22nd National Conference of the Australian Society for Operations Research, Adelaide, Australia, 1–6 December 2013
www.asor.org.au/conferences/asor2013

Modelling and simulation of seasonal rainfall

Jon Borwein [a], **Phil Howlett** [b], **Julia Piantadosi** [b]

[a] *Centre for Computer Assisted Research Mathematics and its Applications (CARMA), University of Newcastle, Callaghan, NSW 2308, Australia.*
[b] *Scheduling and Control Group (SCG), Centre for Industrial and Applied Mathematics (CIAM), University of South Australia, Mawson Lakes, SA, 5095, Australia.*
Email: phil.howlett@unisa.edu.au

Abstract: We propose and justify a model for seasonal rainfall using a copula of maximum entropy to model the joint distribution and using gamma distributions to model the marginal monthly rainfalls. The model allows correlation between individual months and thereby enables a much improved model for seasonal variation. A central theme is the principle of maximum entropy which we use to find the most parsimonious representation for the underlying distributions—using the minimum possible number of parameters to model the relevant physical characteristics. A particular emphasis is the use of the gamma distribution to model the marginal monthly rainfalls.

We wish to simulate monthly and seasonal rainfall at Kempsey, NSW during February–March–April. Our first task is to explain why we choose to model monthly rainfall totals using the gamma distribution. The principle of maximum entropy (Jaynes, 1957a, 1957b) states that, subject to precisely stated prior data, the probability distribution which best represents the current state of knowledge is the one with largest entropy. We will use this principle to argue that the gamma distribution is the best distribution to represent monthly rainfall totals provided the means of the observed monthly totals and the natural logarithm of the observed monthly totals are both well-defined and finite. This is true if the observed totals are always strictly positive. Our second task is to devise a graphical representation that displays the gamma distribution in the simplest possible way—as a straight line. We will use this procedure to compare simulated data from the chosen gamma distribution to the observed data. Our third task is to use the simple graphical representation above to compare the observed monthly rainfall to simulated monthly rainfall generated by the chosen gamma distribution. Our conclusion will be that there is no significant statistical difference between the simulated data and the observed data. Our fourth task is to to demonstrate the *goodness-of-fit* for the observed monthly rainfall data to the selected gamma distributions for each month. To do this we used two kinds of Q-Q plot. Firstly we plot simulated quantiles from the gamma distribution against theoretical quantiles to determine 95% confidence intervals and then plot observed quantiles against theoretical quantiles.

Once it has been decided that the monthly rainfall X_i can be modelled by a gamma distribution $X_i \sim \Gamma(\alpha_i, \beta_i)$ with $F_i(x) = F_{\alpha_i,\beta_i}(x)$ then the observed data set $\{x_{i,j}\}_{j=1,2,\ldots,N}$ can be transformed into a corresponding data set $\{u_{i,j} = F_i(x_{i,j})\}_{j=1,2,\ldots,N}$ for each $i = 1, 2, \ldots, m$. This has the effect of removing seasonal factors from the observed data and also preparing for the use of a copula of maximum entropy to model the joint distribution of the monthly rainfall totals. The next step in the modelling process is to construct a joint probability distribution for the entire three-month time period. Past studies of rainfall accumulations over several months (Katz and Parlange, 1998; Rosenberg et al., 2004; Withers and Nadarajah, 2011) have observed that for models with independent marginal distributions the seasonal variance is often too low. It has been suggested that this happens because there is an overall positive correlation between the individual monthly totals. Since the observed data shows positive correlation for February-March-April at Kempsey our aim will be to construct a joint distribution so that the desired marginal distributions are preserved and so that the grade correlation coefficients match the observed rank correlation coefficients. We construct the joint distribution using a checkerboard copula of maximum entropy (Piantadosi et al., 2012a, 2012b).

Finally we compared the observed rainfall to rainfall generated by three different models (a) a maximum likelihood gamma distribution that models seasonal rainfall but does not generate individual monthly rainfalls. (b) a checkerboard copula of maximum entropy with marginal gamma distributions that preserves the observed rank correlation coefficients and (c) a joint distribution with independent marginal gamma distributions. We conclude that the copula of maximum entropy provides an excellent model for rainfall simulation.

