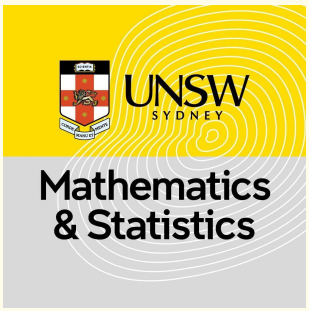



**MATH5805**
**A Modern Introduction to Extreme Value Theory**
**Term 1, 2025**

**Dr. Boris Béranger**

## Week 8 - Lab Solutions

### Exercise 1

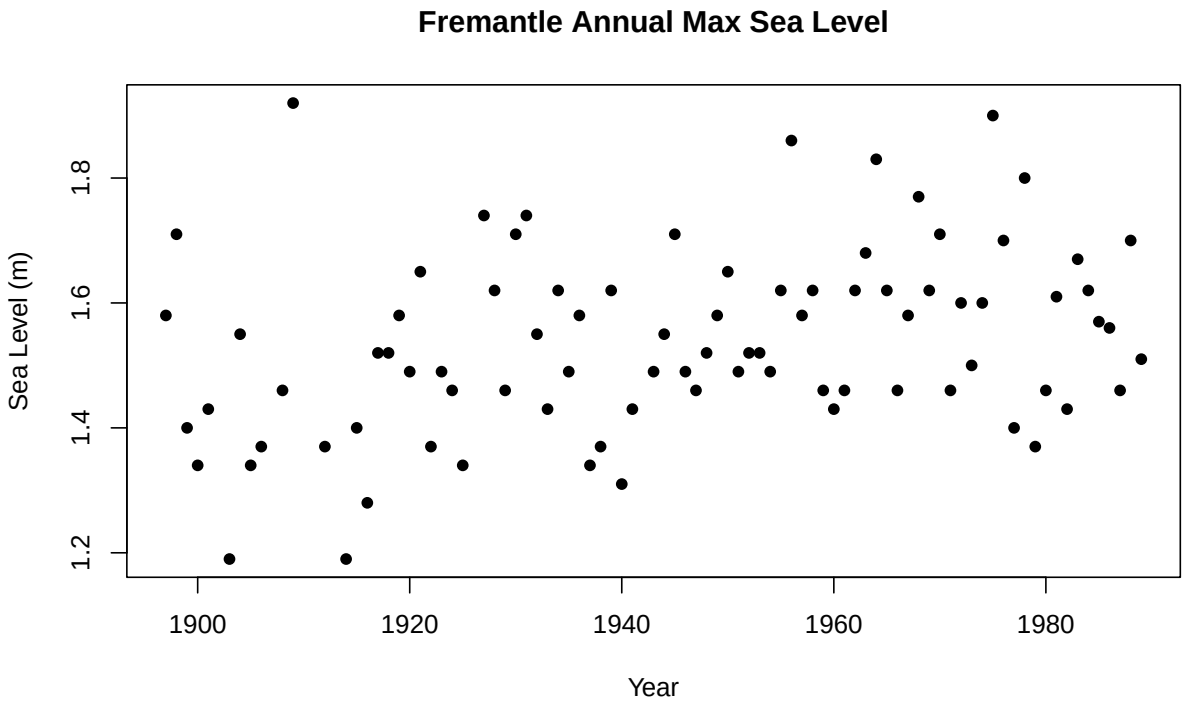
In the Week 3 Lab, we looked at the annual maximum sea levels in metres at Freemante in Western Australia from 1837–1989 (obtained via `data(fremantle)` in the `ismev` package). Similarly, annual maximum sea levels at Port Pirie in South Australia are available via `data(portpirie)`.

a) Plot both datasets and produce a scatterplot of the annual maxima for the years where data are available at both locations. Does it look as though there is any dependence between the extreme levels?

**Solution:**
For the `fremantle` data:

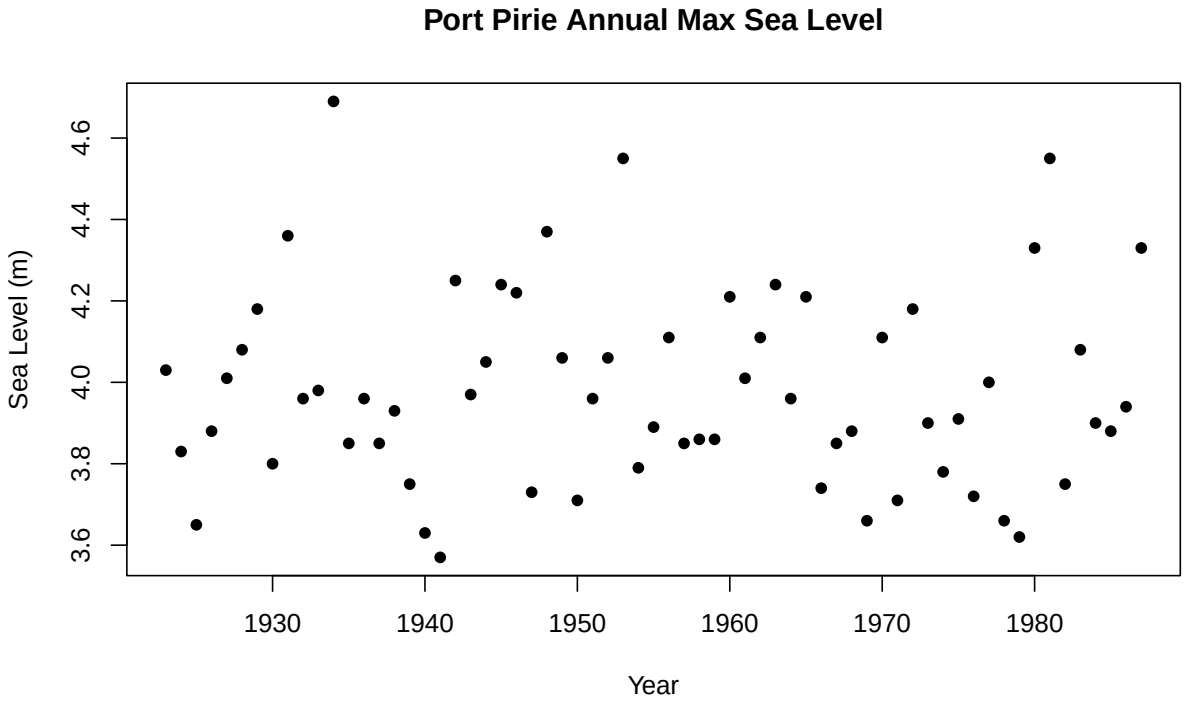
```
library(ismev)
data(fremantle)
fre <- fremantle

plot(fre$Year, fre$SeaLevel, pch=16, xlab="Year", ylab="Sea Level (m)",
     main="Fremantle Annual Max Sea Level")
```




