v^L e^{-\phi_c(x-u+y-v)} dy\, dx \end{array} \right) dv\, du & i=k\,,\ j=l\ (\textit{same cell}) \\[10mm]

\dfrac{1}{L^4} \displaystyle\int_0^L \int_0^L \left(\int_{\Delta X}^{\Delta X+L} \int_0^v e^{-\phi_c(x-u+v-y)} dy\, dx + \int_{\Delta X}^{\Delta X+L} \int_v^L e^{-\phi_c(x-u+y-v)} dy\, dx \right) dv\, du & i=k\,,\ j \neq l\ (\textit{same row}) \\[8mm]

\dfrac{1}{L^4} \displaystyle\int_0^L \int_0^L \left(\int_0^u \int_{\Delta Y}^{\Delta Y+L} e^{-\phi_c(u-x+y-v)} dy\, dx + \int_u^L \int_{\Delta Y+L}^{\Delta Y+L} e^{-\phi_c(x-u+y-v)} dy\, dx \right) dv\, du & i \neq k\,,\ j=l\ (\textit{same col.}) \\[8mm]

\dfrac{1}{L^4} \displaystyle\int_0^L \int_0^L \int_{\Delta X}^{\Delta X+L} \int_{\Delta Y}^{\Delta Y+L} e^{-\phi_c(x-u+y-v)} dy\, dx\, dv\, du & i \neq k\,,\ j \neq l\ (\textit{diagonal}) \end{cases}$$

$$P(R \geq h^-) = \begin{cases} \dfrac{8}{L^4} \displaystyle\int_0^{L/2} \int_0^u \left(\begin{array}{l} \displaystyle\int_0^{u-v} \int_0^{v+u-x} e^{-\phi_c(u-x)} dy\, dx + \int_0^u \int_{v+u-x}^L e^{-\phi_c(y-v)} dy\, dx + \\[2mm] 2\displaystyle\int_{u-v}^u \int_0^{v-(u-x)} e^{-\phi_c(v-y)} dy\, dx + 4\int_{u-v}^u \int_{v-(u-x)}^v e^{-\phi_c(u-x)} dy\, dx + \\[2mm] \displaystyle\int_u^L \int_{v+(x-u)}^L e^{-\phi_c(y-v)} dy\, dx + \int_{u+v}^L \int_0^{v+(x-u)} e^{-\phi_c(x-u)} dy\, dx \end{array} \right) dv\, du & |i-k| = |j-l| = 0 \\[14mm]

\dfrac{1}{L^4} \displaystyle\int_0^L \int_0^L \int_{x=\Delta X}^{\Delta X+L} \int_{y=\Delta Y}^{\Delta Y+L} e^{-\phi_c(x-u)} dy\, dx\, dv\, du & |i-k| > |j-l| \\[8mm]

\dfrac{1}{L^4} \displaystyle\int_0^L \int_0^L \int_{x=\Delta X}^{\Delta X+L} \int_{y=\Delta Y}^{\Delta Y+L} e^{-\phi_c(y-v)} dy\, dx\, dv\, du & |i-k| < |j-l| \\[8mm]

\dfrac{2}{L^4} \displaystyle\int_0^L \int_0^u \left(\begin{array}{l} \displaystyle\int_{\Delta X}^{\Delta X+u-v} \int_{\Delta Y}^{\Delta Y+L} e^{-\phi_c(y-v)} dy\, dx + \\[2mm] \displaystyle\int_{\Delta X+u-v}^{\Delta X+L} \int_{\Delta Y}^{v+x-u} e^{-\phi_c(x-u)} dy\, dx + \\[2mm] \displaystyle\int_{\Delta X+u-v}^{\Delta X+L} \int_{v+x-u}^{\Delta Y+L} e^{-\phi_c(y-v)} dy\, dx \end{array} \right) dv\, du & |i-k| = |j-l| \neq 0 \end{cases}$$

$$(9.4)$$

$$P(R \geq h^+) = \begin{cases} \dfrac{4}{\phi_c^4 L^4}\left(1 - 2e^{-\phi_c L} + e^{-2\phi_c L} - 2\phi_c L + 2\phi_c L e^{-\phi_c L} + \phi_c^2 L^2\right) & i=k,\ j=l\ (same\ cell) \\[2em] \dfrac{2e^{-\phi_c \Delta X}}{\phi_c^4 L^4}\left(3 - 3e^{-\phi_c L} + e^{-2\phi_c L} - e^{\phi_c L} - 2\phi_c L + \phi_c L e^{\phi_c L} + \phi_c L e^{-\phi_c L}\right) & i=k,\ j \neq l\ (same\ row) \\[2em] \dfrac{2e^{-\phi_c \Delta Y}}{\phi_c^4 L^4}\left(3 - 3e^{-\phi_c L} + e^{-2\phi_c L} - e^{\phi_c L} - 2\phi_c L + \phi_c L e^{\phi_c L} + \phi_c L e^{-\phi_c L}\right) & i \neq k,\ j = l\ (same\ col.) \\[2em] \dfrac{e^{-\phi_c(\Delta X + \Delta Y)}}{\phi_c^4 L^4}\left(6 - 4e^{-\phi_c L} + e^{-2\phi_c L} - 4e^{\phi_c L} + e^{2\phi_c L}\right) & i \neq k,\ j \neq l\ (diagonal) \end{cases}$$

$$P(R \geq h^-) = \begin{cases} \dfrac{4}{L^4 \phi_c^4}\left(6 - 6e^{-\phi_c L} - 6\phi_c L + 2\phi_c^2 L^2 + \phi_c^2 L^2 e^{-\phi_c L}\right) & |i-k|=|j-l|=0 \\[2em] \dfrac{1}{L^2 \phi_c^2}\left(e^{-\phi_c(\Delta X + L)} + e^{-\phi_c(\Delta X - L)} - 2e^{-\phi_c \Delta X}\right) & |i-k|>|j-l| \\[2em] \dfrac{1}{L^2 \phi_c^2}\left(e^{-\phi_c(\Delta Y + L)} + e^{-\phi_c(\Delta Y - L)} - 2e^{-\phi_c \Delta Y}\right) & |i-k|<|j-l| \\[2em] \dfrac{2e^{-\phi_c L(\Delta+1)}}{L^4 \phi_c^4}\left(-3 - \phi_c^2 L^2 e^{\phi_c L} + \phi_c^2 L^2 - 6\phi_c L e^{\phi_c L} + 3e^{2\phi_c L}\right) & |i-k|=|j-l| \neq 0 \end{cases}$$

$$(9.5)$$

Figure 9.7 shows an example of the accuracy achieved by the approximation given in Eq. 9.5 for typical values of ϕ_c and L. The lighter shade values give the averaged probability and the darker points give the probability using either the upper or the lower estimate (note that shorter distances give higher probabilities). There is scatter about the $45°$ line due to the different combinations of pixels on a rectangular grid. There is greater variability at probabilities closer to zero since these occur at larger distances and there are more combinations of pixels which can give the same distance. However, this is not clearly shown in Figure 9.7 due to having a smaller region with only the pixels in opposite diagonals being sufficiently distant to give correlations close to zero. The averaged estimated has a mean error of 1.5% and a worst observed error for this example at 3.5%.
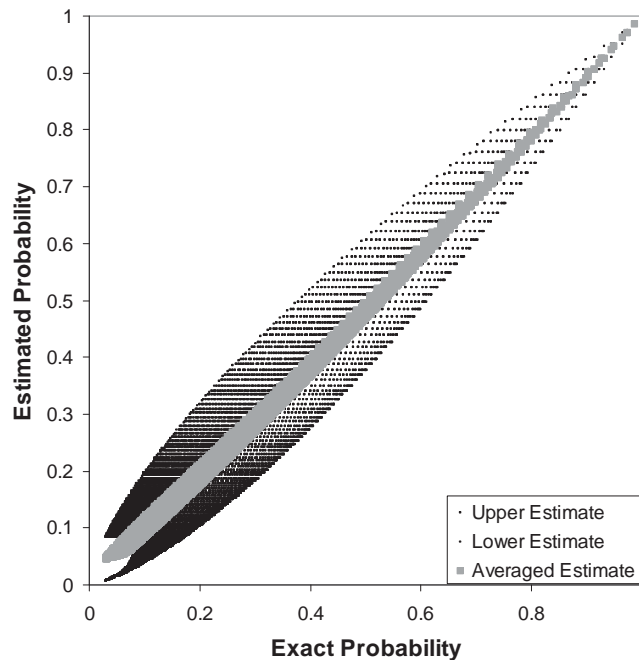
Figure 9.7   Effectiveness of hypotenuse approximations for various distances.

9.3.1   *Temporal process*

Consider a region where storms arrive with rate λ with uniform probability over that region. As outlined in Chapter 7, the storm covers a circular area given by the parameter ϕ_s. The storm gives rise to a number of cells which have properties of intensity, duration and areal extent given by the parameters α, η and ϕ_c. The simplest extension of the model is to fix these parameters as constant over the region and allow only the number of cells to vary over the region. There are several reasons for limiting the extension to the number of cells in preference to other parameters:

- The cells can still be conceptually thought of as a disc of rain. This is not the case if the raincell properties (lifetime, intensity, start time) are changed in space.

- Varying multiple parameters across a region makes the calibration methodology impractical.

- Cells have a smaller spatial scale than storms, thus they allow for steeper gradients in the rainfall statistics as compared to making the storm process inhomogeneous.

The inhomogeneous model is conceptualised on a two dimensional lattice where each pixel has a specified spatial rate, φ. The spatial rate gives the number of cells from a Poisson distribution that originate inside a pixel and the coordinates of each cell are located with uniform probability inside the pixel. The model is inhomogeneous because the spatial rate is made to vary in each cell, thus some regions will have more cells than others. In this way the model mimics overall inhomogeneities within the observed data (e.g. 'when it rains it always rains more in the hilly region').

The average number of cells μ_c overlapping a pixel D_{kl} is given by the rate of cells in each originating pixel $\varphi_{i,j}$ multiplied by the average probability those cells extend to D_{kl},

$$\mu_c^{(k,l)} = \sum_{i=1}^{m} \sum_{j=1}^{n} \varphi_{i,j} P(R > ||(i,j) - (k,l)||) \tag{9.6}$$

where the lattice domain is rectangular with m rows and n columns and the derivation of the overlap probability is given in Eq. 9.5.

It is possible to model the region of interest using a finite lattice because the cell radii decay exponentially such that the probability of a cell overlapping at large distances is negligible. The lattice domain will however need to be bigger than the region of interest by including a buffer to avoid boundary effects. If the values of $\mu_c^{(k,l)}$ are known for each pixel then the $\varphi_{i,j}$ parameters can be determined by solving the inverse problem obtained from Eq. 9.6. The parameters $\varphi_{i,j}$ from each pixel form the underlying intensity surface of raincells and the spatial locations of cells can be simulated.

Average pixel-wise model properties can be written by substituting the parameter μ_c which is the number of cells overlapping a point with $\mu_c^{(k,l)}$ which is the average number of cells overlapping the pixel domain $D_{(k,l)}$. This substitution is valid for the temporal statistics of the $SNSRP$ model since all other parameters are constant across the region and since the spatial variation is captured solely by the different $\mu_c^{(k,l)}$ for each pixel. Specifically, these equations are the mean (Eq. 3.3), the variance and auto-covariance (Eq. 3.4), the skewness (Eq. 3.5), and the probability of a dry interval (Eq. 3.8). A direct substitution is not as straightforward for the spatial cross-correlation.

The effect of inhomogeneity in the raincell process implies that the spatial correlation will differ for every pair of pixels. That is, for a given distance d between any two points the homogeneous model will give only one correlation, but the inhomogeneous model will give many different correlations. The correlation property of the homogeneous model appears as a single decaying line, whereas the same property of the inhomogeneous model has inherent scatter. To determine the covariance between two points it is necessary to consider three separate cases that cause rain to occur at the two points (i) two independent

cells from two independent storms, (ii) two independent cells from a common storm and (iii) a common cell from a common storm. Regarding case (i), the independence of cells implies that the contributions are pixelwise and can be obtained by substituting $\mu_c^{(i,j)}\mu_c^{(k,l)}$ for the term μ_c^2. Regarding case (ii) it is only the cell process that is inhomogeneous, not the storm process. For this reason, the number of storms overlapping a point, μ_s, is constant at all points, hence the derivation of this case remains unchanged. Since the cells cover the two points are independent the substitution of $\mu_c^{(i,j)}\mu_c^{(k,l)}$ for the term μ_c^2 is valid. Regarding case (iii) a single cell is large enough to cover the two points and the resulting term is location dependent. It is not enough to consider the number of cells overlapping D_{ij} and D_{kl} separately as was possible for the previous cases. The probability that a cell originating in some pixel domain D_{gh} jointly overlaps pixels D_{ij} and D_{kl} needs to be evaluated. This term is evaluated as

$$\sum_{a=1}^{m}\sum_{b=1}^{n}\lambda E[X^2]\frac{1}{\eta}e^{-\eta\tau}\varphi_{a,b}\,min(P(R>||(i,j)-(a,b)||),P(R>||(k,l)-(a,b)||)),$$

(9.7)

where the summation considers the effect of the cell originating in any given pixel, the minimum condition allows for the probability that the raincell overlaps the furthermost pixel (as it implicitly will overlap the closer pixel) and the other terms are unchanged from the existing derivation allowing for the cell intensity and duration. This equation has been presented as a discrete summation so that the efficient evaluation of the overlap probabilities can be exploited. In a proportionally small number of cases, if the two pixels are of similar distance to the pixel that the raincell originates, there is a slight approximation due to it assuming that the probability of overlap is solely governed by one pixel. This approximation will cause a slight underestimation of the cross-correlation but the bias is not significant for pixels having a small length scale ($1km$) when compared to typical distances between gauges. The cross-covariance term therefore becomes

$$\gamma_{h,\tau}(D_{i,j},D_{k,l}) = \lambda(\mu_x\beta\mu_c)^2 A(h,\tau)/(\eta^3(\beta^2-\eta^2)) - \lambda(\mu_x\mu_c)^2 B(h,\tau)/(\beta(\beta^2-\eta^2)),$$

$$+2\lambda E[X^2]A(h,\tau)/\eta^3\sum_{a=1}^{m}\sum_{b=1}^{n}\varphi_{a,b}\,min(P(R>||(i,j)-(a,b)||),$$
$$P(R>||(k,l)-(a,b)||)),$$

(9.8)

where $A(h,\tau)$ and $B(h,\tau)$ are as defined previously in Eq. 3.4. The cross-correlation is obtained from the cross-covariance by dividing by the standard deviation at each pixel.

Table 9.1 Computational requirements for matrix inversion for symmetric decompositions on an $n \times n$ square lattice. Based on *Gneiting et al.* [2006].

	Matrix Factorization (flops)	Inversion (flops)	Storage (real values)
Standard Cholesky	$\frac{1}{3}n^6$	n^4	$\frac{1}{2}n^4$
Block Banded Matrix[a]	n^4m^2	n^3m	n^3m
Block Toeplitz Matrix[b]	$6n^5$	n^4	$2n^3$
Block Circulant Matrix[c]	$40n^2\log_2 2n$	$80n^2\log_2 2n$	$8n^2$

[a]The Block Banded Matrix refers to limiting the process to a $(2m+1) \times (2m+1)$ neighbourhood [*Rue*, 2001]
[b]The Block Toeplitz Matrix algorithm as discussed by *Deitrich* [1993]
[c]The Block Circulant Matrix as discussed by *Deitrich and Newsam* [1997]

9.3.2 Matrix Inversion to Find Spatial Rate Surface

The inhomogeneous model arises from an underlying spatial rate that is variable according to some surface. Determining this surface however is an inverse problem. This inverse problem arises from Eq. 9.6 when it is written in a matrix format. A weighting matrix specifies the probability that a cell arising at a given location will overlap a location at some distance away. This matrix must be inverted to obtain the underlying spatial-rate surface.