Keywords: *Rainfall modelling and simulation, maximum entropy, gamma distribution, checkerboard copula.*

1 MODELLING MONTHLY RAINFALL

The gamma distribution is often used to model rainfall accumulations. A common justification is that it is sufficiently *flexible* to model a wide range of observed data. While this may be true the argument is essentially subjective. The principle of maximum entropy (Jaynes, 1957a, 1957b) provides an objective justification.

1.1 Maximum entropy and the gamma distribution

We assume that X is a random variable and that the observed values $\{x_n\}_{n=1}^N$ are strictly positive. We wish to find a probability density $f : (0,\infty) \to (0,\infty)$ such that the differential entropy

$$h(f) = (-1)\int_0^\infty f(x)\log_e f(x)dx \tag{1.1}$$

is maximized subject to the additional constraints imposed by the observed means

$$E[X] = \overline{x} = \frac{1}{N}\sum_{n=1}^N x_n \quad \text{and} \quad E[\log_e X] = \overline{\log_e x} = \frac{1}{N}\sum_{n=1}^N \log_e x_n. \tag{1.2}$$

We can formulate this problem as a convex optimization with linear constraints. From the theory of Fenchel duality and the Fenchel-Young inequality (Borwein and Vanderwerff, 2010 [pp. 171-178]) we have

$$\begin{aligned}
p &= \inf_{f\in L^1[(0,\infty)]}\left\{-h(f)-1 \mid E[1]=1, E[X]=\overline{x}, E[\log_e X]=\overline{\log_e x}\right\}\\
&\geq \sup_{(\alpha,\beta,\kappa)\in\mathbb{R}^3}\left\{\log_e\kappa - \overline{x}/\beta + (\alpha-1)\overline{\log_e x} - \kappa\int_0^\infty x^{\alpha-1}e^{-x/\beta}dx\right\}\\
&= \sup_{(\alpha,\beta,\kappa)\in\mathbb{R}^3}\left\{\log_e\kappa - \overline{x}/\beta + (\alpha-1)\overline{\log_e x} - \kappa\,\Gamma(\alpha)\beta^\alpha\right\}\\
&= \sup_{(\alpha,\beta,\kappa)\in\mathbb{R}^3}\varphi(\alpha,\beta,\kappa)\\
&= -\log_e[\Gamma(\alpha)\beta] + (\alpha-1)\psi(\alpha) - (\alpha+1) = d
\end{aligned} \tag{1.3}$$

where the parameters α, β and κ are determined by the equations

$$\log_e\beta + \psi(\alpha) = \overline{\log_e x}, \quad \alpha\beta = \overline{x}, \quad \kappa(\alpha,\beta) = \frac{1}{\Gamma(\alpha)\beta^\alpha} \tag{1.4}$$

and where $\psi(\alpha) = \Gamma'(\alpha)/\Gamma(\alpha)$ is the digamma function. The supremum and the conditions (1.4) are found simply by solving the equations $\partial\varphi/\partial\alpha = 0$, $\partial\varphi/\partial\beta = 0$ and $\partial\varphi/\partial\kappa = 0$. The function

$$f_{\alpha,\beta}(x) = \frac{1}{\Gamma(\alpha)\beta^\alpha}x^{\alpha-1}e^{-x/\beta}$$

which arises naturally in (1.3) when solving the dual optimization problem to find d is the probability density on $(0,\infty)$ for the gamma distribution with parameters α and β. If X is a random variable with this distribution we write $X \sim \Gamma(\alpha,\beta)$. In the case where α and β are determined by (1.4) then the additional constraints (1.2) are also satisfied. Since it is easy to show that $-h(f_{\alpha,\beta}) - 1 = d$ it follows that $p = d$ and that $f_{\alpha,\beta}$ is the unique solution to our original convex optimization problem. Note that the equations (1.4) are also the maximum likelihood equations used to estimate α and β if one has decided *a priori* to fit a gamma distribution.