For the `portpirie` data:

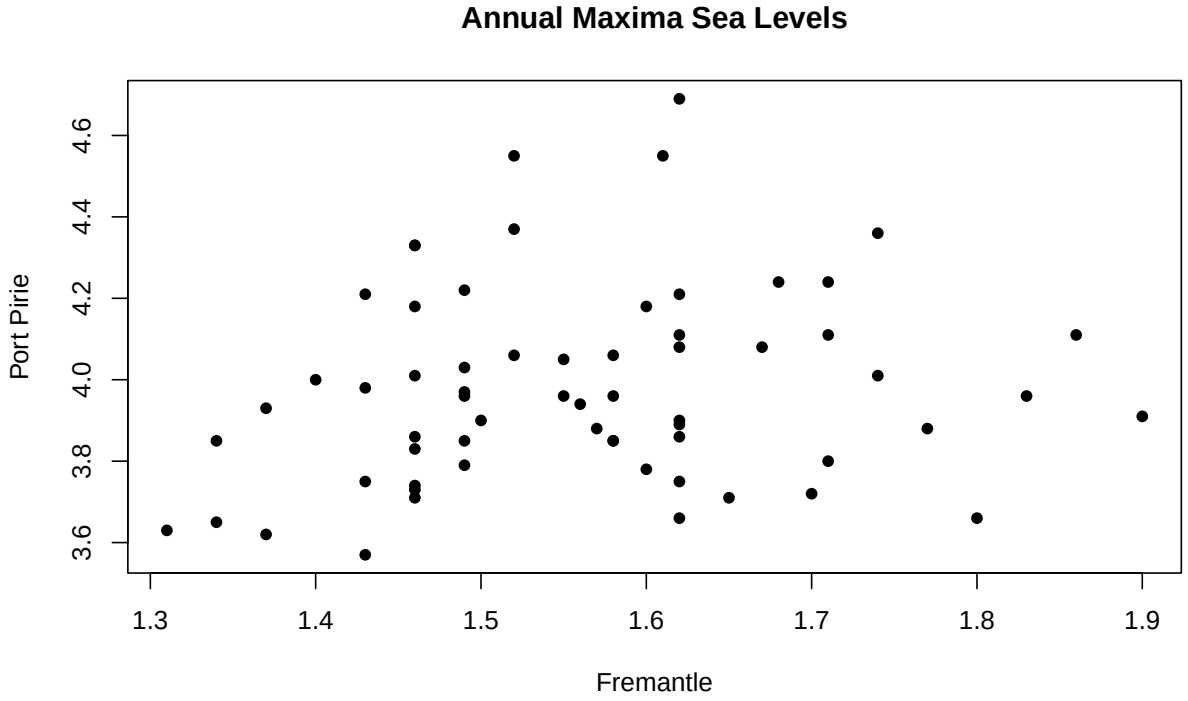
```
data(portpirie)
pp <- portpirie
plot(pp$Year, pp$SeaLevel, pch=16, xlab="Year", ylab="Sea Level (m)",
     main="Port Pirie Annual Max Sea Level")
```




Putting both data together

```r
u <- unique(c(fre$Year,pp$Year))
out <- NULL
for (i in 1:length(u)) { # is there a better way to do this?
  temp <- c(u[i], fre$SeaLevel[fre$Year==u[i]], pp$SeaLevel[pp$Year==u[i]])
  if(length(temp)==3) out <- rbind(out,temp)
}
dat <- data.frame(year=out[,1], fre=out[,2], pp=out[,3])

plot(dat[,2], dat[,3], pch=16, xlab="Fremantle", ylab="Port Pirie",
     main="Annual Maxima Sea Levels")
```




Over the 63 common years of data it seems that there may be some weak dependence between the two locations, though it doesn't appear to be particularly strong.

b) The `evd` package allows for fitting of various types of bivariate extreme value distributions, while also allowing linear modelling of marginal parameters. It works slightly differently to the `extRemes` package, however.

Using the function `fbvevd()` fit a (symmetric) logistic function to the bivariate annual maxima with constant (GEV) marginal parameters. Based on the estimate of $\alpha$, does extremal dependence under this model appear to be strong? Formally test this by fitting the model with $r = 1$ and using the `anova` function. Explore the model fit diagnostics via `plot(...)`.

*Note:* For the `evd` package, the distribution function of the bivariate logistic model is given in the form:

$$G(z_1, z_2) = \exp\{-[z_1^{1/r} + z_2^{1/r}]^r\},$$

which differs from the parameterisation in the lecture notes, which has $r = 1/\alpha$.