Even for moderate sized grids the weighting matrix becomes too large for traditional numerical techniques that calculate inverse matrices, such as Gauss-Jordan and Gauss-Siedel methods, are inappropriate. Of the conventional techniques, computing the Cholesky Decomposition of a matrix in order to obtain the inverse is widely regarded as the most efficient method (it is twice as quick as LU decomposition). A Cholesky decomposition is a symmetric decomposition of the form $W = AA^T$ where the matrix A is lower-triangular. Despite the relative merit of conventional algorithms for computing the Cholesky factorisation, inspecting 9.1 reveals that standard Cholesky techniques are highly restrictive on the size of matrices they can handle. For example, consider a $500\ km \times 500\ km$ region to be modelled using a lattice with $1\ km$ intervals in both directions (n=500). As the weighting matrix, W, requires the influence of every point (n^2) with every other point (n^2), the required computer memory to store W is n^4. At 4 bytes per number this approximately 250 Gigabytes. The computational time is also prohibitive as it is proportional to n^6 floating-point operations.

It is however possible to invert a matrix more efficiently if it has a structure that can be exploited. These structures will be discussed with respect to the matrix W by-way-of a one dimensional example. The two-dimensional cases extend easily but are not as conceptually clear because of a 'block' effect, which gives rise to tiered/nested patterns in the

matrix W. For example, block-banding, block-Toeplitz, block-circulant effects. Consider then a one-dimensional example given in Figure 9.8, where there are n grid-cells, each separated by a distance δx and having an area $A = \delta x^2$. The i^{th} grid-cell has a spatial location x_i and there is an average rate of cell occurences over the grid-cell, $\bar{\varphi}(x_i)$. For convenience, consider the distance between cells to be one unit, $\delta_x = 1$, $A = 1$.
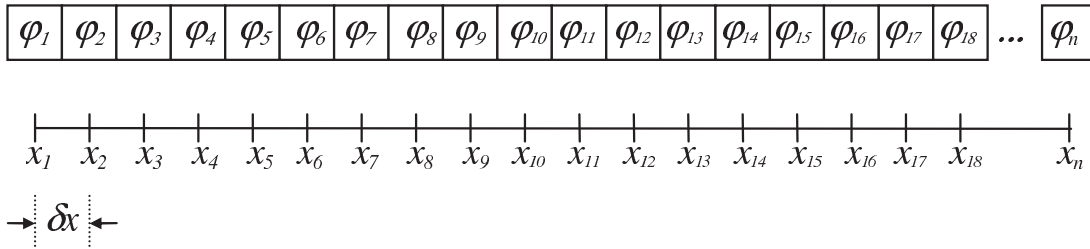

Figure 9.8   1-D line example of spatial cell process.

Consider an exponential distribution such that the overlap probability of a rain-cell from the centre of grid-cell i covering the centre of grid-cell j is given by $w_{i,j}(\tau) = e^{-\phi_c \tau}$, where $\tau = |i - j|$ is the absolute distance between the grid-cells. Letting $\rho = e^{-\phi_c}$ it can be seen that $w_{i,j}(\tau) = \rho^\tau$. This gives rise to the structure of W in Eq. 9.9 where the weight values decrease exponentially with distance away from the diagonal (equivalent to an exponential correlation matrix). Note that effects of spatial averaging over the pixel domain, as outlined in the previous section, have been ignored. This averaging does not change the fundamental structures of the matrices, and it is conceptually easier to think of the matrix as an exponential correlation matrix.

$$
W = \begin{bmatrix}
1 & \rho & \rho^2 & \dots & \rho^n \\
\rho & 1 & \rho & \dots & \rho^{n-1} \\
\rho^2 & \rho & 1 & \dots & \rho^{n-2} \\
\vdots & & & \ddots & \vdots \\
\rho^n & \rho^{n-1} & \rho^{n-2} & \dots & 1
\end{bmatrix}
\tag{9.9}
$$

Observing that after some distance m the weight values will have negligible contribution, a reasonable approximation would allow values beyond m lags to be equal to zero. This is equivalent to a truncated correlation matrix or Markov property where the influence on a grid-cell is determined only by grid-cells within a neighbourhood of size m relative to that cell. The resulting sparse matrix will have a banded structure that can be exploited to give improved factorisation algorithms. *Rue* [2001] discusses the

two-dimensional case which gives rise to a block-banded structure that has an improved performance over the standard Cholesky decomposition as shown in Table 9.1. The improvement primarily relies on the neighbourhood distance being significantly smaller that the size of the entire grid.

An alternative observation that can be made about Eq. 9.9 is that the diagonal terms have identical values. This is referred to as a Toeplitz matrix. *Deitrich* [1993] discuss an algorithm that exploits block-Toeplitz structures for the two-dimensional case, with relative computational improvement summarised in Table 9.1. Whilst this improvement is beneficial it can be improved by noting that a Toeplitz matrix can be turned into a circulant matrix via the process of circulant-embedding. Circulant matrices are diagonalised by the Fourier transform and by exploiting this the significant reductions in computational time and memory storage outlined in Table 9.1 are observed.

Figure 9.9 shows an extension of the line process so that it is reflected about the first grid-cell. The first cell is not duplicated and to preserve a length that is a multiple of two (for the FFT algorithm) an additional grid-cell is inserted on the opposite side of the circle, labelled ϵ. The value given to ϵ is arbitrary and for convenience can be taken as zero.
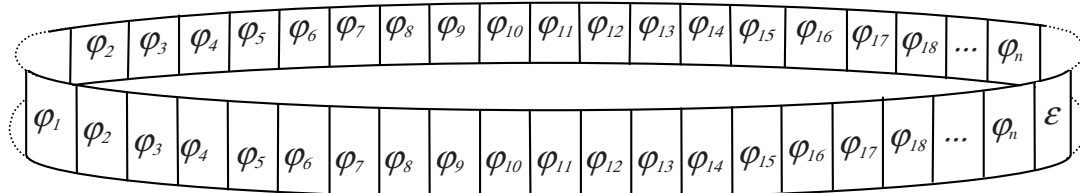

Figure 9.9 Circle schematic of spatial cell process extended from 1-D line example for explaining Circulant matrix structure.

The matrix corresponding to the circle in Figure 9.9 is computed in a similar fashion to Eq. 9.9, noting that there are twice as many grid-cells. Inspecting the first row which corresponds to weight values contributed to grid-cell x_1, the weight values decay from 1 to ρ^n corresponding to the behaviour of the process on the line $x_1 \dots x_n$. The weight values then begin to increase back to ρ as a revolution of the circle is completed back to the comparison of two points x_1 and x_2 separated by only one distance (recall that ϵ is an arbitrary value). Beginning this process from any other point on the circle gives exactly the same sequence of correlations such that each row of matrix C in Eq. 9.10 is generated as a 'right-shift' operation on the row above. This is the definition of a circulant matrix,

where it is completely characterised by the first row and all other rows are obtained as an offset. The top-left quadrant Eq. 9.10 is the matrix of the original line-process, the bottom-left quadrant is the matrix of the reflected portion of the line and the top-right and bottom-left quadrants are the weights between the original and reflected part. The Toeplitz matrix of the original line process, W, is said to be embedded in the circulant matrix, C, as beginning on any main diagonal element of C it is possible to extract an $n \times n$ matrix that is the matrix W. For a two-dimensional process a surface is reflected in both directions and embedded into a torus which creates a matrix that is block-circulant.

$$
C = \left[\begin{array}{cccc|cccc}
1 & \rho & \dots & \rho^n & \epsilon & \rho^n & \dots & \rho \\
\rho & 1 & \dots & \rho^{n-1} & \rho^n & \epsilon & \dots & \rho^2 \\
\vdots & & \ddots & \vdots & \vdots & & \ddots & \vdots \\
\rho^n & \rho^{n-1} & \dots & 1 & \rho & \rho^2 & \dots & \epsilon \\
\hline
\epsilon & \rho^n & \dots & \rho & 1 & \rho & \dots & \rho^n \\
\rho^n & \epsilon & \dots & \rho^2 & \rho & 1 & \dots & \rho^{n-1} \\
\vdots & & \ddots & \vdots & \vdots & & \ddots & \vdots \\
\rho & \rho^2 & \dots & \epsilon & \rho^n & \rho^{n-1} & \dots & 1
\end{array} \right]
\tag{9.10}
$$

A special property of circulant matrices is that they can be diagonalised by Fourier transforms. Moreover, the eigenvalues of a circulant matrix C can also be determined using a discrete Fourier transform of the first row of that matrix, denoted $C_{1,:}$. Denoting the discrete Fourier transform of a vector x as a product with the Fourier transformation matrix F, the diagonal matrix of eigenvalues is obtained as $\Delta_C = diag(FC_{1,:})$. The Fourier transformation matrix is given by *Kreyszig* [2005, pp. 525] as,

$$
F_n = \frac{1}{\sqrt{n}} \left[\begin{array}{ccccc}
1 & 1 & 1 & \dots & 1 \\
1 & \omega^1 & \omega^2 & \dots & \omega^{n-1} \\
1 & \omega^2 & \omega^4 & \dots & \omega^{2(n-1)} \\
\vdots & \vdots & \vdots & & \vdots \\
1 & \omega^{n-1} & \omega^{2(n-1)} & \dots & \omega^{(n-1)(n-1)}
\end{array} \right]
\tag{9.11}
$$

where $\omega = e^{\frac{-2\pi}{n}i}$, and has the property $F_n^{-1} = F_n^H$, that the inverse transformation is equal to the conjugate transpose (i.e. the adjoint or Hermitian transpose). Correspondingly, the circulant matrix can be decomposed as,

$$
C = F^H \Delta_C F.
\tag{9.12}
$$

This decomposition can be exploited to give highly efficient means for computing either the Cholesky decomposition or the inverse of the circulant matrix. While the circu-

lant embedding technique requires the size of the grid to be doubled in both dimensions, the ability to exploit the FFT algorithm significantly compensates for this. Furthermore, it is only necessary to store the first row of the circulant matrix in memory for the calculation of the eigenvalues, which is significantly less than storing the entire matrix W. For example a grid size of 512×512 cells, requires approximately 10 Megabytes and is 1×10^7 times quicker as compared to a standard Cholesky factorisation.

Despite this, the utility of circulant matrices is not readily obvious for inverting the Toeplitz matrix W which is embedded inside the circulant matrix. While the inverse C^{-1} is easy to determine, the relationship between C^{-1} and W^{-1} is non-trivial. *Chan and Ng* [1996] observe that in many instances iterative solutions for the inversion of Toeplitz matrices can be more efficient than direct methods. They provide a comprehensive review of conjugate gradient techniques for inverting Toeplitz systems including one method that uses circulant embedding [*Chan and Ng*, 1996, §1.3]. This method is outlined below.

For the system of linear equations $m = Ws$, the conjugate gradient algorithm is given below. Recall that W is a large matrix storing the correlation between each pair of pixels, m is the known mean in each pixel and s is the spatial rate. The conjugate gradient technique is advantageous for finding the vector s since it does not require the matrix W to be stored (only products with W that result in vectors need to be stored). The algorithm is outlined in Table 9.2 below.

Table 9.2   Conjugate Gradient Algorithm.

$k = 0$	(iteration counter)
$s_0 = 1$	(initial estimate solution vector)
$r_0 = m - Ws_0$	(residual vector)
$p_0 = r_0$	(conjugate gradient vector)
loop	
$\quad \alpha_k = r_k^T r_k / (p_k^T W p_k)$	(scalar coefficient for conjugate vector)
$\quad s_{k+1} = s_k + \alpha_k p_k$	(solution vector update)
$\quad r_{k+1} = r_k + \alpha_k W p_k$	(residual vector update)
$\quad$ **if** $r_k < tol$ **exit loop**	(exit criteria)
$\quad \beta_k = r_{k+1}^T r_{k+1} / (r_k^T r_k)$	(scalar coefficient for vector adjustment)
$\quad p_{k+1} = p_k + \beta_k p_k$	(conjugate gradient vector update)
$\quad k = k + 1$	
end loop	

From Table 9.2 it can be seen that there are several steps which involve the matrix W. The product Ws_0 must be computed once to determine the initial vector of residuals

and the product Wp_k is evaluated once each iteration as it can be stored and reused. To determine these products, consider the following equation

$$\begin{bmatrix} \mathbf{W} & \mathbf{B} \\ \mathbf{B} & \mathbf{W} \end{bmatrix} \begin{bmatrix} \mathbf{p_k} \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{Wp_k} \\ * \end{bmatrix} \tag{9.13}$$

where B is some matrix that satisfies the circulant property and $*$ represents values of no resulting interest. From Eq. 9.12 the circulant matrix can be replaced with its Fourier decomposition. The eigenvalues for this decomposition can be computed once and stored. Thus the product in Eq. 9.13 which is required each iteration can be achieved with two FFTs. This technique is highly computationally efficient and only requires a small amount of memory storage. The number of iterations required by the conjugate gradient technique can be made to be small by using a preconditioning matrix such as Strang's preconditioner. *Chan and Ng* [1996] discuss several different preconditioning techniques.

By specifying the problem as a matrix inversion it is not possible to enforce the constraint that the spatial rate is positive at all locations. An example of where this may not occur is if the surface of the number of cells (specified by μ_c) has steep changes in value over a short distance. The inversion may compensate for this by causing some grid-cells to have negative values in the neighbourhood close to a grid-cell having a high spatial rate. It is important to note that the cell radius parameter ϕ_c has a large influence over this as it controls the values in the matrix W. As the cell radii follow an exponential distribution, for a quantile Q of this distribution, the weight value will be significant at the distance $d = -\frac{1}{\phi_c}ln(1 - Q)$. For a typical value, $\phi_c = 0.1$ and a high quantile $Q = 0.99$ the distance is approximately $50 \ km$, beyond which cells will have a low contribution. This length-scale gives a qualitative scale for the fluctuation of the μ_c surface.