1.2 A linear representation of the gamma distribution

Let $X \sim \Gamma(\alpha,\beta)$ be a gamma random variable with probability density $f(x) = f_{\alpha,\beta}(x)$ for $x \in (0,\infty)$. Therefore

$$(\log_e f)(x) = (\alpha-1)\log_e x - x/\beta - \alpha\log_e\beta - \log_e\Gamma(\alpha)$$

and hence by differentiating both sides and substituting $y = 1/x$ we obtain

$$(\log_e f)'(1/y) = (\alpha-1)y - 1/\beta. \tag{1.5}$$

Note that the reciprocal random variable $Y = 1/X$ has probability density given by $g(y) = f(1/y)/y^2$ for $y \in (0,\infty)$. Choose $h > 0$ and consider the collection of half-open intervals $y \in (2rh, 2(r+1)h]$ for

$r = 0, 1, \ldots, R$. Write $\eta_r = (2r+1)h$ for $r = 0, 1, \ldots, R$ to denote the midpoints of these intervals. Let $\{y_n\}_1^N$ denote the observed values of Y in an experiment with N independent trials and let n_r be the number of outcomes lying in the interval $(2rh, 2(r+1)h]$. The corresponding observed frequency is $\widehat{p}_r = n_r/N$. The actual probability is given by

$$p_r = \int_{2rh}^{2(r+1)h} g(y)dy \approx 2hg(\eta_r) = 2hf(1/\eta_r)/\eta_r^2 = \frac{2f(1/\eta_r)}{(2r+1)^2 h}$$

which we can rewrite as $f(1/\eta_r) \approx (2r+1)^2 h p_r/2$. If $\xi_r = 1/(2rh)$ then $1/\eta_r < \xi_r < 1/\eta_{r-1}$ for $r = 1, \ldots, R$ and it follows that

$$(\log_e f)'(\xi_r) \approx \frac{(\log_e f)(1/\eta_{r-1}) - (\log_e f)(1/\eta_r)}{1/\eta_{r-1} - 1/\eta_r} = (4r^2-1)h \log_e \frac{2r-1}{2r+1}\sqrt{\frac{p_{r-1}}{p_r}} = w(r, p_{r-1}, p_r).$$

Let $y_r = 2rh \Leftrightarrow 1/y_r = \xi_r$ in which case the above approximation and equation (1.5) give

$$w(r, p_{r-1}, p_r) \approx (\alpha - 1)2rh - 1/\beta \tag{1.6}$$

for each $r = 1, 2, \ldots, R$. For a particular experiment when N is large the probability p_r is estimated by the modified observed frequency $\widehat{p}_r + \epsilon_r = (n_r+1)/(N+R+1)$ for each $r = 0, 1, \ldots, R$ where the modification avoids the possibility of a zero value. If we define $w_r = w(r, \widehat{p}_{r-1} + \epsilon_{r-1}, \widehat{p}_r + \epsilon_r)$ then the points (y_r, w_r) for each $r = 1, 2, \ldots, R$ will lie close to the line $w = (\alpha - 1)y - 1/\beta$.

1.3 A model for monthly rainfall at Kempsey in NSW

All observed monthly totals at Kempsey (NSW) in February, March and April for 1889–2011 are positive. Hence a gamma distribution provides the best model. Thus $X_i \sim \Gamma(\alpha_i, \beta_i)$ where α_i and β_i are determined by maximum likelihood. To simulate the monthly rainfall we used gamma distributions with parameters

$$\boldsymbol{\alpha} = (1.5502, 2.0134, 1.2735) \quad \text{and} \quad \boldsymbol{\beta} = (100.4753, 77.2556, 91.1034).$$

For each month we conducted 1000 simulation trials each over $N = 123$ years. The plots of (y_r, w_r) for each trial are shown in Figure 1. In each case the trials were used to find empirical 95% confidence intervals for w_r for each $r = 1, \ldots, 8$. If the observed monthly rainfall can be represented by the selected gamma distribution then the observed data (y_r, w_r) should be quasi-linear. The departure from true linear will mainly be due to random error. In Figure 2 we plotted the observed data against the empirical 95% confidence intervals from the simulated data. In all cases the observed points lie within the relevant confidence intervals.