**Solution:**

```r
library(evd)
# Note this overwrites "portpirie" with a poorer version of the dataset
fit <- fbvevd(dat[,2:3], model="log")
fit
```

```
##
## Call: fbvevd(x = dat[, 2:3], model = "log")
## Deviance: -90.57894
## AIC: -76.57894
## Dependence: 0.1270806
##
## Estimates
##     loc1    scale1    shape1      loc2    scale2    shape2       dep
##  1.50786   0.11682  -0.12191   3.87064   0.19726  -0.01669   0.90529
##
## Standard Errors
##    loc1   scale1   shape1     loc2   scale2   shape2      dep
## 0.01662  0.01160  0.09971  0.02829  0.02036  0.10685  0.09855
##
## Optimization Information
##   Convergence: successful
##   Function Evaluations: 72
##   Gradient Evaluations: 13
```

```r
fit2 <- fbvevd(dat[,2:3], model="log", dep=1)
fit2
```

```
##
## Call: fbvevd(x = dat[, 2:3], model = "log", dep = 1)
## Deviance: -89.67481
## AIC: -77.67481
## Dependence: 0
##
## Estimates
##     loc1    scale1    shape1      loc2    scale2    shape2
##  1.50890   0.11741  -0.14908   3.87112   0.19747  -0.04292
##
## Standard Errors
##    loc1   scale1   shape1     loc2   scale2   shape2
## 0.01663  0.01181  0.09269  0.02831  0.02055  0.10017
##
## Optimization Information
##   Convergence: successful
##   Function Evaluations: 9
##   Gradient Evaluations: 1
```

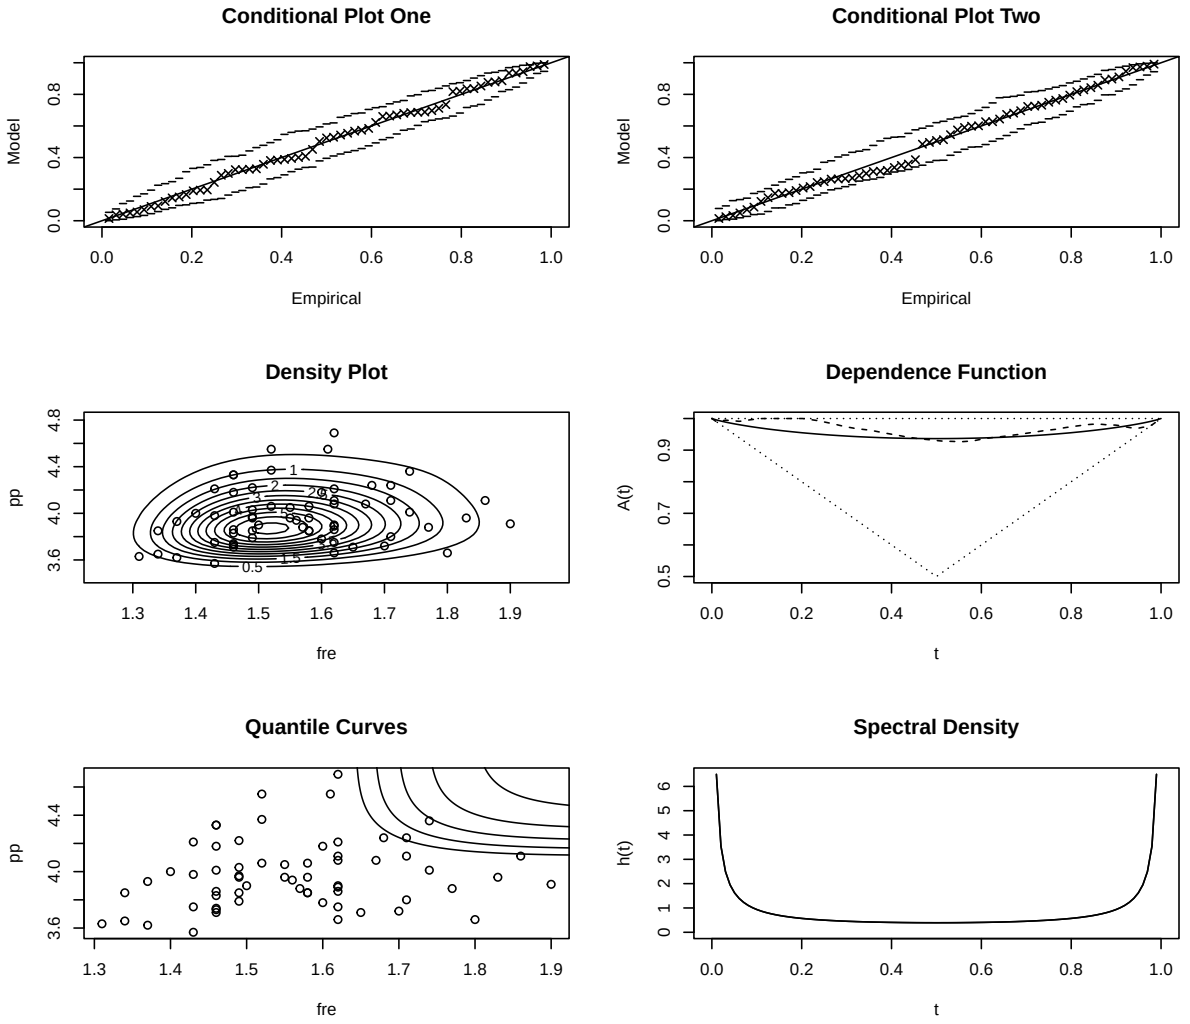
```r
anova(fit, fit2)
```

```
## Analysis of Deviance Table
```

```
##
##      M.Df Deviance Df  Chisq Pr(>chisq)
## fit     7  -90.579
## fit2    6  -89.675  1 0.9041     0.3417
```

The estimated dependence parameter is $\hat{r} = 0.90528882$ with a standard error of 0.09855 ($r = 1$ still corresponds to independence, but $r \to 0$ indicates complete dependence). This indicates pretty weak dependence. The anova test indicates ($p$-value=0.3417) that there is little evidence $r$ (or $\alpha = 1/r$) is different from 1.

```
par(mfrow=c(3,2))
plot(fit)
```




The first plot shows $qq$-plots for each margin (both fits look ok, though with some deviations). The second shows a fitted density plot, plus an estimate of Pickands' dependence function (the dashed line seems to be an empirical estimate with the solid curved line the model estimate). The final plot shows quantiles of the fitted upper model quantiles and the spectral density. All bivariate measures lend support to very weak (if any) dependence.

c) Define $\bar{x}_1$ and $\bar{x}_2$ to be the sample means of the Fremantle ($X_1$) and Port Pirie ($X_2$) annual maxima sea levels from the joint dataset.

i) Calculate $\mathbb{P}(X_1 \leq \bar{x}_1, X_2 \leq \bar{x}_2)$, the joint probability of simultaneously observing annual maxima sea levels at the two locations below their respective average levels.