9.4 SIMULATION TECHNIQUE

Chapter 4 highlighted the need to simulate cells outside of a target region in order to preserve statistics within that region. However, due to the varying nature of the spatial-rate surface it is not possible to analytically derive equations to efficiently simulate the cells landing outside of the target region and it is necessary to use a buffer instead. Chapter 4 also emphasized that if the region is large with respect to the buffer, then the computational overhead will be minimal. This is of benefit to the inhomogeneous model as it is typical that inhomogeneity will be significant over larger regions. Section 7.5 outlined a technique for extending the efficient simulation of cells to include the efficient simulation

of a storm envelope. For the inhomogeneous model the rate of storm arrivals is uniform over the region and it is possible to retain the efficient simulation algorithms for this aspect of the model. This is also beneficial as a length scale for storm extent ($\sim 100\ km$) is an order of magnitude larger than the spatial scale of rain-cells ($\sim 10\ km$). Thus using a buffer for the storm-extent would require a region many times larger than the target region and would be computationally prohibitive. As a result of these factors, the most efficient technique uses a combination of simulation regions, as shown below in Figure 9.10.
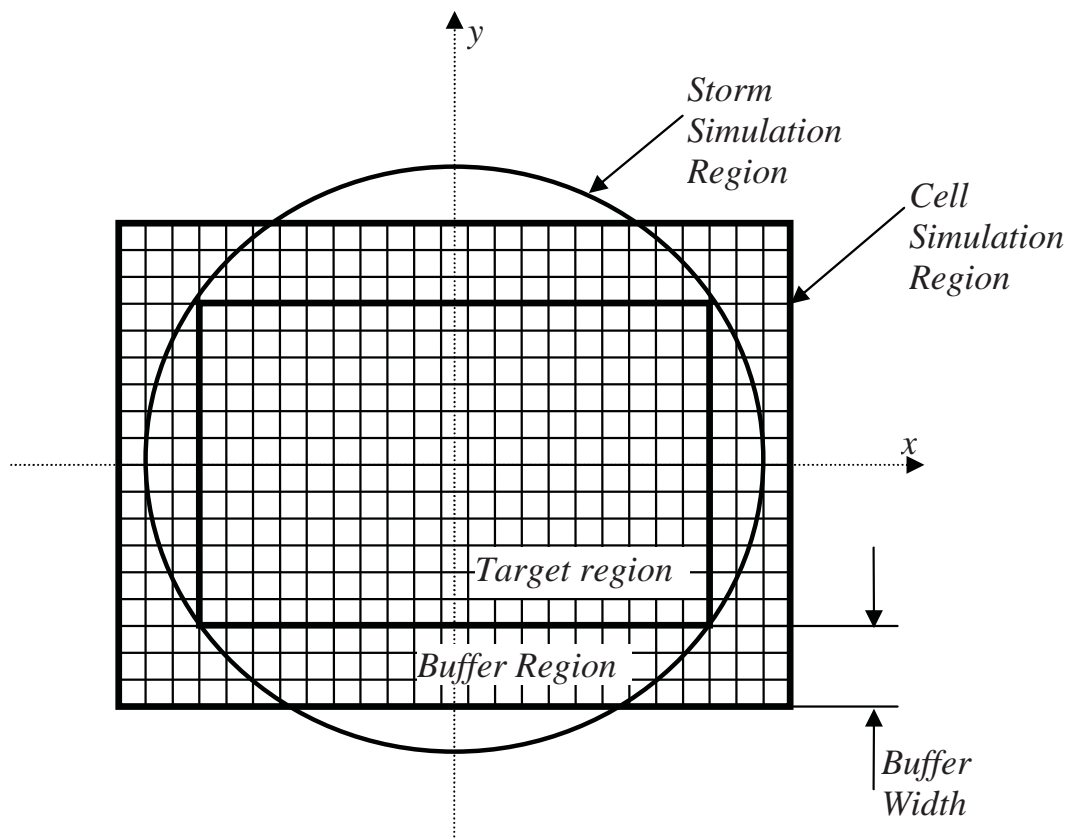
Figure 9.10   Schematic of simulation regions for inhomogeneous model.

Figure 9.10 shows a circular region overlaying the target region for simulating storm origins and areal extent. The reason for this region being circular is so that it can take advantage of the efficient simulation algorithm in Section 4 when simulating the storm properties. The circle needs to be sufficiently large to cover all points within the target region, otherwise the statistics within the target region will not be preserved. Note that there is a slight inefficiency in this arrangement as it is possible for storms to be simulated

to overlap a point within the circular region that does not overlap the target region. Given the arrival of a storm, the cell properties are simulated according to a rectangular gridded mesh over the target region with a surrounding buffer. Cells that land within the buffer but do not overlap a point within the target region are rejected. The size of the buffer is linked to the rain-cell parameter such that it is sufficiently large to include the majority of rain-cells that will influence the outmost grid-points (say the 0.99 quantile of the cell radius distribution). The statistics of points within the buffer region are not preserved. The use of a buffer implies that the modelled region should include grid-cells within the buffer region.
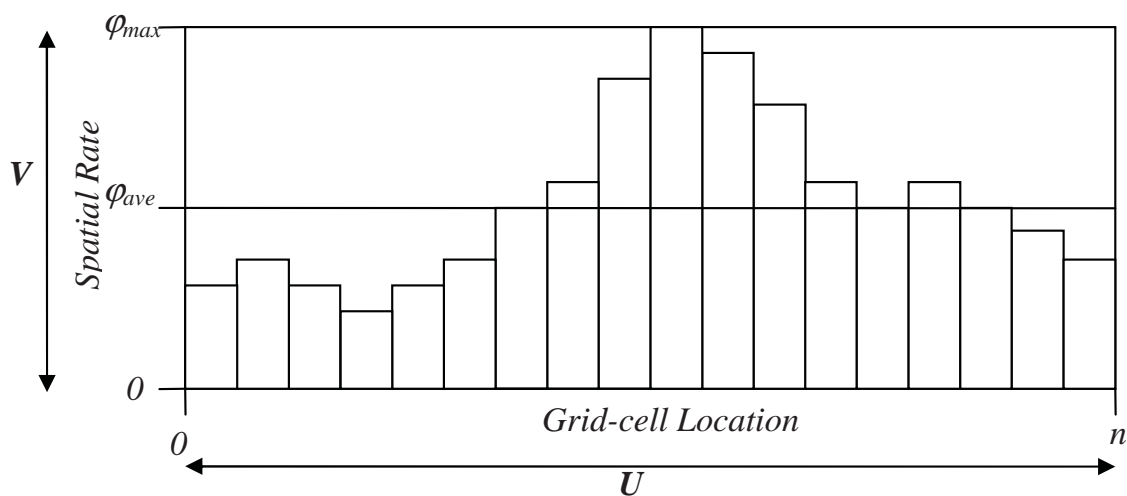
Figure 9.11   Accept-reject schematic for inhomogeneous model.

The simulation of the rain-cells according to a variable spatial-rate requires an accept-reject based technique, as shown for a one-dimensional example in Figure 9.11. The total number of cells is sampled using the average spatial rate φ_{ave}. For a given cell, a potential location for that cell is sampled uniformly over the domain. A separate random number is generated uniformly and scaled to the range of the maximum and minimum spatial rates over the region. At the proposed location, if the sampled random number is less than the spatial rate curve, the location is accepted, otherwise it is rejected and a new location is proposed. The efficiency of this algorithm is given by the proportion of area underneath the spatial-rate curve, which simplifies to $\varphi_{ave}/\varphi_{max}$. This relation also holds for a volume under a two-dimensional surface. For example, if the peak spatial-rate is 4 times higher than the average, the efficiency of the inhomogeneous model is 0.25 as

compared to the homogeneous model.

9.5   CASE STUDY

9.5.1   *Observed Data*

A total of 422 daily rainfall gauges and 10 pluviograph gauges were selected across the south-west region of western Australia for inclusion in this study. To match the study of *Jothityangkoon et al.* [2000], only those gauges with an unbroken record over the period 1974 to 1984 were chosen. The locations of these gauges are shown in Figure 9.12, where the large symbols are the pluviograph sites and the smaller symbols are daily sites. The bounding box on this region is provided so that subsequent plots of statistics can be easily inspected with respect to their latitude and longitude.
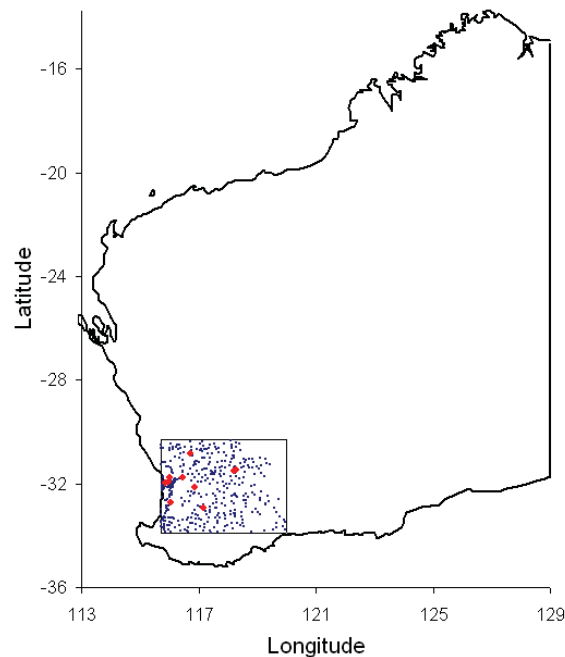
Figure 9.12   Location of gauges for the period 1974-1984 within the study region covering the Avon basin, Western Australia. Small points indicate daily rainfall gauges and larger points indicate pluviograph gauges.

Figure 9.13 shows a plot of the mean 1 hour rainfall depth over this region with respect to each month. The blue regions show areas of low rainfall and the red regions higher rainfall on a log z-axis. From these plots it is clear that there is a strong seasonal trend, with low rainfall during the summer months (December to January), higher rainfall during the winter months (May to July) and intermediate values in the transition periods

of Autumn and Spring. Rainfall gradients across this region are also evident. The Perth region between the coastline and the Darling escarpment (on the left hand side) shows consistently higher rainfall during all months. This is also where the highest density of gauges lie. Moving inland from this escarpment (visible by the sudden sparseness of gauges) there is a gentle trend of decreasing rainfall amounts for all months except during the summer period. During summer the rainfall amount appears to be relatively uniform over the region. There is also a subtle trend with higher rainfalls on the Southern coastline that is most noticeable during the spring months, September to November.

Figure 9.14 shows a plot of the coefficient of variation at the daily timescale with respect to each month. The blue regions show areas of low variability and the red regions higher variability. From these plots it is clear that there is again a strong seasonal trend, with more variable rain during the summer months and less variable rain during the winter months. During summer, the Perth region has consistently more variable rainfall than the inland region, but in winter this trend is reversed so that Perth's rainfall is less variable. There appears to be more 'scatter' in this statistic (meaning adjacent sites have dissimilar values) than when inspecting the mean rainfall values in the previous figure.

Figure 9.15 shows a plot of the portion of a dry days in a month. The blue regions show the portion of dry days near 0.5 and the red regions near 1.0. Similar trends to the previous two statistics can be observed. Summer is very dry across the entire basin, with less than 5% of days having a rainfall amount recorded. Winter shows a more consistent amount of rain. During winter there is also a strong rainfall gradient, with approximately half of the days being rainy in the south-western corner and two thirds being dry in the inland region. The intervening months, notably April and October, show the transition between the summer and winter regimes.

Figure 9.16 shows the daily lag-1 autocorrelation statistic across the region on a scale from 0.0 to 0.5. This statistic does not show as strong a trend as the previous statistics, although it is evident that the consistent winter rainfall in the south-western corner gives higher autocorrelations and that the inland region has consistently low correlations reflecting the sporadic occurrence of rain in this region.

The aim of the inhomogeneous model is to be able to reproduce these trends by varying only the rate of raincell occurrence across the region. Other parameters such as the intensity of the cells, the lifetime of cells and the occurrence of storms will remain constant.
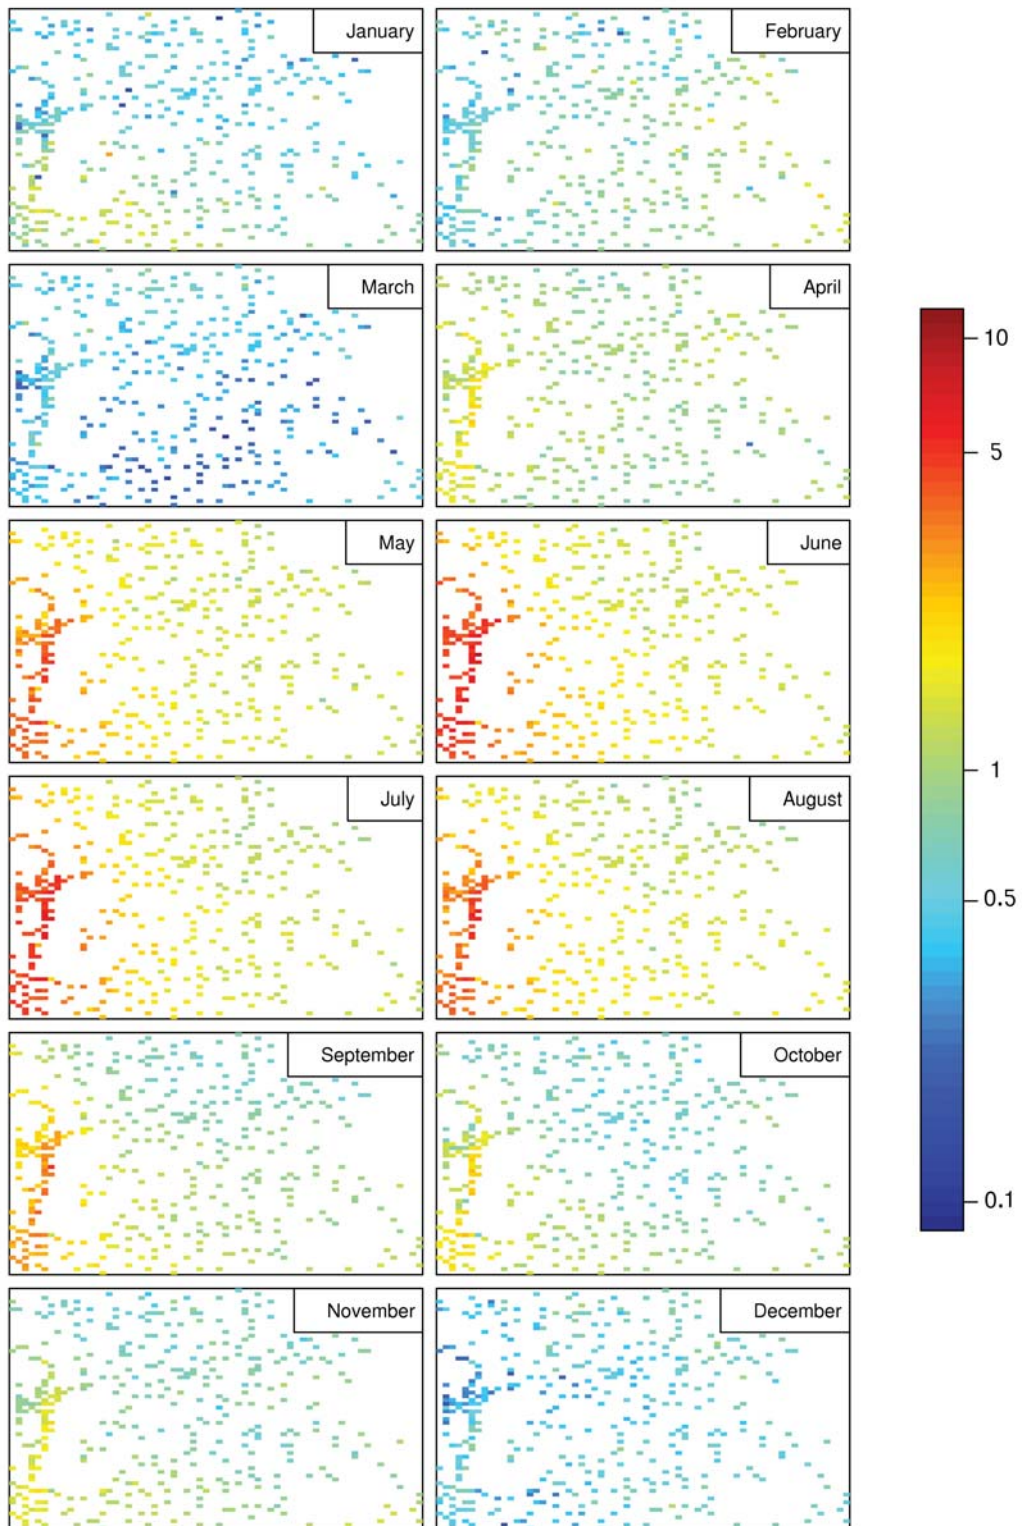
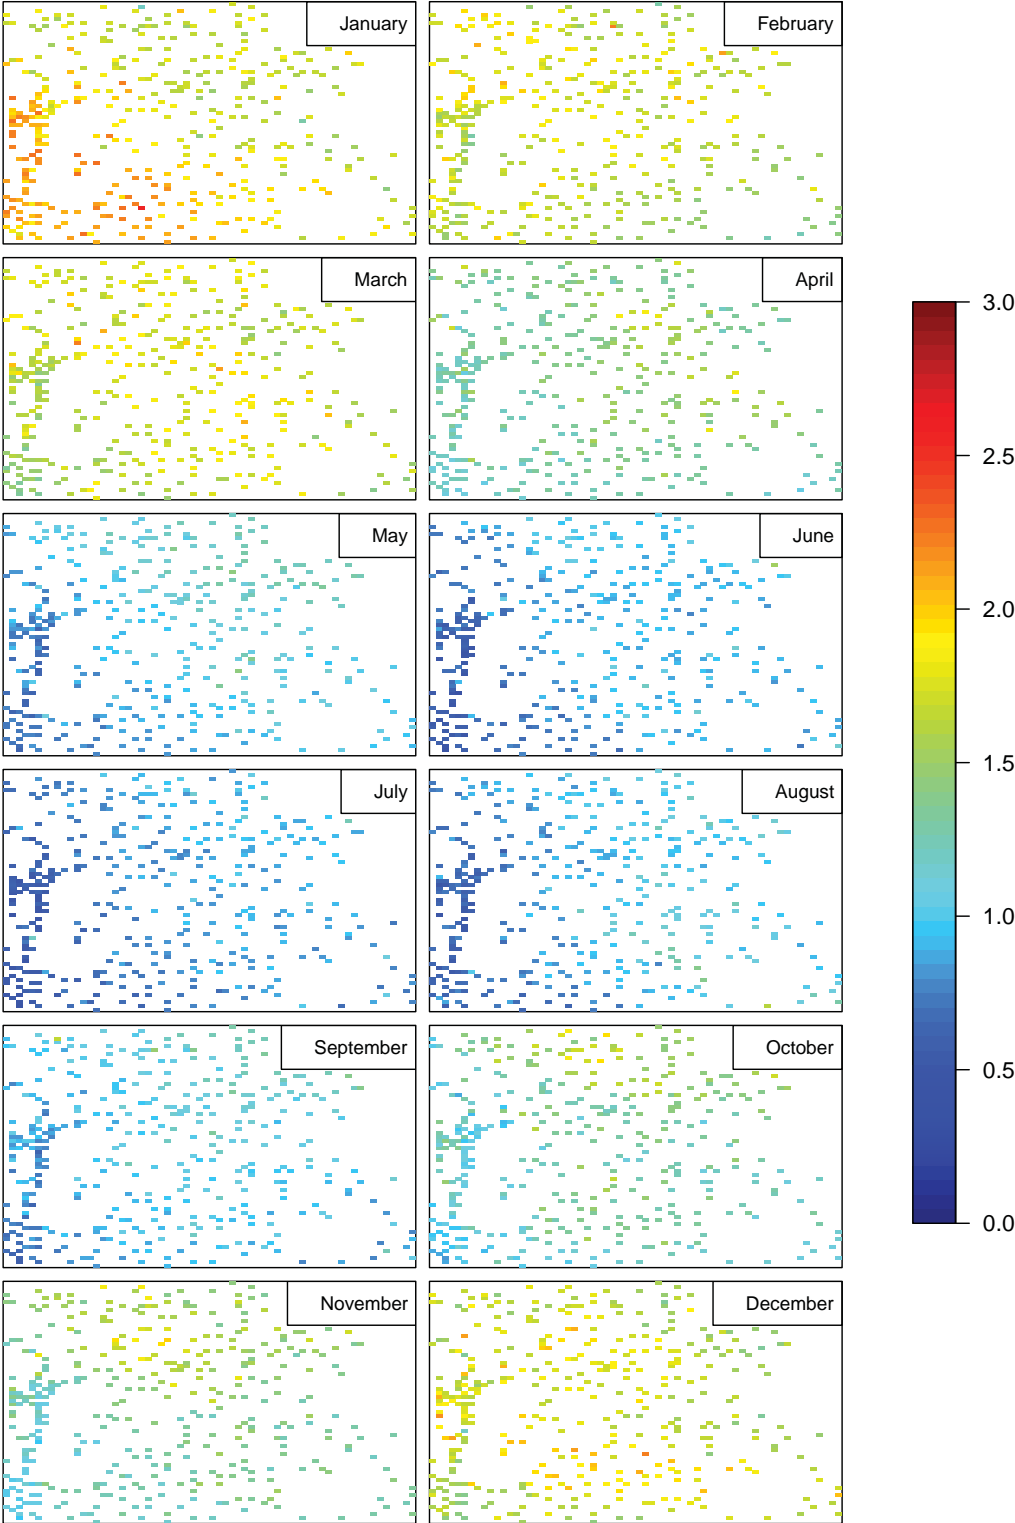**Figure 9.13** Variation in the mean rainfall depth at a 1 hour aggregate over the Avon Basin for 12 months of the year.