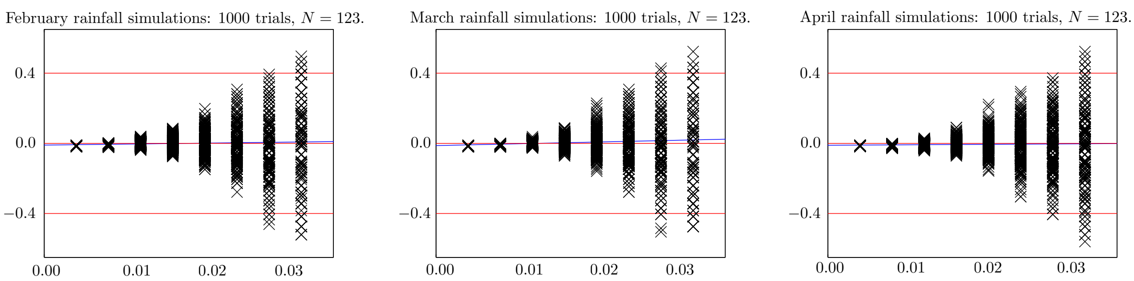

Figure 1: Plots (y_r, w_r) for simulated monthly rainfall at Kempsey with $X_i \sim \Gamma(\alpha_i, \beta_i)$ for 1000 trials with each trial covering $N = 123$ years for February (left), March (centre) and April (right).

1.4 The Q-Q plots—theory, simulation and observations

Firstly we used the gamma distribution to generate 1000 simulation trials with each trial covering a period of $N = 123$ years. Then we plotted the simulated quantiles against the theoretical quantiles. The results

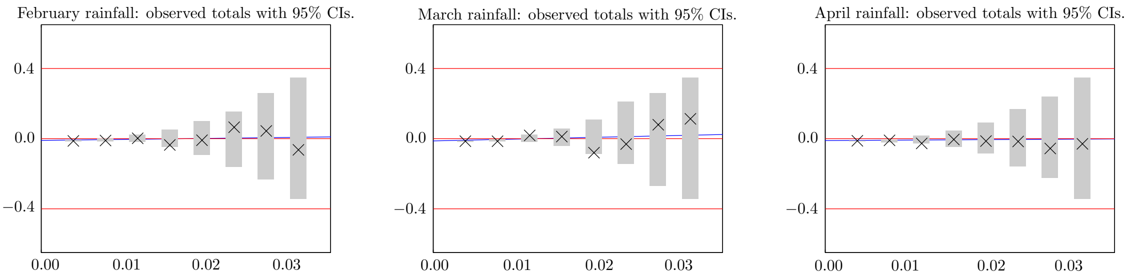

Figure 2: Plots (y_r, w_r) for observed monthly rainfall at Kempsey with empirical 95% confidence intervals for the respective gamma distributions for February (left), March (centre) and April (right).

are shown in Figure 3. These plots show the full range of random variation that one should expect from observations of the gamma random variable over a period of 123 years . By discarding the bottom 25 and top 25 values for each quantile from the simulated data we found empirical 95% confidence intervals. Secondly we plotted the observed quantiles against the theoretical quantiles for the gamma distributions. The results are shown in Figure 4. We used grey bars on these plots to show the empirical 95% confidence intervals for the simulated quantiles of each gamma distribution. In all cases the observed values lie within the desired intervals. Thus there is not enough statistical evidence to reject the hypothesis that the monthly rainfall totals can be modelled by the maximum likelihood gamma distributions.

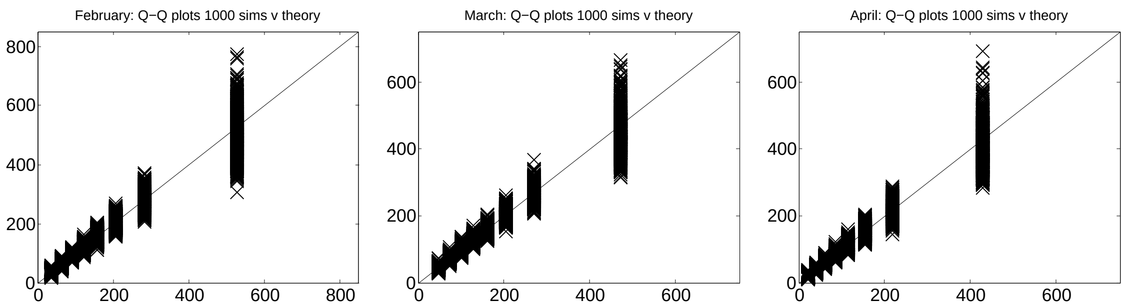

Figure 3: Q-Q plots of the quantiles for 1000 simulated trials with each trial covering a period of 123 years versus the corresponding theoretical quantiles for the respective maximum likelihood gamma distributions used to model monthly rainfall at Kempsey in February (left), March (centre) and April (right).