**Solution:**

```
p <- fit$estimate
z <- c(mean(dat[,2]), mean(dat[,3]))
aa <- pbvevd(z, mar1=p[1:3], mar2=p[4:6], dep=p[7], model="log")
aa
```

```
##       dep
## 0.324973
```

The joint probability is 0.324973.

ii) Calculate $\mathbb{P}(X_1 \leq \bar{x}_1, X_2 \leq \bar{x}_2) = \mathbb{P}(X_1 \leq \bar{x}_1)\mathbb{P}(X_2 \leq \bar{x}_2)$, the same probability but assuming complete independence between the two locations. (You can use the marginal model parameter estimates from fitting the joint model.) How does this compare to the empirical estimates of this probability?

**Solution:**

```
bb <- pgev(z[1], loc=p[1], scale=p[2], shape=p[3])
bb
```

```
##    shape1
## 0.5373122
```

```
cc <- pgev(z[2], loc=p[4], scale=p[5], shape=p[6])
cc
```

```
##    shape2
## 0.5604439
```

```
bb*cc
```

```
##    shape1
## 0.3011334
```

The independent probabilities give $0.5373122 \times 0.5604439 = 0.3011334$.

iii) Can you reconcile the two answers?

**Solution:**

```
# compare aa to bb*cc
sum((dat[,2]<z[1]) & (dat[,3]<z[2]))/length(dat[,1])
```

```
## [1] 0.3015873
```

We have joint probability > product of independent probabilities. This makes sense as with positive dependence we have a greater chance of events occurring in joint lower or joint upper tail regions. (Thereby justifying multivariate modelling!) The empirical

estimate is 0.3015873 which is closer to the independence estimate, but there are very few data points here, so this estimate is quite variable.

d) For $(x_1^*, x_2^*) = (1.8, 4.5)$ what is the probability $\mathbb{P}(X_1 > x_1^*, X_2 > x_2^*)$ assuming i) the fitted joint model, and ii) independence between the two locations. Compare this to the empirical estimate and reconcile your answers.

**Solution:**

```
z <- c(1.8,4.5)
aa <- pbvevd(z, mar1=p[1:3], mar2=p[4:6], dep=p[7], model="log")
bb <- pgev(z[1], loc=p[1], scale=p[2], shape=p[3])
cc <- pgev(z[2], loc=p[4], scale=p[5], shape=p[6])
1-bb-cc+aa # joint tail probability
```

```
##      shape1
## 0.006892943
```

```
(1-bb)*(1-cc) # independence tail probability
```

```
##      shape1
## 0.001826453
```

```
sum((dat[,2]>z[1]) & (dat[,3]>z[2]))/length(dat[,1]) # empirical
```

```
## [1] 0
```

The joint probability is 0.006892943 and the independence tail probability is 0.001826453, and so again we are underestimating the joint tail probability. The empirical estimate is zero, highlighting that extreme value theory allows us to credibly extrapolate into regions where we don't have any observed data (even though each marginal value of $x^*$ is not the maximum marginally. I.e. jointly the event is extreme, but marginally the events are merely large.). Ignoring (positive) dependence between extremes results in underestimating the chance of future extremes, which, as we might imagine, might have undesirable consequences!

**Exercise 2**

Often in an analysis we are interested in the extremal behaviour of a *structure variable*, which is some function of the observed variables, e.g. $Z = f(X, Y)$, where we have observations $(X_i, Y_i)$, $i = 1, \ldots, n$ and where $f$ is some known function of interest. There are two ways that we can model the extremal behaviour of Z:

- *Univariate:* For each pair $(X_i, Y_i)$ we calculate $Z_i = f(X_i, Y_i)$. This produces a univariate dataset $\{Z_i\}$ which we then analyse using univariate models for extremes.
- *Multivariate:* We analyse $(X, Y)$ using multivariate models for extremes, and then use the multivariate model to predict for the univariate function $f$ we are interested in.

Consider again the joint Fremantle/Port Pirie annual maximum sea level dataset from Exercise 1. Suppose we are interested in the structure variable

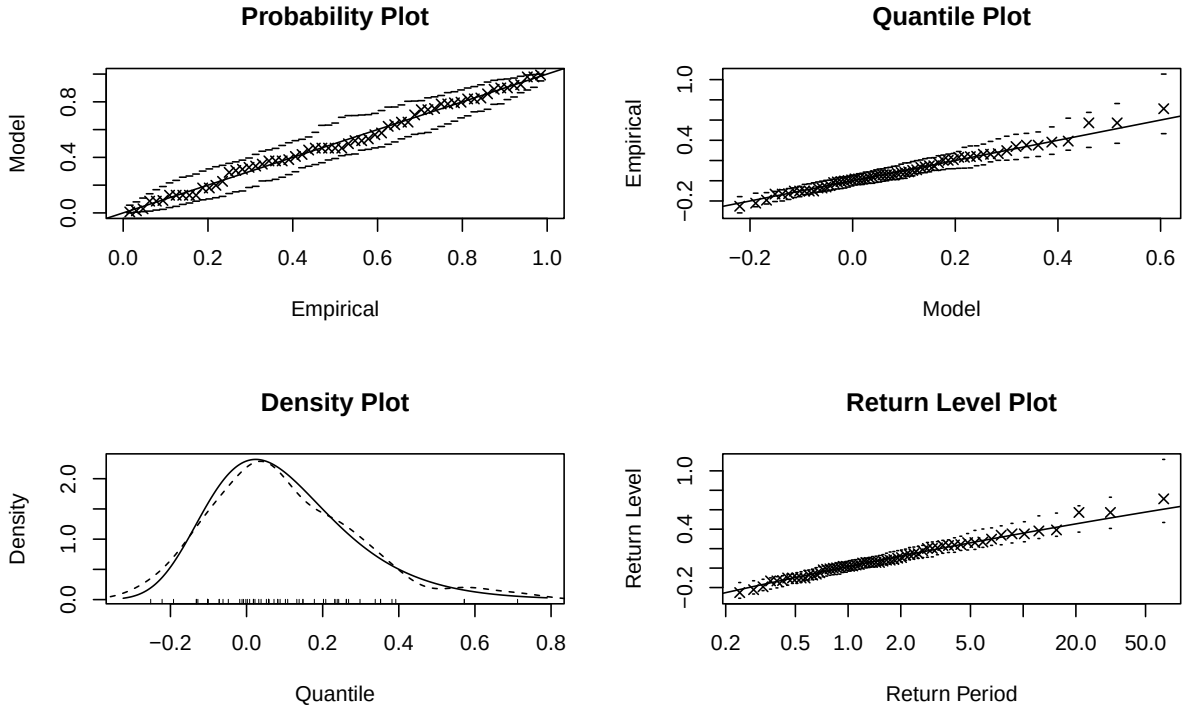
$$Z = \max\{X - \bar{X}, Y - \bar{Y}\}$$