Figure 9.14   Variation in the coefficient of variation at a daily aggregate over the Avon Basin for 12 months of the year.
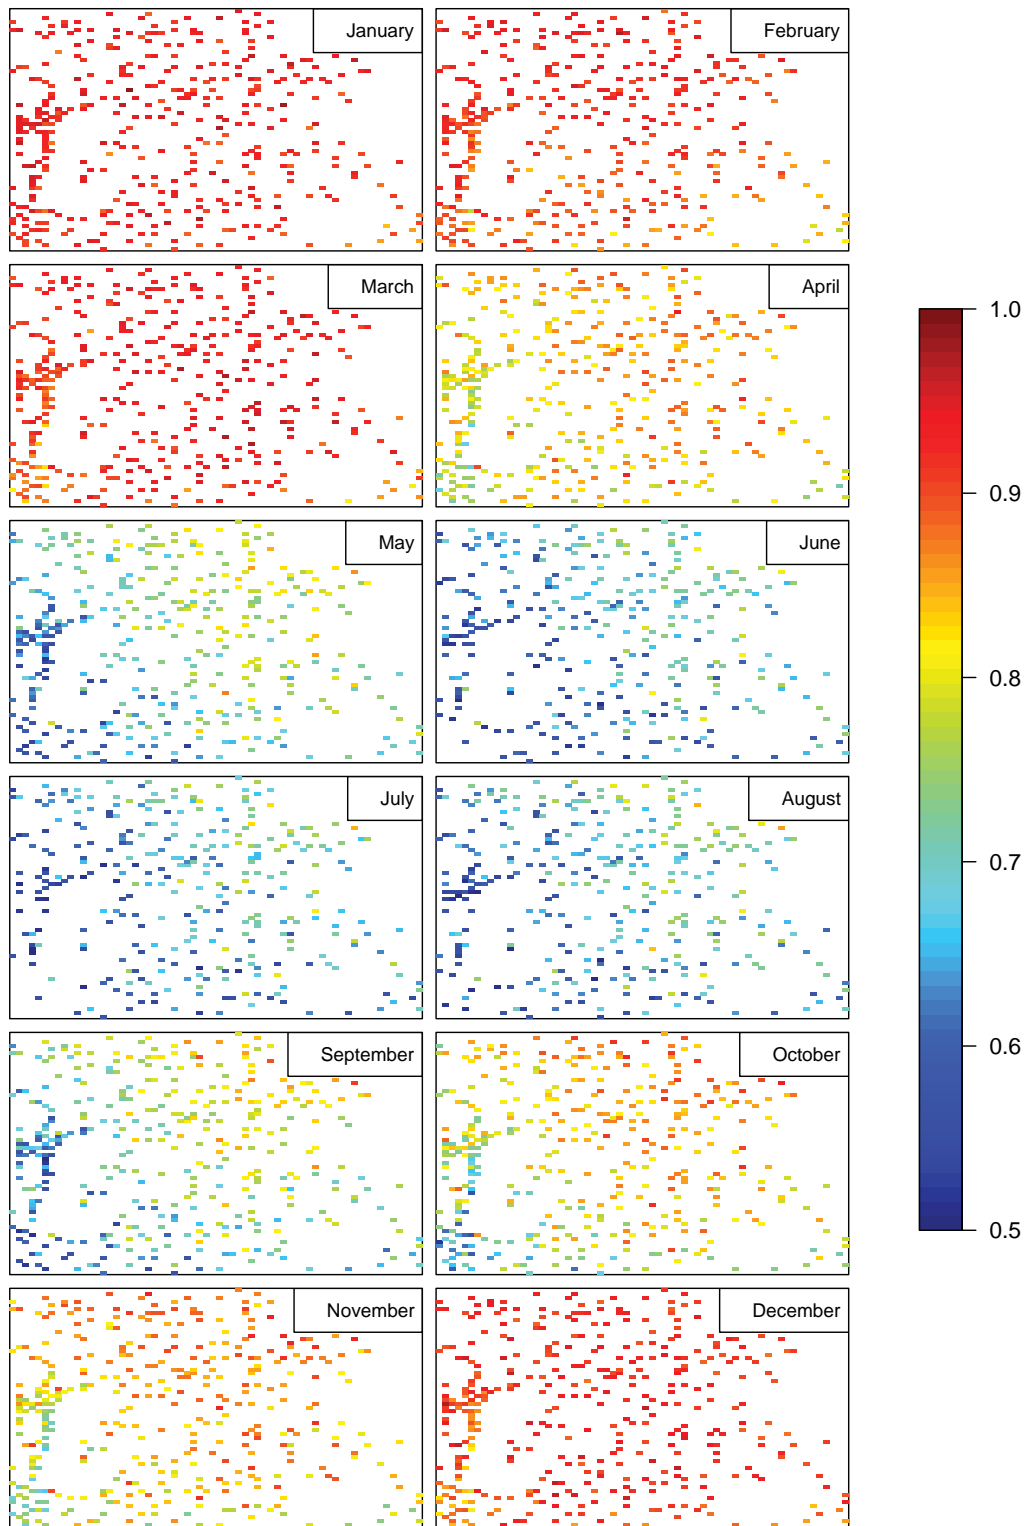
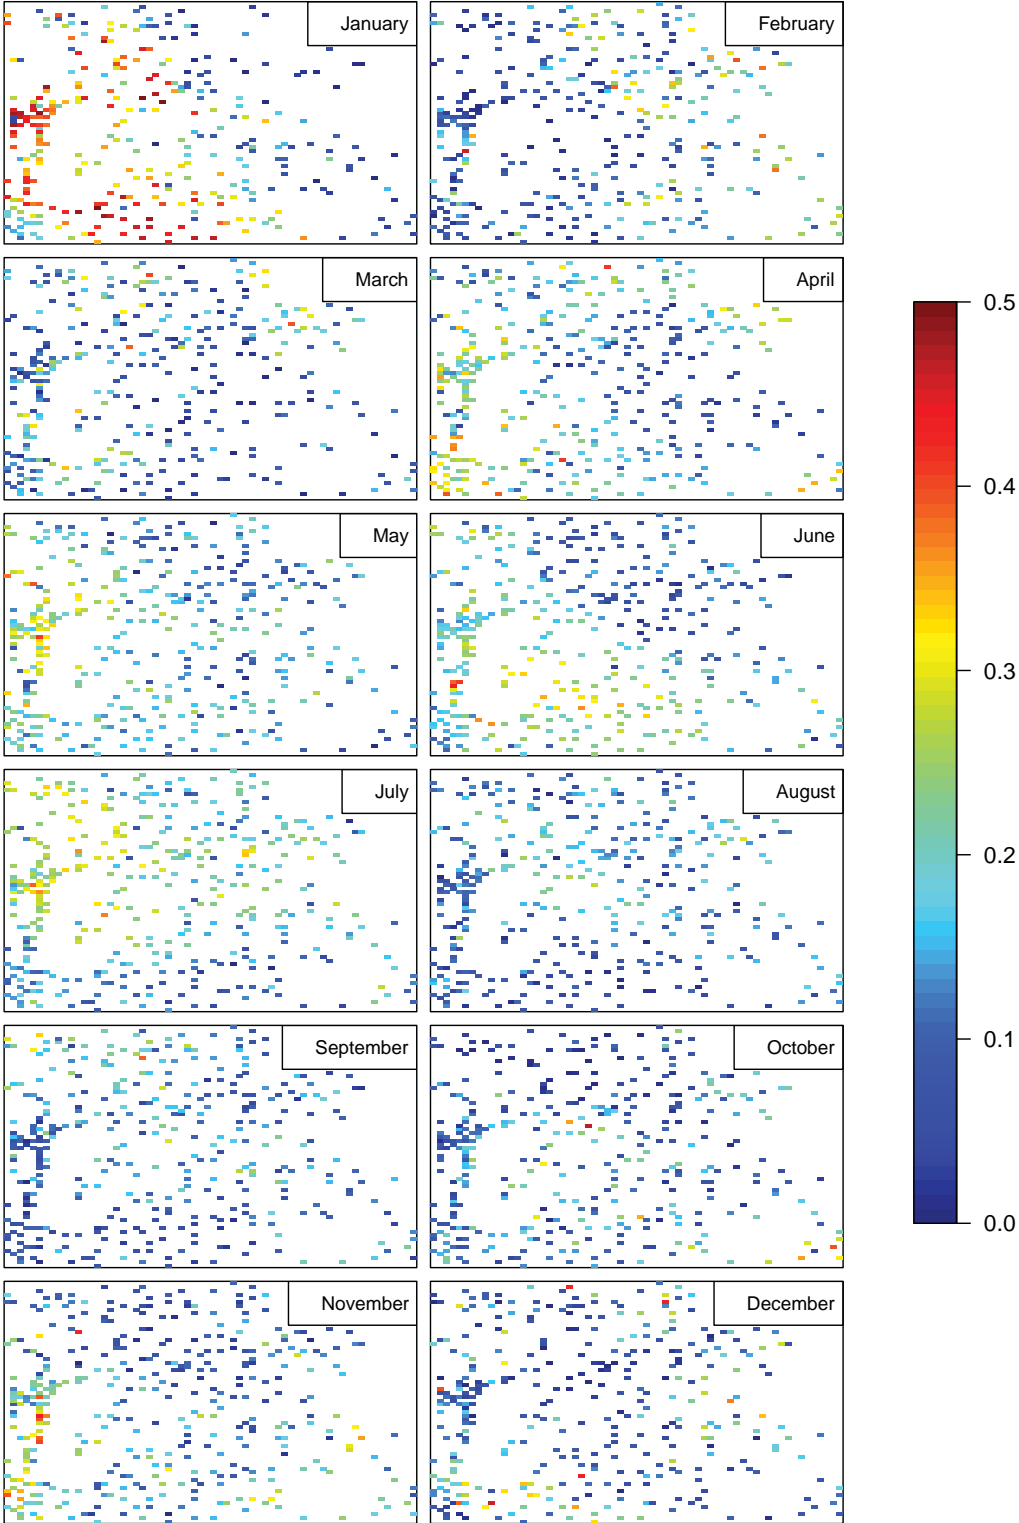**Figure 9.15** Variation in the probability of a dry day over the Avon Basin for 12 months of the year.

Figure 9.16 Variation in the daily auto-correlation over the Avon Basin for 12 months of the year.