1.5 The transformed data—removal of seasonal effects

If $X_i \sim \Gamma(\alpha_i, \beta_i)$ where $F_i(x) = F_{\alpha_i,\beta_i}(x)$ then the transformed monthly rainfalls $U_i = F_i(X_i)$ are uniformly distributed on $(0, 1)$ and so seasonal factors are removed from the observed data. Figure 5 shows histograms of the observed data for transformed monthly rainfall totals at Kempsey. We used the binomial distribution with $N = 123$, $p = 0.1$ and $q = 0.9$ to calculate approximate 95% confidence intervals $I = (p - 1.96\sqrt{pq/N}, p + 1.96\sqrt{pq/N}) = (0.047, 0.153)$ for the heights of the bars. All but 1 of the 30 observed frequencies lie within these limits—a failure rate of 3.3% which is consistent with what one should expect.

2 COPULAS WITH PRESCRIBED CORRELATION

An m-dimensional *copula*, where $m \geq 2$, is a cumulative probability distribution $C(\boldsymbol{u}) \in [0, \infty)$ defined on the m-dimensional unit hypercube $\boldsymbol{u} = (u_1, u_2, \ldots, u_m) \in [0, 1]^m$ for a vector-valued random variable

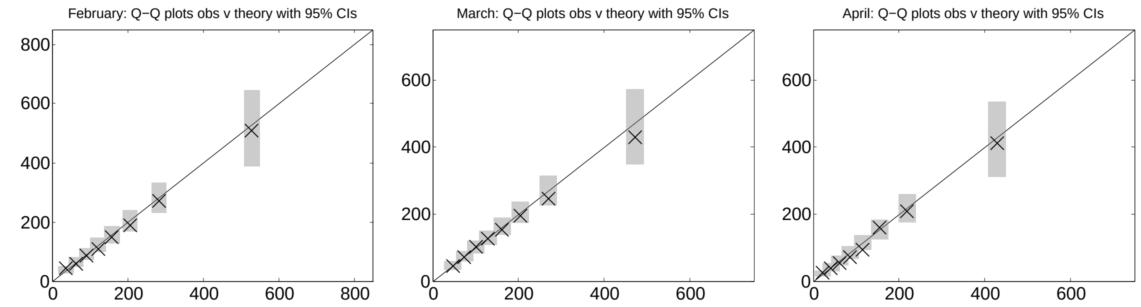

Figure 4: Q-Q plots for the observed quantiles versus the corresponding theoretical quantiles for the monthly rainfall model at Kempsey in February (left), March (centre) and April (right). The vertical grey bars denote the empirical 95% confidence intervals for the simulated quantiles generated by the gamma distributions.

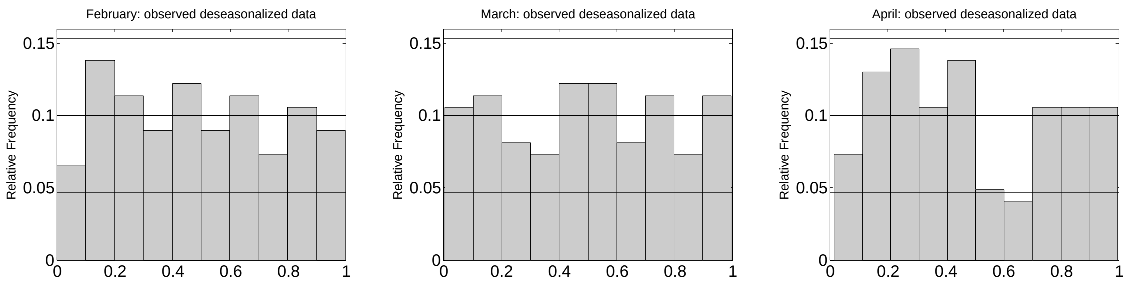

Figure 5: Histograms of the transformed monthly rainfall totals $u_{i,j} = F_i(x_{i,j})$ in February (left), March (centre) and April (right) at Kempsey where $F_i = F_{\alpha_i,\beta_i}$ is the cumulative distribution for $\Gamma(\alpha_i, \beta_i)$. The horizontal lines show the mean value and the upper and lower bounds for the 95% confidence intervals.