which represents the largest exceedance of the local mean measurement at either location.

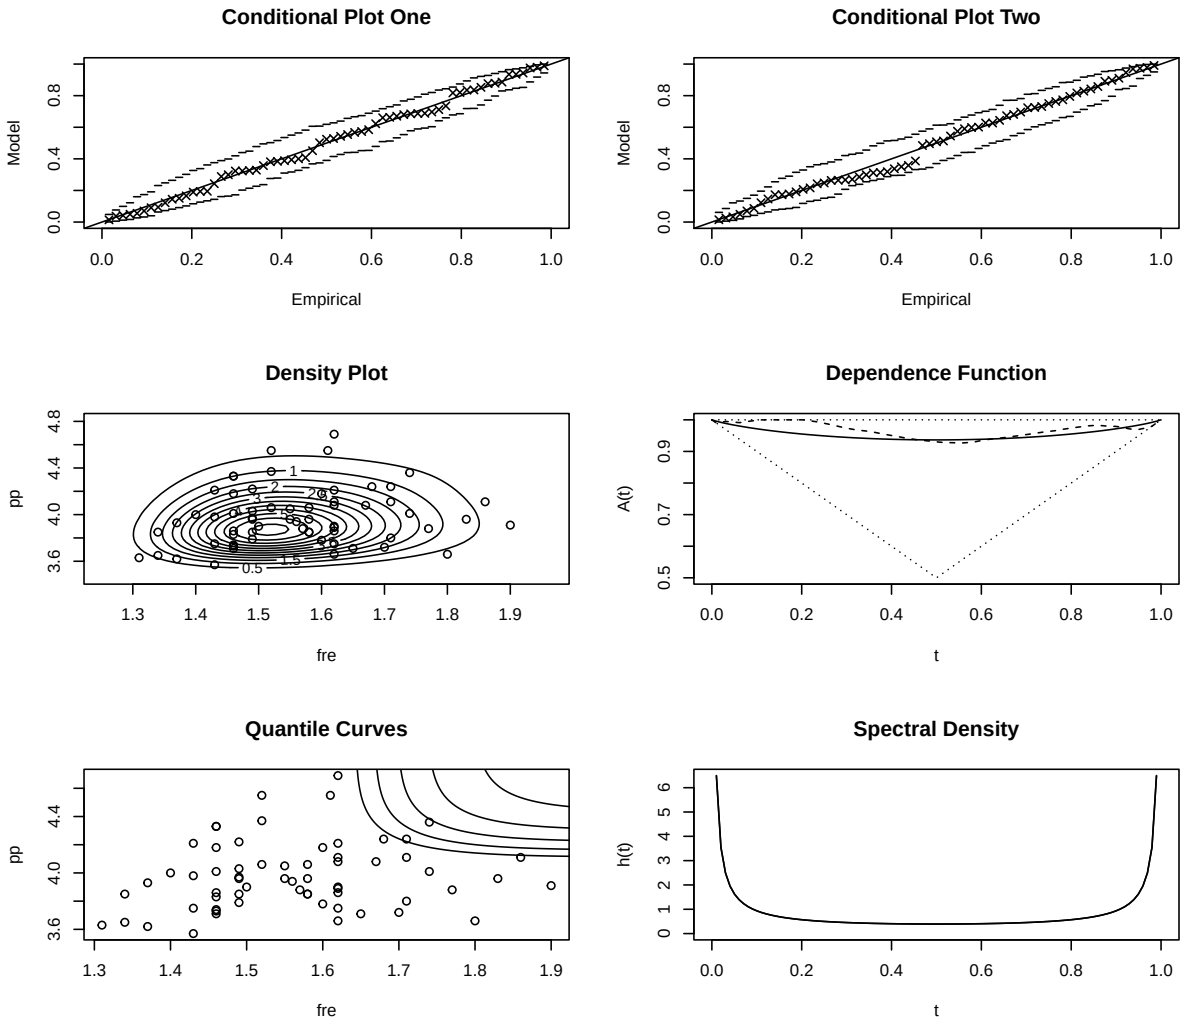
a) Construct the univariate samples for $Z_1, \ldots, Z_{63}$ and fit a (univariate) GEV model with constant parameters. Also fit a bivariate logistic model for the joint maxima $(X_i, Y_i)$. Assess the fit of both of these models.