9.5.2 Calibration to Regional Statistics

To calibrate the model the data was first pooled across the region and regional statistics were calculated. The procedure for this step is the same as in previous case studies for homogeneous regions. Only a single storm type was fitted otherwise it would be necessary to develop surfaces of raincell rates for each storm type, and it is not clear that there is enough information in the data to reliably do this.
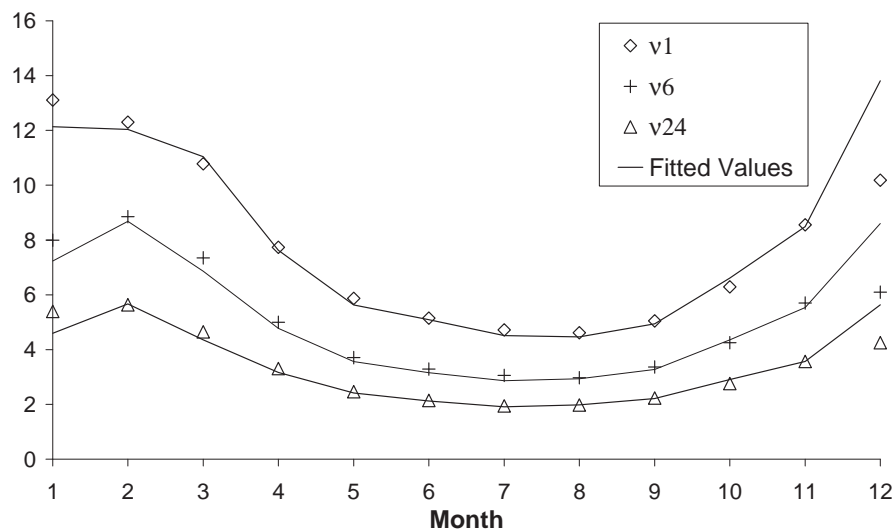

Figure 9.17   Fitted regional coefficient of variation, Avon Basin.
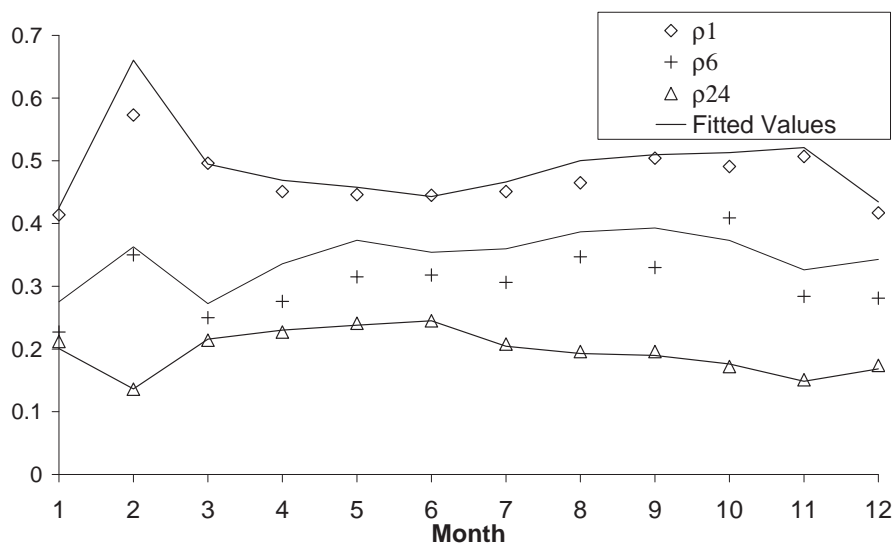

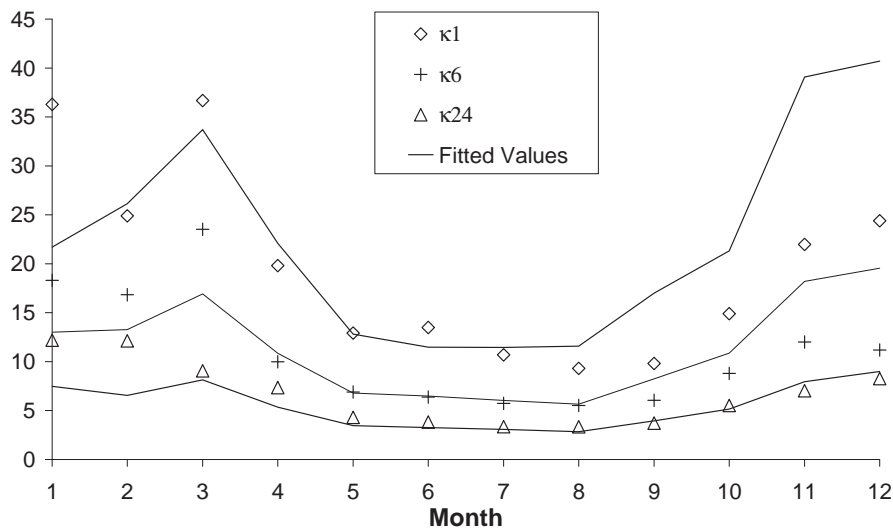**Figure 9.18**   Fitted regional autocorrelation, Avon Basin.

Figure 9.19   Fitted regional skewness, Avon Basin.

Figure 9.17 compares the fitted coefficient of variation to the observed regional values for each month, Figure 9.18 compares the fitted auto-correlation and Figure 9.19 compares the fitted skewness. As with previous case studies, the skewness statistic is the most variable and is the hardest to match. It should also be noted that there were significantly less pluviograph gauges than daily gauges in this region, so less weighting was given to the sub-daily statistics in the calibration. Table 9.3 reports the regional parameters for each month.

9.5.3   *Estimation of number of cells*

Having fitted the regional parameters, the model was again fitted at every site allowing only the parameter for the number of cells μ_C to be reestimated. All other parameters were held constant. The regional estimate of μ_C for each month was used as an initial estimate of this value for each gauge. A numerical complication was noticed during this procedure, that the optimiser would pursue unrealistic values of μ_C for trivial gains in the goodness-of-fit to the at-site statistics. To remedy this, an arbitrary constraint of 300 cells was placed on this parameter. Figure 9.20 shows the variation in the estimated μ_C parameter for each month and for each site. The dark-red sites are most likely sites where the constraint of 300 cells was reached. This constraint seems to be consistently reached in the south-western corner of the region. This suggests that this subregion is not well represented by the regional parameters that were held constant and that if some of the other parameters were allowed to vary, then more realistic and stable values of μ_C could

Table 9.3   Avon Basin regional parameter estimates for each month.

k	$\hat{\lambda}$ h^{-1}	$\hat{\beta}$ h^{-1}	$\hat{\eta}$ h^{-1}	$\hat{\mu}_C$ -	$\hat{\alpha}$ -	$\hat{\phi}_c$ km^{-1}	$\hat{\phi}_s$ km^{-1}
1	0.00058	0.077	1.5	32.8	0.809	0.13	0.01
2	0.00117	0.21	2.96	50.2	0.549	0.01	0.01
3	0.00188	0.183	3.22	20	0.575	0.3	0.01
4	0.00251	0.106	2.01	20.2	0.628	0.17	0.017
5	0.00523	0.121	2.51	36	0.568	0.15	0.01
6	0.00732	0.12	1.62	41	0.488	0.15	0.014
7	0.00823	0.11	2.17	73	0.433	0.2	0.01
8	0.00731	0.182	2.31	105.1	0.389	0.15	0.015
9	0.0058	0.165	3.66	44	0.493	0.17	0.02
10	0.00345	0.194	2.04	32	0.495	0.17	0.018
11	0.00293	0.124	3.4	28	0.478	0.16	0.013
12	0.00163	0.187	3.4	24	0.521	0.17	0.01

be achieved. As many of the gauges in this region have a high number of cells, it can be concluded that this area does have more frequent rainfall. Inspecting the months of April and October in Figure 9.20 there are numerous scattered sites that reached the constraint of 300 cells. It is unlikely that these sites truly experience more frequent rainfall, but that it is an artefact of the at-site statistics making it difficult for the optimiser to locate a reliable set of parameters. Comparing the spurious sites in Figure 9.20 for the months of April and October to the same months and sites in Figure 9.15, it seems that sites having a spuriously high number of cells correspond to the sites having an unusually low portion of dry days when compared to their immediate neighbours. Further checking of the data would be required to assess whether the discrepancy in the number of dry-days is genuine or an issue of data quality.

Having estimated the μ_C parameter at each gauge, the next step is to interpolate these values over the entire region. This is necessary because the simulation requires a spatial rate at every point in the region, and since the spatial rate must be determined from μ_C, μ_C must be estimated at every point in the region. Thin plate spines were used for interpolation where the covariance structure was estimated using a simple exponential model. More detailed modelling of the covariates such as distance from the coast and elevation was not considered as the main intention was to demonstrate the overall approach. If these covariates were included they would act to reduce the variability of the residuals between the estimated μ_C and the fitted surface at each gauge. The interpolated surfaces can be
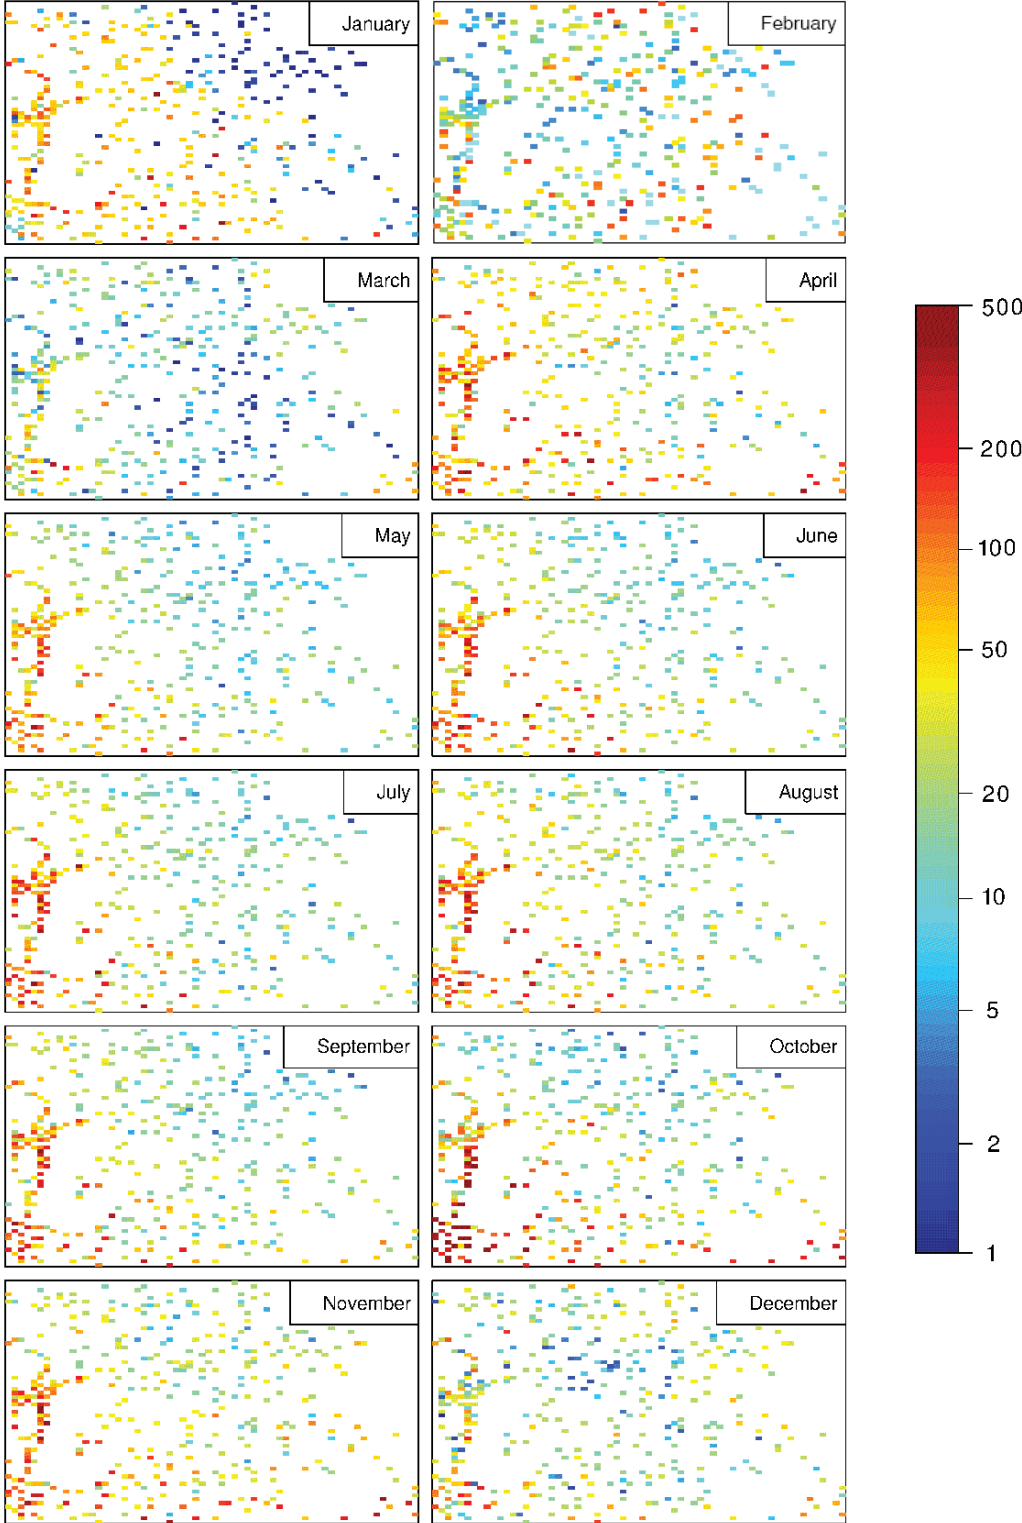
Figure 9.20   Variation in the number of cells fitted independently at each site over the Avon Basin for 12 months of the year.

seen in Figure 9.21. The boundary on the interpolated surfaces shows the boundary of the interpolation. It was also necessary to obtain estimated values outside of this boundary and these values were achieved via extrapolation. The reason for the extrapolation is that the simulation region requires a buffer zone. As a long section of the boundary to the region is coastline, extrapolation is necessary as there are no alternative estimates of rainfall.

The spline function smooths the surfaces so that the μ_C parameter is not as variable in space as the at-site estimates. Figure 9.22 shows a plot of the absolute values of the residuals from the interpolation procedure. The residuals have been standardised by diving by the average value of μ_C for each month. Some of the discrepancies are large which is a result of the high degree of smoothing in the interpolation. This reflects a conflict between the overall behaviour of the region (the need for a meaningful smooth surface) with the local observation of rainfall (subject to statistical variability).

9.6 RESULTS

Taking the estimates of μ_C as point estimates only, each gauge was simulated independently for a period of 11 years to assess the influence of smoothing. Figure 9.23 shows the simulated (y-axis) and observed portion of dry days (x-axis), where each symbol represents a different gauge. Agreement between the two would give values along the $45°$ diagonal line and note that a homogeneous model would give a horizontal line, the same value for all sites. It can be seen that before the spline, the winter months have many sites with good agreement (red symbols aligned on the diagonal), but that in the summer months there less agreement as not many points lie on the diagonal. This suggests that irrespective of any smoothing, the variation of μ_C parameter was not effective during the summer months and that other parameters should be varied (either in addition to, or in place of μ_C). The effect of the spline is represented by the blue symbols and when compared to the red symbols shows a higher degree of scatter. The artefact of a near-horizontal line in the lower portions for the blue symbols is a result of the smoothed surface removing the peaks of high numbers of cells (μ_C=300) observed at some gauges. Although less obvious, the same effect can be noticed for the higher dry portions due to increases in low values of μ_C at some sites as a result of smoothing.