$\boldsymbol{U} = (U_1, U_2, \ldots, U_m)$ with uniform marginal probability distributions for the real-valued random variables $U_1, U_2, \ldots, U_m$. See (Nelsen, 1999). The correlation coefficients for the joint distribution are defined by

$$\rho_{r,s} = \frac{E[(U_r - 1/2)(U_s - 1/2)]}{\sqrt{E[(U_r - 1/2)^2]\, E[(U_s - 1/2)^2]}} = 12E[U_r U_s] - 3 \tag{2.7}$$

for each $1 \le r < s \le m$. In order to model the joint probability distribution for a vector-valued random variable $\boldsymbol{X} = (X_1, X_2, \ldots, X_m) \in (0, \infty)^m$ with known marginals $u_i = F_i(x_i)$ we simply construct uniformly distributed random variables $U_i = F_i(X_i) \in (0, 1)$ for each $i = 1, 2, \ldots, m$ and use the m-dimensional copula $C(\boldsymbol{u}) = C(\boldsymbol{F}(\boldsymbol{x})) = C(F_1(x_1), F_2(x_2), \ldots, F_m(x_m))$. We say that the grade correlation coefficients for $\boldsymbol{X}$ are simply the correlation coefficients for $\boldsymbol{U}$ defined above. That is

$$\rho_{r,s} = \frac{E[(F_r(X_r) - 1/2)(F_s(X_s) - 1/2)]}{\sqrt{E[(F_r(X_r) - 1/2)^2]\, E[(F_s(X_s) - 1/2)^2]}} = 12E[F_r(X_r)F_s(X_s)] - 3 \tag{2.8}$$

for each $1 \le r < s \le m$. We distinguish between the Spearman rank correlation coefficients $\widehat{\rho}_{r,s}$ (Nelsen, 1999) for the observed data $\{u_{i,j}\}_{j=1,\ldots,N}$ and the grade correlation coefficients $\rho_{r,s}$ defined by (2.8).

3 Modelling the joint probability with a checkerboard copula

We construct a joint distribution using a checkerboard copula of maximum entropy (Piantadosi et al., 2012a, 2012b). A trivariate *checkerboard* copula is a probability distribution defined by subdividing the unit cube into n^3 congruent small cubes with constant density on each one. If the density on I_{ijk} is defined by $n^2 h_{ijk}$ then the marginal distributions will be uniform if

$$\sum_{j,k} h_{ijk} = 1 \text{ for all } i, \quad \sum_{i,k} h_{ijk} = 1 \text{ for all } j, \quad \sum_{i,j} h_{ijk} = 1 \text{ for all } k.$$

In such cases we say that $\boldsymbol{h} = [h_{ijk}]$ is triply-stochastic. We wish to construct a joint density in this form with the desired correlations. For sufficiently large n there are many ways that this can be done. *The principle of maximum entropy suggests that the best such distribution is the most disordered or least prescriptive solution—the triply-stochastic hyper-matrix $\boldsymbol{h}$ which has the most equal subdivision of probabilities but still allows the required correlations.*

Problem 1 (The primal problem). *Find the hyper-matrix $\boldsymbol{h} = [h_{\boldsymbol{i}}] \in \mathbb{R}^\ell$ where $\boldsymbol{i} = (i_1, \ldots, i_m)$ and $\ell = n^m$ to maximize the entropy*

$$J(\boldsymbol{h}) = (-1)\left[\frac{1}{n}\sum_{\boldsymbol{i}\,\in\,\{1,\ldots,n\}^m} h_{\boldsymbol{i}} \log_e h_{\boldsymbol{i}} + (m-1)\log_e n\right] \tag{3.9}$$

subject to the multi-stochastic constraints

$$\sum_{j\neq r,\, i_j \in \{1,\ldots,n\}} h_{\boldsymbol{i}} = 1 \tag{3.10}$$

for all $i_r \in \{1, \ldots, n\}$ and each $r = 1, \ldots, m$ and $h_{\boldsymbol{i}} \geq 0$ for all $\boldsymbol{i} \in \{1, \ldots, n\}^m$ and the grade correlation coefficient constraints

$$12\left[\frac{1}{n^3}\cdot\sum_{\boldsymbol{i}\,\in\,\{1,\ldots,n\}^m} h_{\boldsymbol{i}}(i_r - 1/2)(i_s - 1/2)\right] - 3 = \widehat{\rho}_{r,s} \tag{3.11}$$

for $1 \leq r < s \leq m$ where $\rho_{r,s}$ is known for all $1 \leq r < s \leq m$.