**Solution:**

```
# Run the code from Q1(a)
datz <- dat[,2:3]-cbind(rep(mean(dat[,2]),63), rep(mean(dat[,3]),63))
fit <- fbvevd(dat[,2:3], model="log")
p <- fit$estimate
Z <- apply(datz,1,max)
par(mfrow=c(2,2))
fit.Z <- fgev(Z)
plot(fit.Z)
```




```
par(mfrow=c(3,2))
plot(fit)
```




The model fit for $Z$ doesn't look perfect - and we know that there is some time-variability in the GEV parameters for Fremantle from our previous analyses - but the current fit isn't terrible. For the bivariate analysis, our conclusions are the same as for Exercise 1: the fit is reasonable, although we do know that, at least for Fremantle, some of the marginal GEV model parameters can be expressed as functions of time or other predictors (SOI). So we might imagine that we can improve both univariate and bivariate modelling here, but for now we keep things simple.

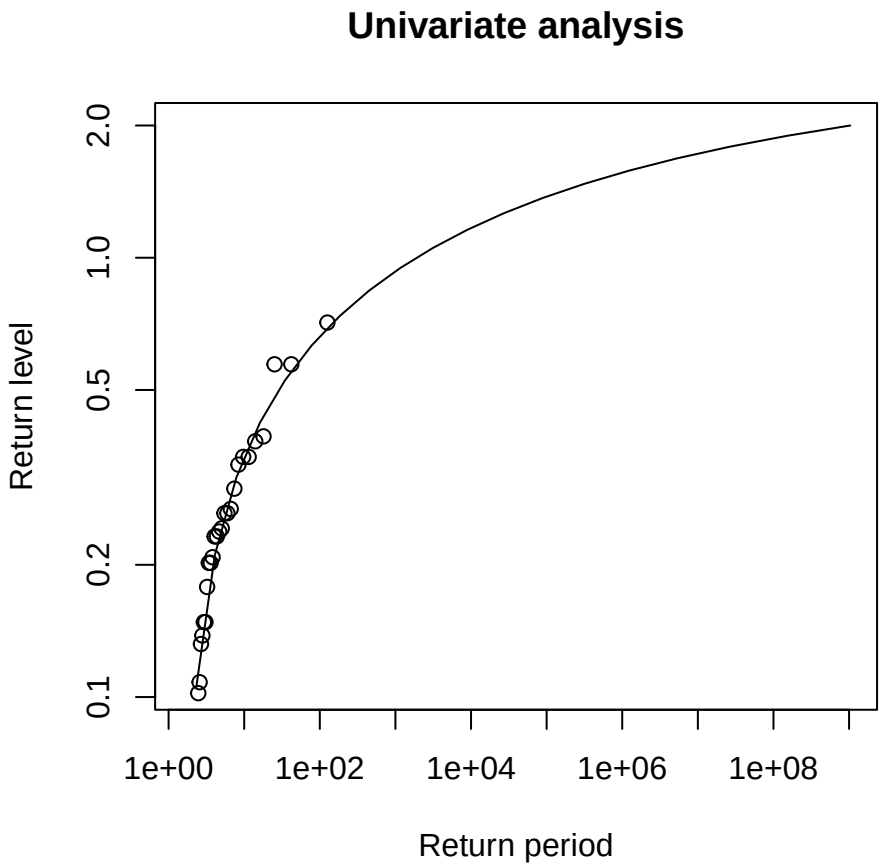
b) For the univariate analysis, construct a return period plot (and superimpose empirical estimates) for the structure variable return levels between 0 and 2.

**Solution:**

```
p.Z <- fit.Z$estimate
r <- seq(0, 2, length=20)
rp.Z <- 1/(1-pgev(r, loc=p.Z[1], scale=p.Z[2], shape=p.Z[3]))

plot(rp.Z, r, log="xy", type='l', xlab="Return period", ylab="Return level",
     main="Univariate analysis")
```

```
L <- length(Z)
points(1/(1-(1:L-0.5)/L), sort(Z))
```




Here we use the `pgev` function from the `evd` package for the cdf of the GEV distribution. The fitted GEV model appears to capture the empirical data well. Note that on this log-log scale plot the return level is concave. This reflects that we have a negative shape parameter $\hat{\xi} = -0.05389817$. (Such plots would be convex for Fréchet distributions, and linear for Gumbel.)

c) For the multivariate analysis, similarly construct a return period plot (with superimposed empirical estimates) for the structure variable return levels between 0 and 2.

*Hint:* We can write the distribution function of $Z$ in terms of $(X, Y)$ as

$$\mathbb{P}(Z \leq z) = \int_{A_z} g(x, y) dx dy$$

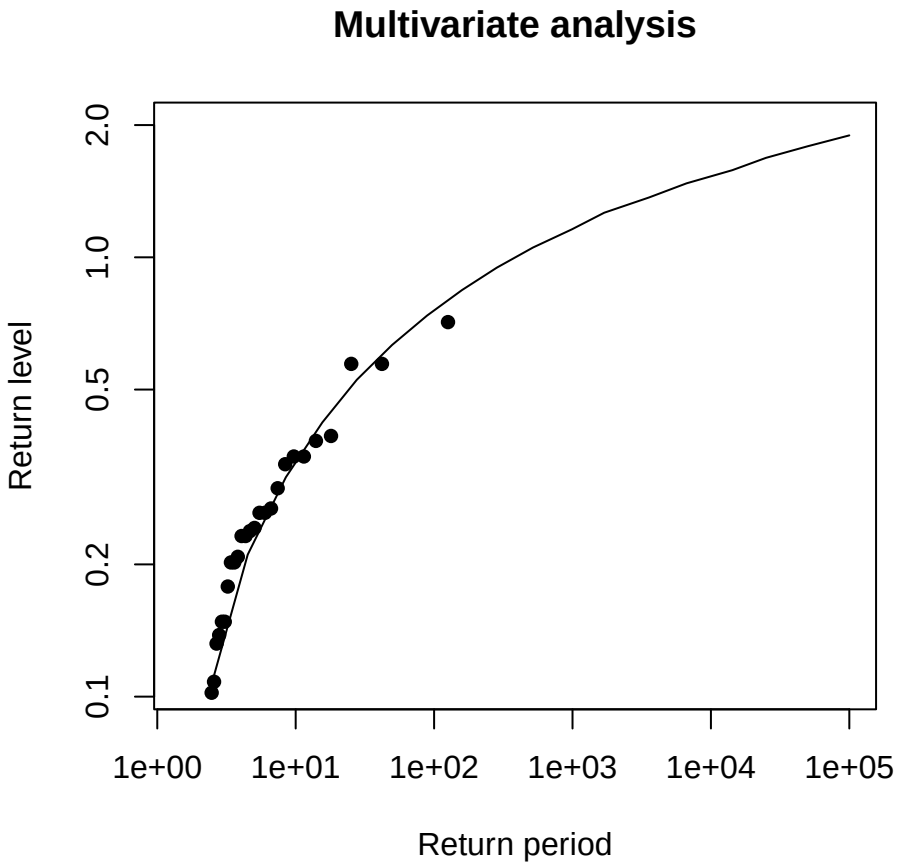
where $g(x, y)$ is the joint model for $(X, Y)$, and $A_z = \{(x, y) : f(x, y) \leq z\}$. When the integral is too difficult analytically, we can aways approximate it via Monte Carlo simulation so that e.g.