Figure 9.24 provides a similar comparison between the simulated and observed daily coefficient of rainfall. Agreement is represented by the diagonal line. Overall, this statistic is not matched as well as the portion of dry days when allowing only the μ_C parameter to vary. The winter months tend to show more values that are scattered about the diagonal
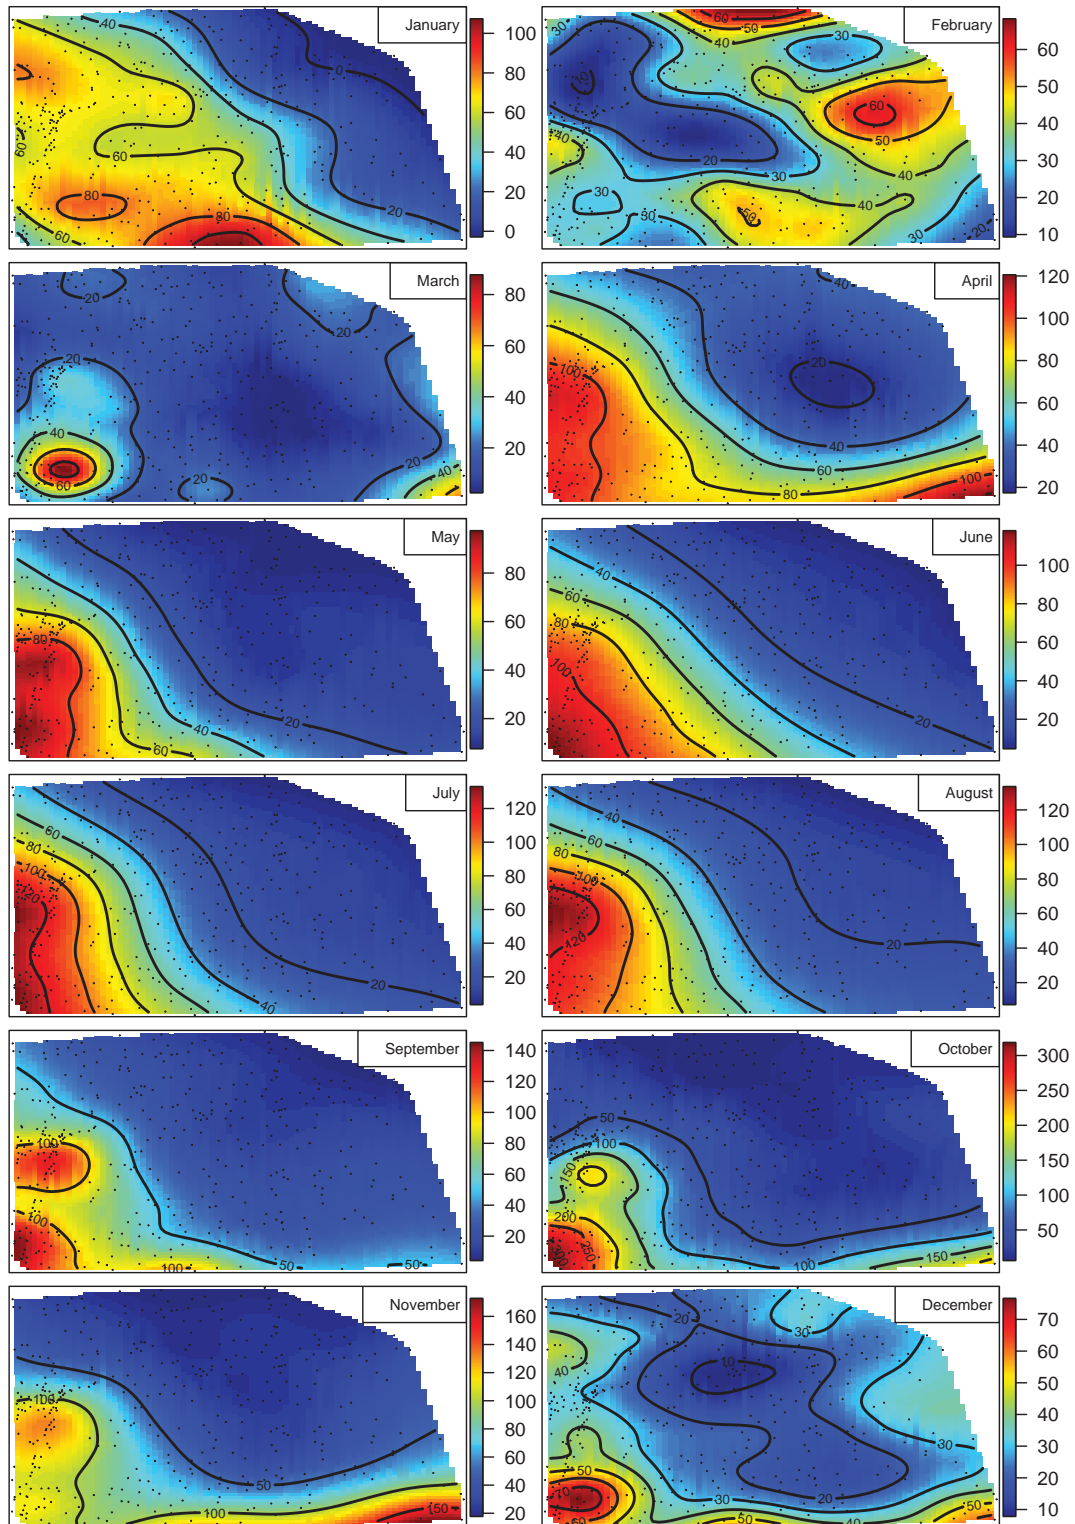
Figure 9.21 Fitted variation in the number rate of cells using thin-plate splines over the Avon Basin for 12 months of the year.
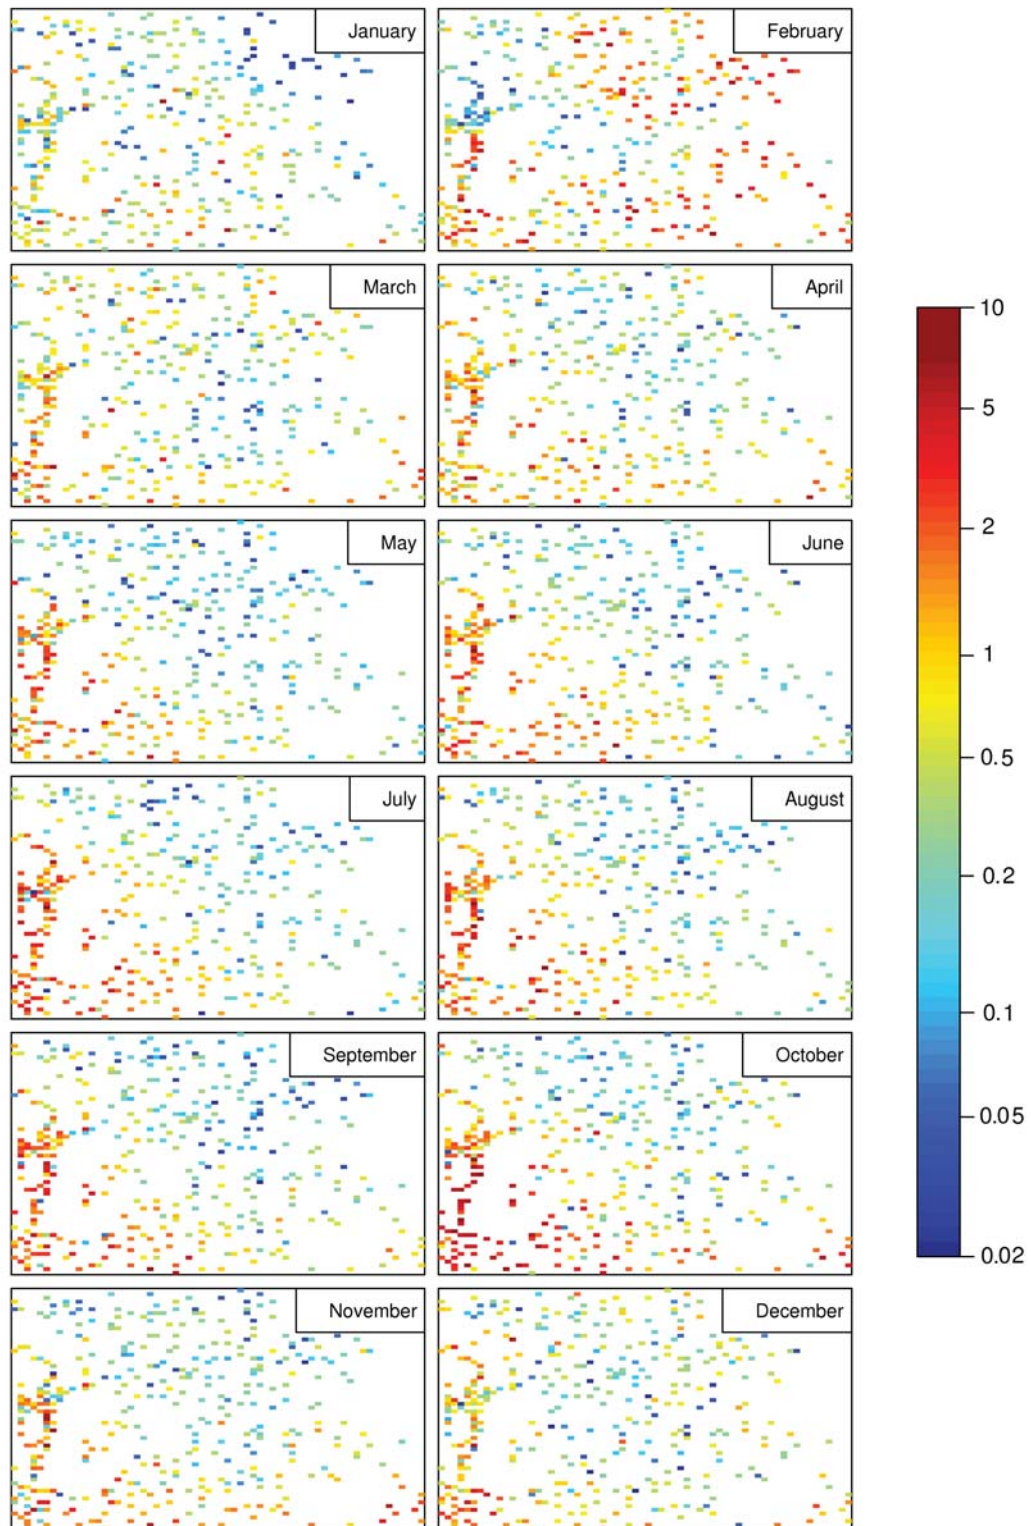
Figure 9.22   Absolute value of standardised residuals of the number of cells between the independent calibration and the thin-plate spline over the Avon Basin for the 12 months of the year.
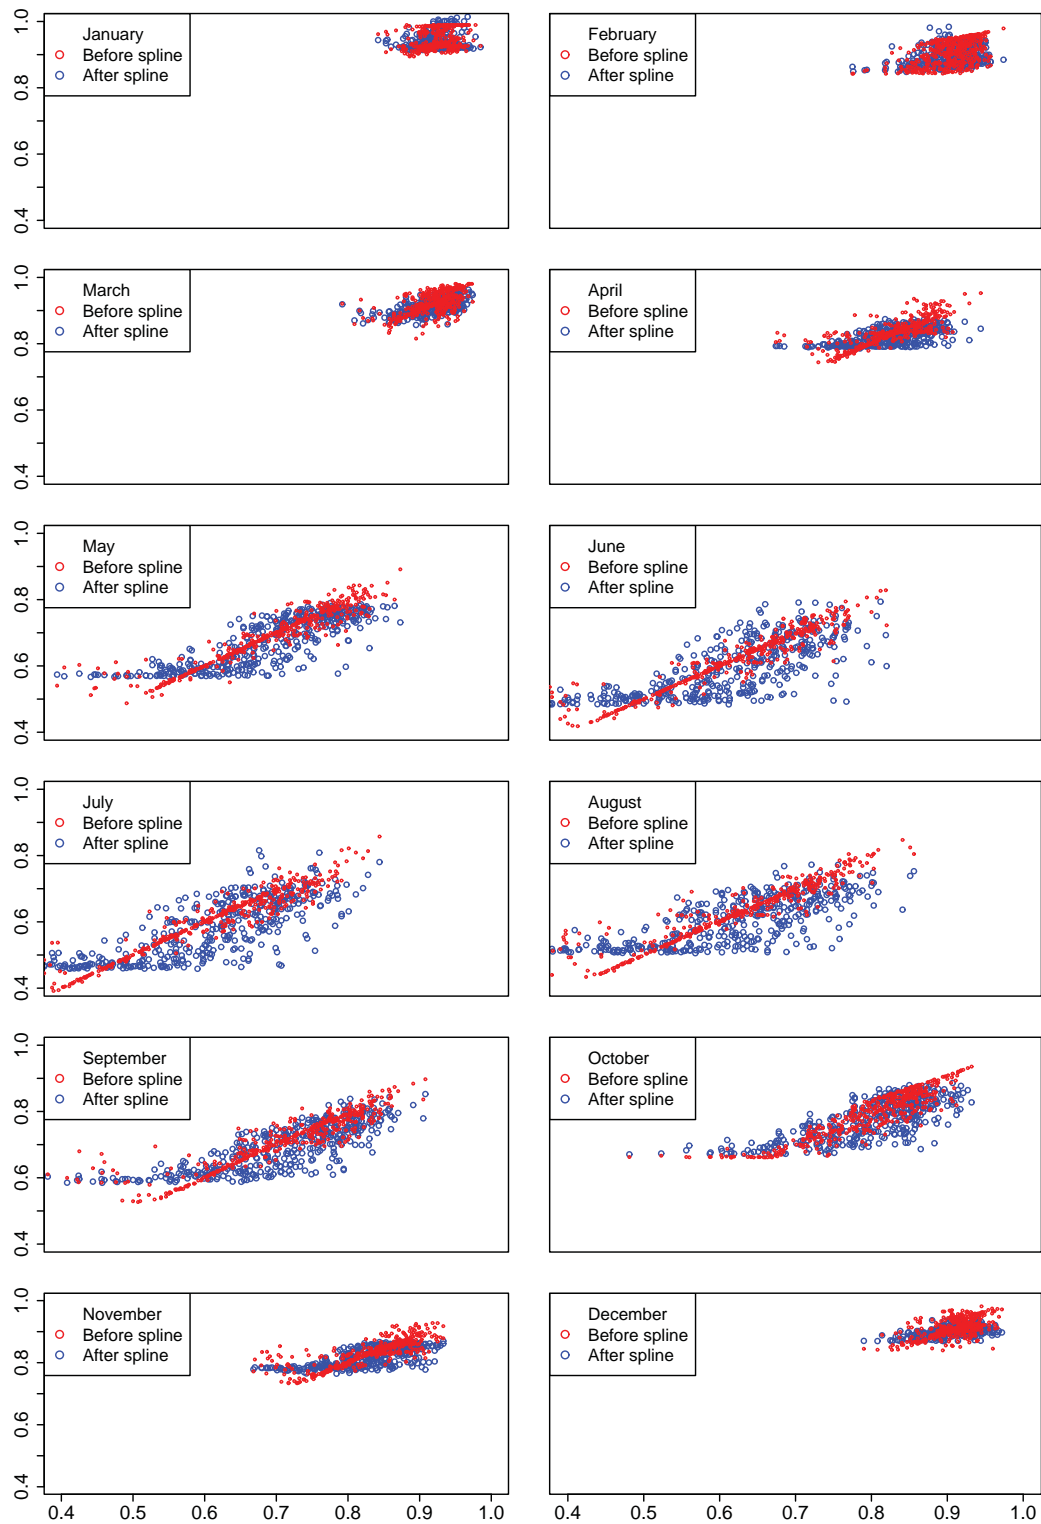
Figure 9.23 Comparison of simulated and observed daily dry probability for all points in the Avon Basin for the 12 months of the year.

line. By way of contrast, the summer months and January in particular show that the simulated coefficient of variation is nearly always the same despite variation in the observed value. Again, the spline increases the scatter in the simulated statistics so that there is less agreement with the observed estimates.

Simulating the inhomogeneous model requires the inverted surfaces of the spatial rate to be determined from the interpolated μ_C surface. The results of this inversion can be seen in Figure 9.25. The inhomogeneous simulation is very slow owing to the unavoidable need for a buffer region. The cross-correlation has been the main statistic used to assess the performance of the $SNSRP$ model. It is again used here to assess the effect of the inhomogeneous model. Figure 9.26 shows the uncertainty in correlation statistics at the 5% significance level when estimated from 330 datapoints (11 years, with 30 days per month). So, for example, a correlation that is truly 0.0, when estimated from 11 years of data for a given month, has a 95% chance of lying between ± 0.1. It is useful to remember this when comparing the simulations to observed data.