The problem can be solved using the theory of Fenchel duality. See (Piantadosi et al., 2012a, 2012b). The m-dimensional copula of maximum entropy is defined by $m(m-1)/2$ real parameters—the grade correlation coefficients—defined in equation (2.8).

3.1 A model for February-March-April rainfall at Kempsey

We set $m = 3$ and $n = 4$. The multi-stochastic hyper-matrix $\boldsymbol{h} \in \mathbb{R}^{4\times4\times4}$ describing the trivariate checkerboard copula of maximum entropy is shown below. The copula was constrained by setting $\rho_{12} = \widehat{\rho}_{12} = 0.202$, $\rho_{13} = \widehat{\rho}_{13} = 0.112$ and $\rho_{23} = \widehat{\rho}_{23} = 0.152$. We calculate

$$\boldsymbol{h}_1 \approx \begin{bmatrix} 0.1262 & 0.0975 & 0.0733 & 0.0536 \\ 0.0870 & 0.0756 & 0.0639 & 0.0525 \\ 0.0567 & 0.0554 & 0.0527 & 0.0487 \\ 0.0350 & 0.0384 & 0.0411 & 0.0427 \end{bmatrix}, \boldsymbol{h}_2 \approx \begin{bmatrix} 0.0920 & 0.0765 & 0.0618 & 0.0486 \\ 0.0750 & 0.0701 & 0.0637 & 0.0563 \\ 0.0578 & 0.0608 & 0.0621 & 0.0618 \\ 0.0422 & 0.0499 & 0.0573 & 0.0641 \end{bmatrix},$$

$$\boldsymbol{h}_3 \approx \begin{bmatrix} 0.0641 & 0.0573 & 0.0499 & 0.0422 \\ 0.0618 & 0.0621 & 0.0608 & 0.0578 \\ 0.0563 & 0.0637 & 0.0701 & 0.0750 \\ 0.0486 & 0.0618 & 0.0765 & 0.0920 \end{bmatrix}, \boldsymbol{h}_4 \approx \begin{bmatrix} 0.0427 & 0.0411 & 0.0384 & 0.0350 \\ 0.0487 & 0.0527 & 0.0554 & 0.0567 \\ 0.0525 & 0.0639 & 0.0756 & 0.0870 \\ 0.0536 & 0.0733 & 0.0975 & 0.1262 \end{bmatrix},$$

where $\boldsymbol{h}_i = [h_{ijk}]$. The entropy is given by $J(\boldsymbol{h}) \approx -0.040714$. Similar results are obtained if the copula of maximum entropy is replaced by a checkerboard normal copula although numerical calculation of the latter is considerably more difficult (Piantadosi et al., 2012b) and the entropy is slightly less.

4 CONCLUSIONS

We compared the observed statistics for total rainfall to statistics for models using (a) the maximum likelihood gamma distribution[1] (b) the checkerboard copula of maximum entropy with marginal gamma distributions and (c) the joint distribution with independent marginal gamma distributions. Details can be found elsewhere (Piantadosi et al., 2012b). All models have mean 427 mm equal to the observed mean. The variances are (a) 49996 mm^2 (b) 47448 mm^2 (c) 38237 mm^2 while the observed variance is 53325 mm^2. Although the copula of maximum entropy appears to under-estimate the seasonal variance we see that the model using a maximum likelihood gamma distribution has a very similar variance. The selected results in Figure 6 from 16 successive trials over a period of $N = 123$ years show that the sample statistics are not necessarily represenative of the

[1] This model generates simulated seasonal rainfall totals directly and does not generate individual monthly rainfall totals.

underlying distribution. More reliable statistics would require observations over a much larger timespan. In Figure 7 we compare the observed totals with a model using the maximum likelihood gamma distribution and also show simulation results using the copula of maximum entropy over a period of 12300 years. The seasonal statistics for both models are quite similar. We conclude that the copula of maximum entropy with marginal gamma distributions provides an excellent model for both seasonal and monthly rainfall.