$$\int_{A_z} g(x, y) dx dy \approx \frac{1}{M} \sum_{i=1}^{M} I\left(f(x^{(i)}, y^{(i)}) \leq z\right),$$

where $(x^{(i)}, y^{(i)}) \sim g(x, y)$ are $M$ iid samples from $g(x, y)$, $I(a)$ is the indicator function equal to 1 if $a$ is true and 0 otherwise, and $M$ is large (depending on the accuracy required).

**Solution:**

```
M <- 100000
xx <- rmvevd(M, dep=p[7], model="log", d=2, mar=list(p[1:3],p[4:6]))
zz <- xx - cbind(rep(mean(dat[,2]),M), rep(mean(dat[,3]),M))
mm <- apply(zz, 1, max)
rp.joint <- rp.Z*NA
for (i in 1:length(rp.joint)) rp.joint[i] <- M/sum(mm>r[i])
plot(rp.joint, r, log="xy", type='l', xlab="Return period",
     ylab="Return level", main="Multivariate analysis")
points(1/(1-(1:L-0.5)/L), sort(Z), pch=16)
```

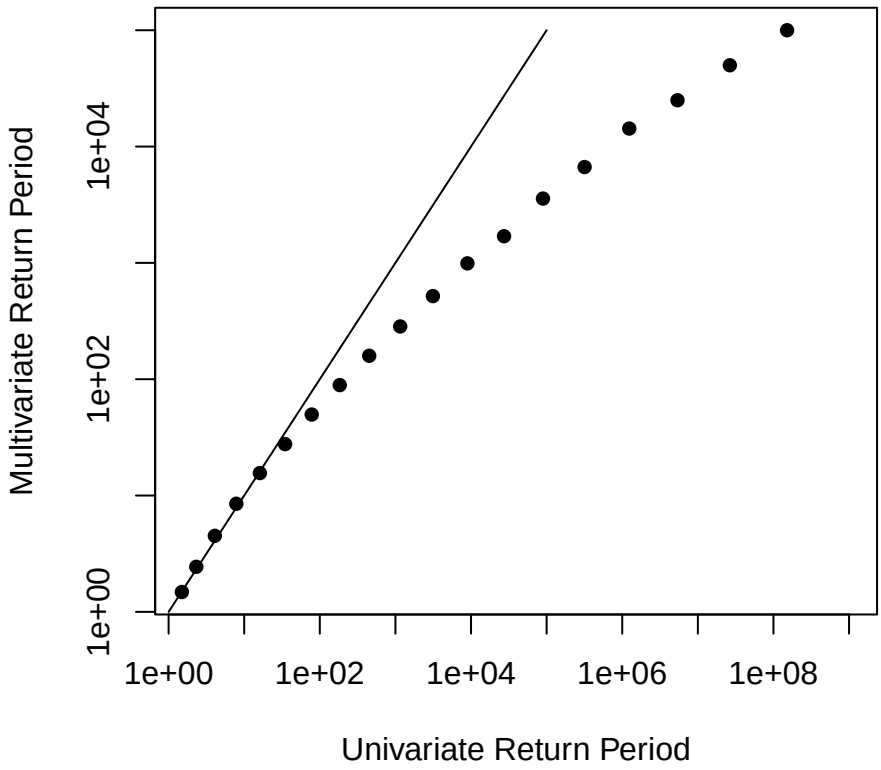



The return period/level plot is similar to before and is broadly consistent with the empirical estimates, though with some slight divergence for moderate values.

d) How do the two return period plots compare?

**Solution:**

```
plot(rp.Z, rp.joint, log="xy", pch=16, xlab="Univariate Return Period",
     ylab="Multivariate Return Period")
lines(c(1,M),c(1,M))
```




The return periods are equivalent for moderate extremes, but diverge for larger extremes, with the univariate approach giving larger return periods than the multivariate approach. (That is, the univariate model gives less weight to larger extremes than the multivariate model. It's difficult to say which return levels are preferable from these plots, but we can discuss the pros and cons of the two approaches below.

e) What are the advantages and disadvantages of the two different approaches to structure variable modelling?

**Solution:**

- *Modelling complexity:* Modelling univariate extremes is much simpler than modelling multivariate extremes. There is only a single limiting distribution for univariate extremes (GEV, GPD), whereas there are many parametric forms for multivariate extremes that must be navigated. Univariate models also require estimation of less parameters, which can improve statistical efficiency for smaller datasets. Inferences are also more complicated in the multivariate setting (e.g. we needed to use Monte Carlo methods for return periods, or composite likelihood techniques for fitting moderate-dimensional models).
- *Data use efficiency:* In the univariate approach a multivariate vector is collapsed down to a single observation, which results in a loss of information. Working with the full vector of observed data makes use of the full dataset and can result in more precise inferences. Relatedly, here we have computed $Z$ as a "maxima" based on a function of two other maxima, $X$ and $Y$. If the original underlying data $(\tilde{X}, \tilde{Y})$ were available (prior to calculating the maxima $X$ and $Y$), then calculating the observation-level $\tilde{Z} = f(\tilde{X}, \tilde{Y})$ may result in $\tilde{Z}$ in a block that are larger than the $Z = f(X, Y)$ that was computed from the marginal bloc maxima $X$ and $Y$. This would be undesirable!