Figure 9.27 compares the simulated and observed daily cross correlation statistics for all pairs of gauges within the region. It can be seen that the inhomogeneous model provides a significant increase in the variability of the cross-correlation beyond what would be expected from the random variation about the homogeneous model. The observed daily correlations display even greater variability than what the inhomogeneous model is capable of producing, which suggests a limitation of the model. This limitation is most likely due to the smooth interpolation surface and also the use of only one parameter to vary rainfall estimates across the region. It is important to note that some of the scatter present in the observed cross correlations is suspect, specifically, the near-zero cross-correlations at short distances for the months of February, March, September, October, November and December.

9.7 CONCLUSIONS

This chapter has presented a methodology for implementing an inhomogeneous Neyman-Scott rainfall model. There are several complications to this methodology:

- The use of circular rainfall discs with a square grid for the spatial rate makes integration of the underlying equations difficult. The use of square raincells could simplify this slightly.

- The overlap of discs requires the solution of an inverse problem to determine the underlying spatial rate. This inversion is non-trivial for large and/or high resolution
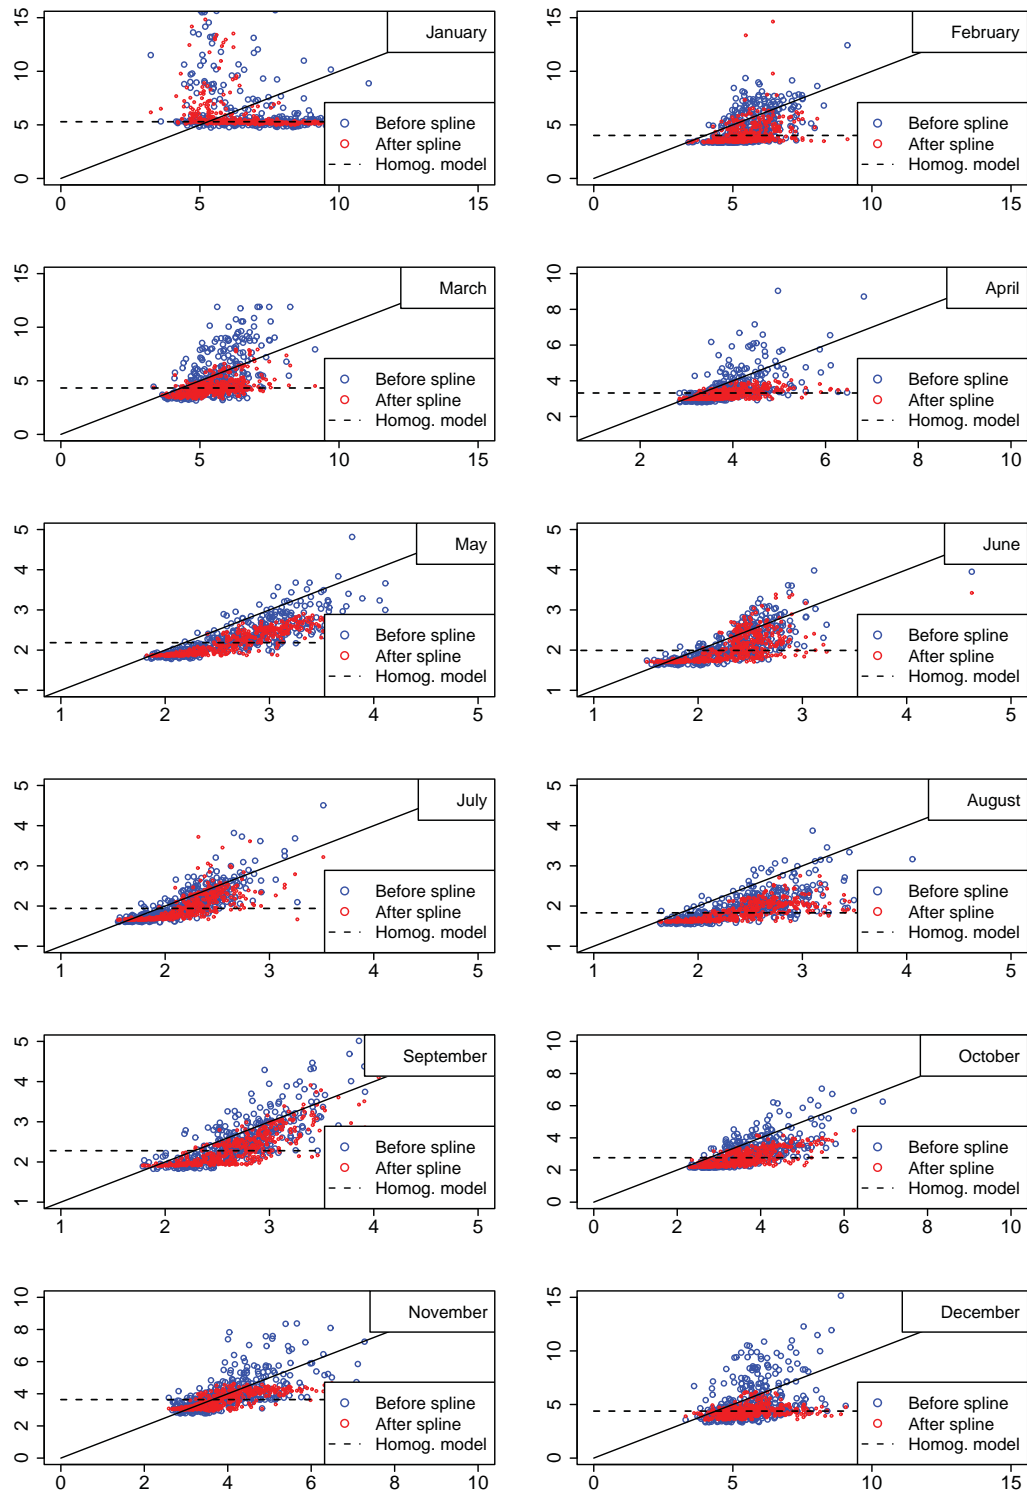
Figure 9.24 Comparison of simulated and observed daily coefficient of variation for all points in the Avon Basin for the 12 months of the year.
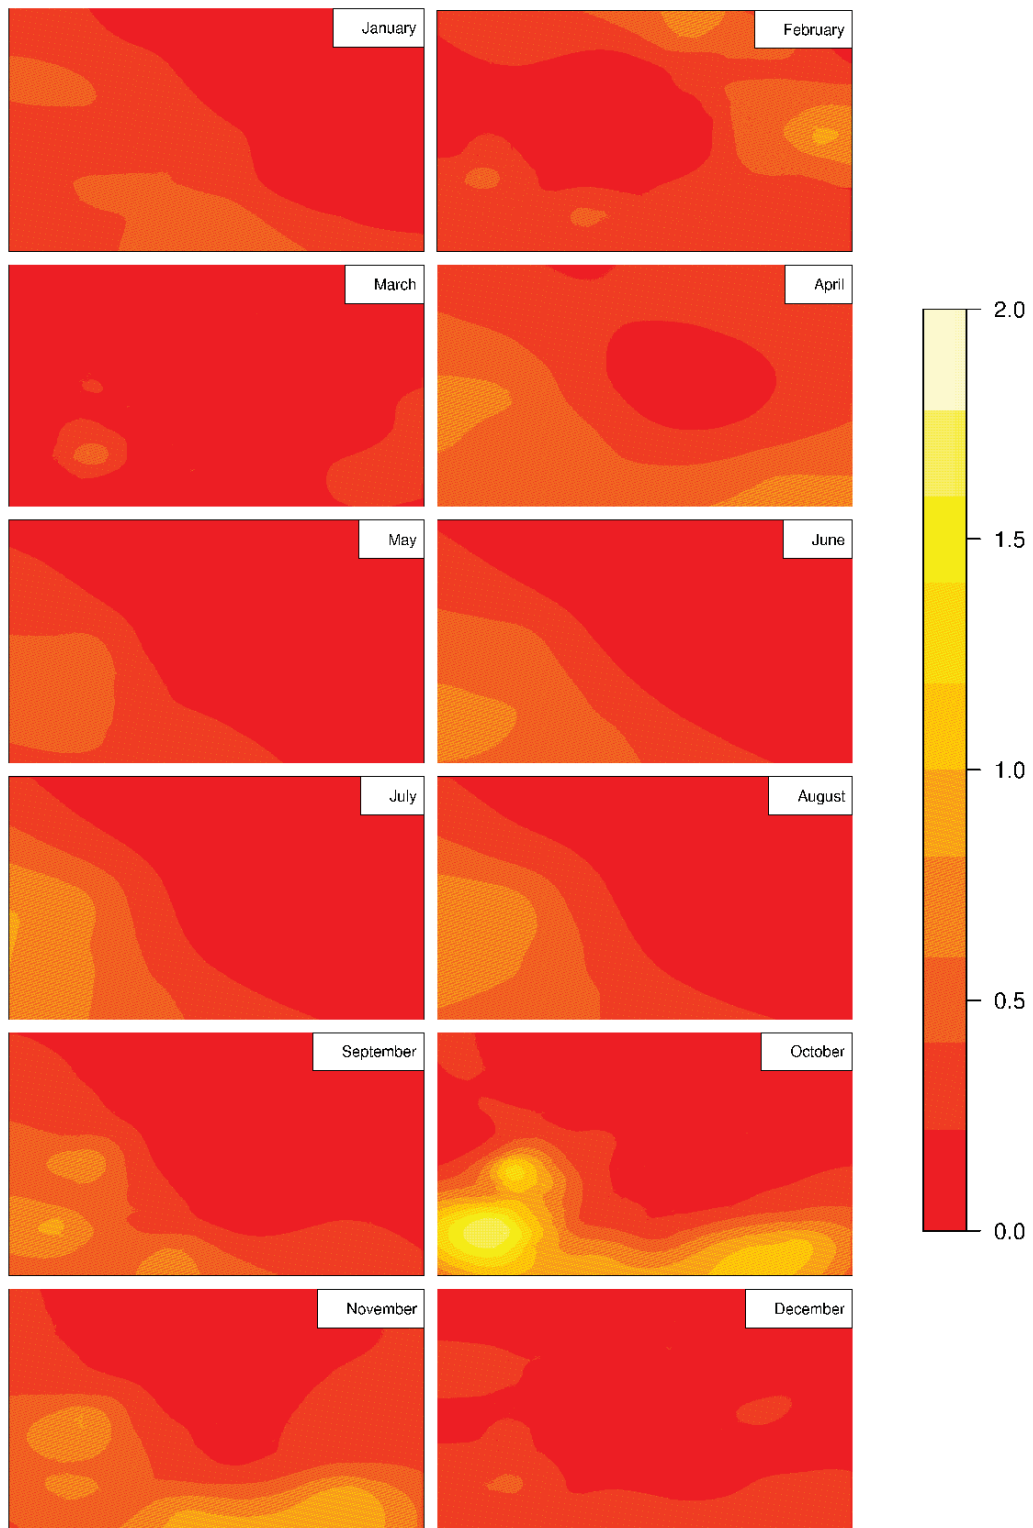
Figure 9.25   Simulated variation in the spatial rate of cells over the Avon Basin for 12 months of the year.
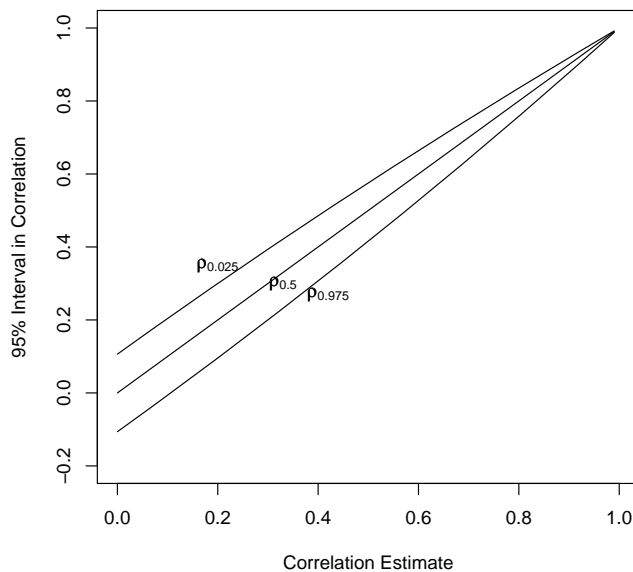
Figure 9.26   95% uncertainty interval of cross-correlation estimates due to finite sample size.

  grids.

- The need for a continuous surface of the spatial rate requires the parameter being varied across the region to be interpolated. This interpolation produces smoothing depending on the assumption of the covariance structure.

- Due to the calibration methodology it was only feasible to vary one of the parameters of the model whilst other parameters remain constant over the entire region. An issue was noticed during calibration that the μ_C parameter pursued unrealistic values at a large number of sites. This suggests that varying only one parameter is insufficient.

- The overlap of raincells requires the use of a buffer region to simulate the model. If the region of interest is near a coastline, then rainfall values will need to be extrapolated over the ocean so that the boundary effects are avoided.