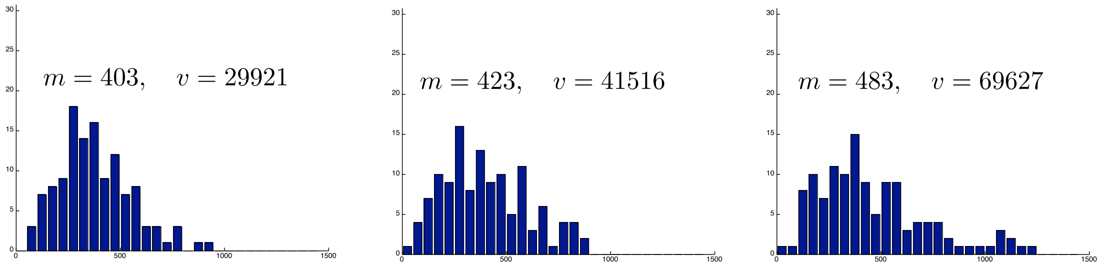

Figure 6: Selected histograms for total rainfall from 16 successive random simulations of total rainfall at Kempsey for February–March–April over a period of $N = 123$ years using the copula of maximum entropy with marginal gamma distributions. The plots—Trial #12 (left), Trial #13 (centre) and Trial #16 (right)—show typical sample variation for $N = 123$ years. The simulated mean m and variance v are shown on the plots.

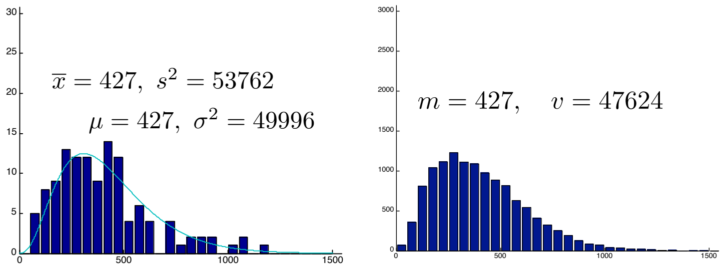

Figure 7: Histogram for observed total rainfall for February–March–April during 1889–2011 with mean $\overline{x}$ and variance s^2 and maximum likelihood gamma distribution $X \sim \Gamma(3.6524, 116.9983)$ with mean μ and variance σ^2 (left) and histogram for a simulation trial over a period of 12300 years using the copula of maximum entropy and marginal gamma distributions with mean m and variance v (right).

REFERENCES

Borwein, J. M. and J. D. Vanderwerff (2010). *Convex Functions: Constructions, Characterizations and Counterexamples*, Volume 109 of *Encyclopedia of Mathematics and its Applications*. CUP.

Jaynes, E. T. (1957a). Information theory and statistical mechanics. *Physical Review 106*(4), 620–630.

Jaynes, E. T. (1957b). Information theory and statistical mechanics ll. *Physical Review 108*(2), 171–190.

Katz, R. W. and M. B. Parlange (1998). Overdispersion phenomenon in stochastic modelling of precipitation. *J. Climate 11*, 591–601.

Nelsen, R. (1999). *An Introduction to Copulas*, Volume 139 of *Lecture Notes in Statistics*. Springer, NY.

Piantadosi, J., P. G. Howlett, and J. M. Borwein (2012). Copulas of maximum entropy. *Optimization Letters 6*(1), 99–125.

Piantadosi, J., P. G. Howlett, J. M. Borwein, and J. Henstridge (2012). Maximum entropy methods for generating simulated rainfall. *Numerical Algebra, Control and Optimization 2*(2), 233–256.

Rosenberg, K., J. Boland, and P. G. Howlett (2004). Simulation of monthly rainfall totals. *ANZIAM J. 46*(E), E85–E104.

Withers, C. S. and S. Nadarajah (2011). On the compound poisson-gamma distribution. *Kybernetika (Prague) 47*, 15–37.