The inhomogeneous model showed some improvement over a homogeneous model for several statistics. The statistics for the portion of dry days and the coefficient of variation was shown to vary over the region, although they did not perfectly reproduce the observed statistics. The daily cross-correlations exhibit a large degree of variability, be-
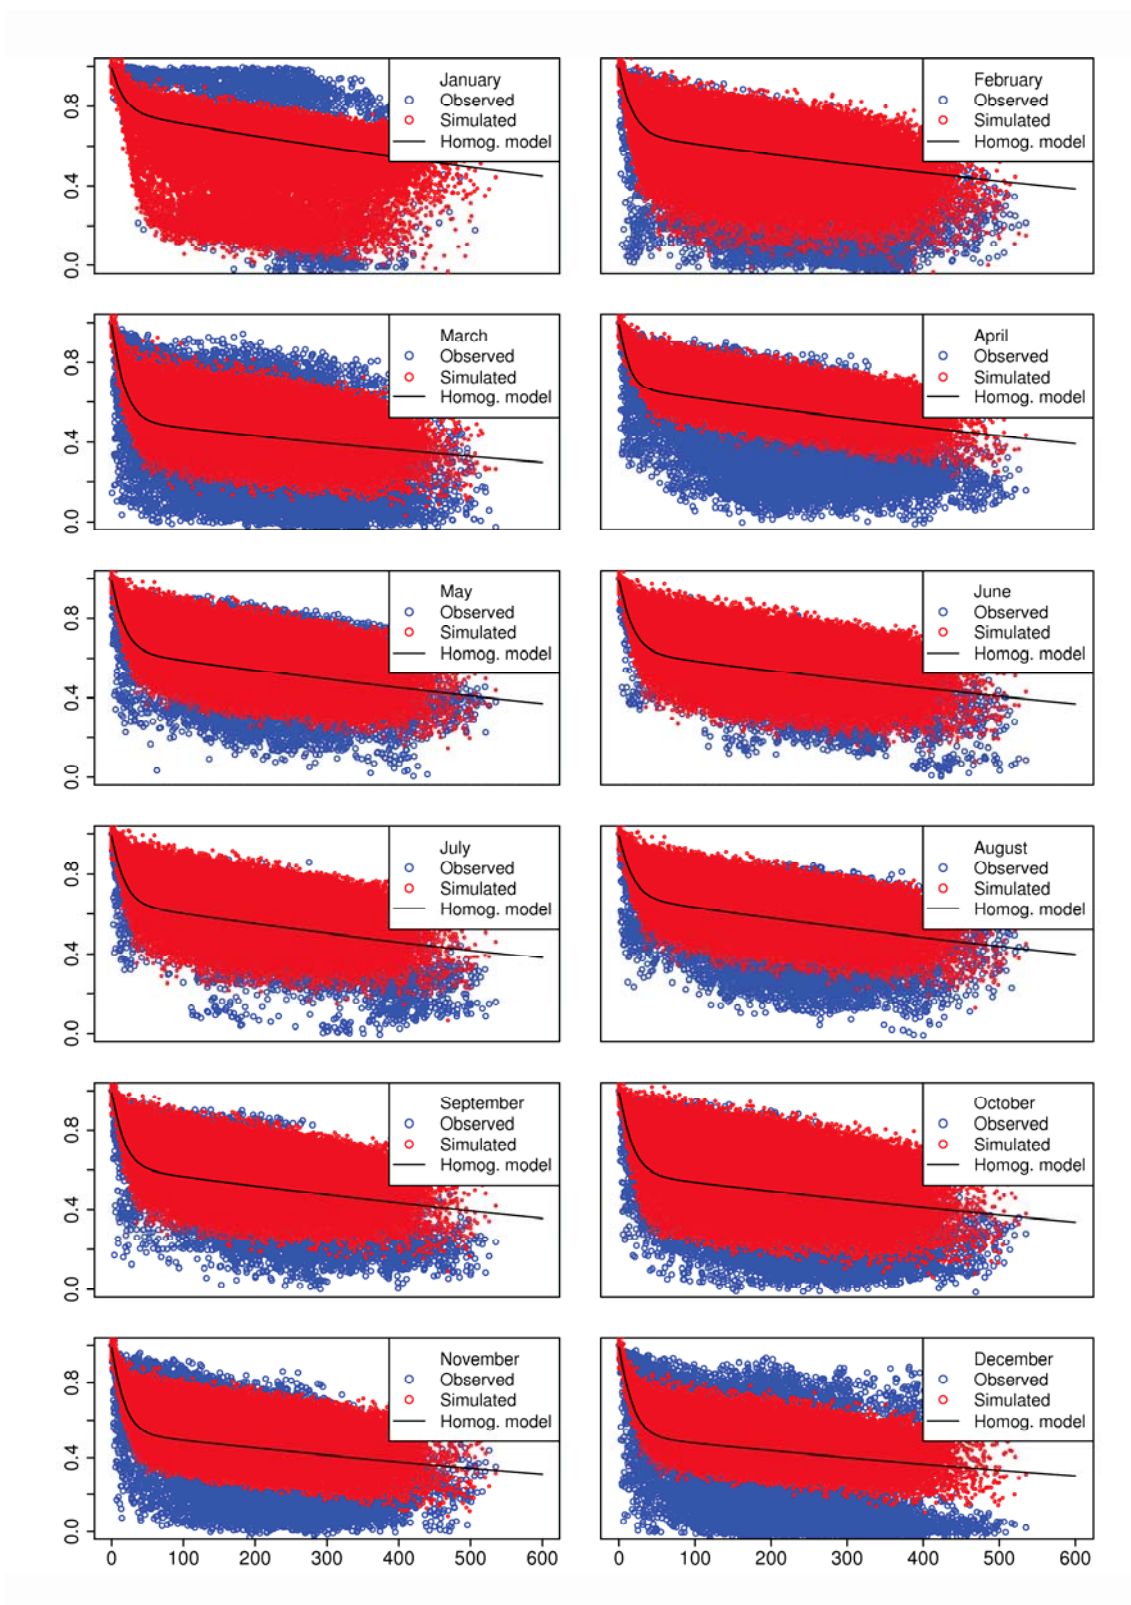
Figure 9.27   Comparison of observed cross-correlations with simulated cross-correlation. Correlations are evaluated over the entire Avon Basin for 12 months of the year.

yond the normal amount expected due to sampling variation. The inhomogeneous model was able to significantly increase the modelled variability of the cross-correlations above those of a homogeneous model. Nonetheless, the observed cross-correlations showed even greater variability than what the inhomogeneous model was able to produce.

Chapter 10
Summary and Recommendations

10.1 MODEL REVIEW

Rainfall is an inherently complex phenomenon. This has prompted a wide variety of approaches to modelling it, from physically ignorant statistical models to conceptual models having idealised physical processes to highly physical (meteorological) descriptions. This variety also reflects the wide range of objectives for rainfall modelling including empirical studies, forecasting, radar calibration, high resolution urban modelling, disaggregation of coarse rainfall inputs in space and/or time, ecological studies, agricultural studies, climatic studies, extreme value analysis and numerous other applications.

With the increasing numerical power of computers, continuous simulation models have become more popular for addressing hydrological questions. The aim of continuous simulation is to mimic the dominant rainfall patterns in observed data across a broad range of temporal and spatial scales. As a result, synthetic rainfall sequences can be used to evaluate risks associated with hydrological quantities. In doing so, continuous simulation models can overcome some of the limitations of more traditional approaches as discussed in Chapter 1.

In step with computing developments, the sophistication of rainfall models has increased to the point of modelling rainfall occurrence jointly in space and time. However, space-time models are not yet mature. On the one hand, models based on radar imagery are restricted in their ability to reproduce temporal statistics, and on the other, models based on a sparse network of rain-gauges have difficulty in describing the spatial process. This thesis focuses on rain-gauge models, and one in particular, the Neyman-Scott model [*Cowpertwait*, 1995; *Cowpertwait et al.*, 2002]. The developments to this model and the corresponding methodology are now summarised chapter by chapter.

Chapter 4: With the spatial $NSRP$ formulation there is a bias due to a spatial boundary. If only the raincells that originate within a region of interest are simulated then there is a boundary effect, since cells originating outside of the region can extend inside. It was concluded that wrapping raincells across a boundary is aesthetically unpleasing and that using a buffer region can be computationally inefficient. Therefore, a method for directly simulating only those cells that will intersect a target region (no matter where they originate) was developed. This method was demonstrated to provide a considerable improvement to the computational efficiency of the model, both in terms of memory and runtime. Chapter 7 extended this observation to the influence of a storm boundary on raincells, that a raincell can originate outside the storm boundary but extend inside. As the storms are themselves a homogeneous Poisson process, the method is applicable to them and can provide further efficiency. The development of an inhomogeneous Poisson process for raincells (Chapter 9) implies that the new simulation method cannot be used. The method therefore only has applicability where a spatial process is homogeneous.

Chapter 5: A common technique when calibrating rainfall models is to use separate parameters for each month to capture the seasonal process. The Neyman-Scott formulation has an offset between the onset of a storm and the arrival of the first rain cell, which is related to the parameter β. As the model is conceptual, it is not strictly necessary that calibrated parameters have a realistic interpretation, only that they can reproduce observed statistics. *Cowpertwait et al.* [2002] publishes a result for the parameter β that lies beyond reasonable interpretation in that the first occurrence of rainfall occurs days and weeks after the inception of the storm. This observation is largely irrelevant to simulated statistics, except when a storm crosses over a monthly boundary where there is a different rainfall regime corresponding to the new parameters in that month.

The amount of rainfall that lands in a subsequent month was derived from the model properties and used to assess the bias in the mean rainfall for a given month. The reported case-study from *Cowpertwait et al.* [2002] was used to demonstrate the magnitude of this bias. A calibration technique was developed to constrain this bias whereby all months were jointly calibrated with the aim that the incoming rainfall from a prior month should balance with the outgoing rainfall to subsequent months. The case-study demonstrated that the outcome of a constraint on the monthly bias was to increase the cell dispersion parameter, $\beta \approx 0.1$. The effect of this is that the amount of rainfall transferred across a month boundary is reduced to a negligible amount. It was concluded that rather than repeat the joint-calibration approach for subsequent model developments, it is more practical to constrain the β parameter directly and maintain the independent monthly calibration

procedure.

Chapter 6: The $NSRP$ model requires the derivation of analytic properties to be compared with observed data during calibration. This limits model development as it requires these model properties to be derived in the first place, which is non-trivial in many cases and also not feasible for many statistics. A related point is that, if an aspect of the model is changed (say a type of distribution for a given variable), this requires the properties to be rederived. By simulating the model for a given set of parameters, the simulated statistics can be compared to observed statistics to assess the quality of the parameters. This technique allows changes to the model formulation and also the inclusion of any relevant statistic in the calibration procedure. The L-skewness is a statistic that is not able to be derived for the $NSRP$ model and so was used as a direct illustration of the benefit of this approach.

The simulation-calibration method is highly computationally demanding, since it is necessary to use long records to reduce uncertainty in the statistics and since it is embedded within an optimisation routine.It was found that even with excessively long simulations a gradient-based optimiser was unable to locate good parameters and that a stochastic optimiser has to be used instead. As stochastic optimisers are themselves computationally demanding, the overall approach is very time consuming (on the order of days). To some extent, the structure of the spatial $NSRP$ can be exploited to mitigate the computational burden since the temporal $NSRP$ model is a special case of the spatial $NSRP$ model. As a result, the temporal parameters were calibrated first to a simulation at a single point, followed by the spatial parameters having optimised all other parameters.

While the technique investigated in this chapter is promising in that it allows new model formulations, it was found that the computational burden of the calibration was prohibitive. This is because of the requirement for stochastic optimisation. The technique was not adopted for use in subsequent chapters, and especially since some of the methods increase the computational demand of the model (e.g. spatial inhomogeneity).

Chapter 7: The spatial process of the *Cowpertwait et al.* [2002] $NSRP$ model is highly idealised. It was demonstrated that this model implicitly assumes that the region of interest is smaller than the storm. The model produces spurious correlations at long distances when larger regions are considered. An improved $NSRP$ model was developed to limit the spatial extent of a storm. As there is limited spatial information in sparse gauge networks, the characterisation of the storm was chosen to be parsimonious. By choosing a circular region and a random radius following the exponential distribution, the storm process required only one additional parameter. A visualisation of the modelled storm

was also presented to highlight the spatial characteristic of the model.

The extended model was demonstrated for the metropolitan Sydney region using 52 daily gauges and 24 sub-daily gauges. The daily cross-correlation statistic was added to calibration procedure in order to calibrate the storm radius parameter. Unless the simulation region is small, the inclusion of a storm radius is important and was retained with model extensions and case studies presented in subsequent chapters.

Chapter 8: Classical analyses of extreme values ignore the variability of extremes with respect to seasonality and climatic states. The term '1 in 100' years reflects this ignorance. This chapter demonstrated the difference in seasonal extremes for rainfall at a selection of major Australian cities. This influence of climatic states on extreme values was demonstrated for the Murray-Darling river system. A methodology was proposed so that climatic and seasonal partitions can be taken into account and related back to the non-partitioned distribution of annual maximums. The method was demonstrated for a design scenario involving a dual-purpose basin. While the method is promising, there were several limitations. Firstly, the presence of correlation among seasonal extremes creates a complication for the methodology. A copula function was used to remedy this, but these functions are not well suited to high dimensional problems (if there are many seasons). Secondly, climatic partitions reduce the effective amount of data to fit extreme value distributions. The resulting uncertainty associated with climatic partitions was investigated but remains a point for ongoing investigation.

A case-study involving the Bourke region of New South Wales was used to highlight the presence of seasonal and climatic variability in rainfall records. The Neyman-Scott model was calibrated to two separate datasets for each month (i) using data when the Southern Oscillation Index (SOI) is positive and (ii) when the SOI is negative. A conventional timeseries model was used to generate synthetic sequences of the SOI and the parameter sets for each simulated month were based on the simulated SOI values.

Chapter 9: An inhomogeneous Poisson random field was introduced for the spatial rate of cells. The main benefit of this approach is that some statistics such as the dry probability cannot be accurately reproduced using the existing homogeneous model. The premise of this model is that it can be more likely for raincells to occur in some regions than in others. The Avon Basin case-study was used since it had previously been investigated by *Jothityangkoon et al.* [2000] who developed a non-stationary fractal model.

The calibration technique first used a regional estimate of parameters and then the parameter for the number of cells, μ_C was recalibrated at each gauge while maintaining the other constant parameters. A thin-plate surface was fitted to the μ_C parameter, which was

subsequently used to find the underlying rate of cell occurrence. The model required a linear approximation to facilitate analytic solutions when integrating the probability that a raincell overlaps a certain point. The main benefit of this was computational efficiency. An advanced numerical technique was implemented to obtain the underlying spatial rate of cells from the mean rainfall image. An issue arises with the inversion technique due to the presents of strong rainfall gradients, which leads to the possibility of spurious negative spatial rates. The model is formulated as a solution to a system of linear equations, however, it is not straightforward to solve this system with an additional constraint that the rate of cells is a positive value. The data was smoothed to overcome this issue.

The statistics demonstrate that the observed daily cross correlations are highly variable and that this cannot be reproduced by a homogeneous model, but it can be reproduced by the inhomogeneous model. While there is some improvement in other statistics, as compared to the homogeneous model, there is scope for further improvement. The main reason for this limitation is that only one parameter was allowed to vary across each site, but this may not be sufficient to capture the complexity of the rainfall structure.

10.2 RECOMMENDATIONS

There are several themes demonstrated in this thesis with regards to space-time rainfall models.

1. That the techniques needed to model rainfall in space and time are computationally demanding.

2. That there is a wide array of statistics that can be used for calibrating and assessing the model.

3. That data are not necessarily homogeneous in space and time.

The model adopted in this thesis is from the family of point-process models, which is only one family amongst many types of models. As space-time models are still in their infancy, it is not yet clear which framework offers the best compromise between the competing objectives of model parsimony, reproduction of key statistics, computational efficiency and ease of calibration. Therefore, it is recommended that the development of space-time rainfall models is pursued on multiple fronts. Several promising alternatives are briefly discussed below

1. Markov random field models - can be formulated for gridded data or for multisite data. A main advantage of these models is that they condition a rainfall amount

only on those values within the immediate neighbourhood, rather than all possible values. They also have robust calibration and simulation methodologies because of their relationship to Markov Chain Monte Carlo techniques. Their main disadvantage is in their computational burden.

2. Cascade models - are popularly used to disaggregate a rainfall mass. Their main advantage is they are conceptually simple. A difficulty is incorporating space-time correlations and advection into them (such as observed in a moving storm) because these concepts do not align with the cascade mechanism.

3. Interpolation of multisite models - is a simple but possibly effective technique for generating spatial images of rainfall in space and time. The main advantages being that there are numerous existing multisite models and that they do not confound the issue of spatial infilling with reproduction of the observations across the gauge network. Spatial interpolation techniques such as thin-plate splines and kriging are well developed.

No matter what model is chosen there is an overarching question on how to best verify that the model is sufficiently reproducing the observed statistics. A region having 10 rain gauges, each with 10 statistics compared at the 3 different timescales for each month leads to 3600 different statistics. These statistics can be thought of as local statistics that assess a specific aspect of the model for a specific location or month. In addition to this are global statistics that are aggregates of the local statistics for assessing the performance over the entire region or, for example, the annual or inter-annual scale. Not only is there a vast number of statistics, but, the local and global statistics can be in competition with each other since a global improvement to the model may cause considerable discrepancies in local aspects of the model. There is the need for fundamental research into the methodology used to assess rainfall models. This assessment should question the assumptions of homogeneity and the basis for defining regions.

Many factors such as the popularity of GIS packages, the advent of distributed runoff models, the progress of continuous simulation and ongoing improvements in computing power suggest that demand for spatial rainfall as an input to hydrological models will continue to grow. One promising direction for future research is to combine continuous simulation rainfall models with distributed runoff models. Since the catchment acts as an aggregator, the statistics of observed streamflow serve as a good assessment of model performance. This is because it can address the trade-offs between high spatial-resolution but possibly biased radar-rainfall estimates with and lower spatial-resolution but more

reliable gauge measurements. Numerical and modelling issues aside, the main advantage of this approach is that it would facilitate a more objective means for assessing the wide range of rainfall models (radar-rainfall, multi-site, single-site, random field models) and the appropriateness of their assumptions.