- *Marginal data heterogeneity:* When there is heterogeneity in the data that is explainable by predictors, it may be simpler to account for the variation in each margin separately, than within some complex function of data from multiple margins. E.g. suppose there is strong predictor influence in one margin (e.g. Fremantle), and no predictors informing for another margin (e.g. Port Pirie). Addressing this in the Fremantle margin would be simpler than trying to resolve this in the conflated structure variable.
- *Use of missing data:* In the above analysis we could only compute $Z$ if both $X$ and $Y$ were observed, which makes the univariate analysis sensitive to (partially) missing data where only one of $X$ or $Y$ is observed. The multivariate approach can make use of partially observed data by (conceptually) using all $X$ data to estimate the parameters in the $X$ margin, similarly for $Y$, and perfectly observed $(X, Y)$ pairs for the dependence parameters.
- *Change of structure variable:* By modelling the multivariate data, the same fitted model can be used for any structure variable with minimal changes. Conversely, if the structure variable changes, the univariate approach must perform the full analysis again from the beginning.
- *Accounting for uncertainty:* The multivariate approach may be better for capturing uncertainty in the data and model than the univariate approach, and propagating this through to the desired inference. Wider confidence intervals are not necessarily worse if they better account for true uncertainty! Better uncertainty modelling can come through more efficient use of data and greater modelling flexibility, as discussed above.

**Exercise 3**

Consider again the historical AUD/xxx daily exchange rates data from, given in the file `AUD-exchangeRates.csv`. In the Week 3 and 4 Labs, we analysed the univariate extremes of the log daily returns of the AUD/USD (column FXRUSD) series. Now we will explore the joint extremes of the 3-dimensional log daily returns of the AUD/(USD, CNY, JPY) exchange rates (columns FXRUSD, FXRCR, FXRJY).

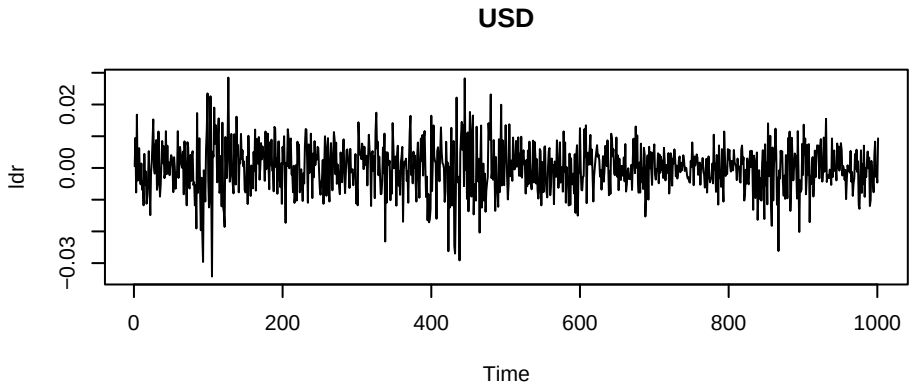
a) Compute and plot the log daily returns of each exchange rate, and produce pairwise scatter plots. What are the dependencies in each currency?

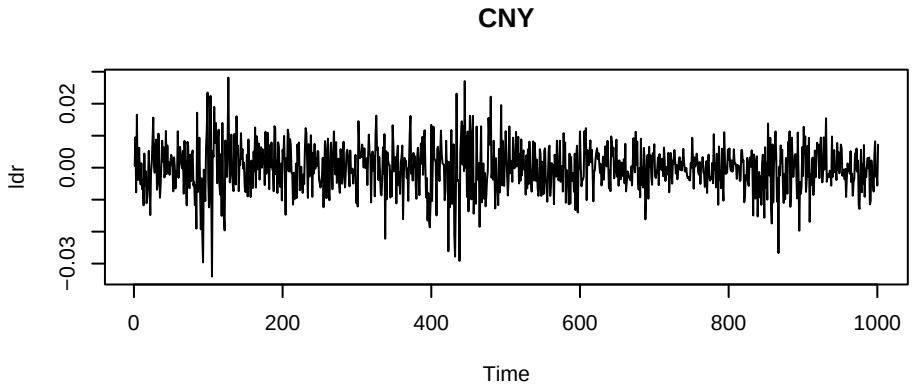
**Solution:**

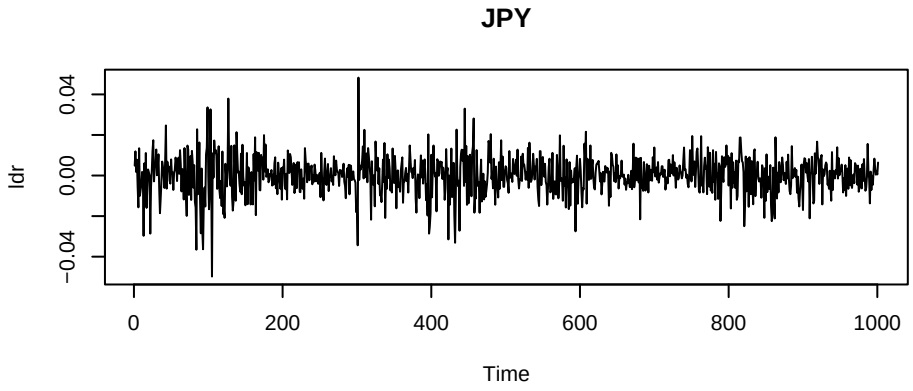
```
dat <- read.table("../data/AUD-exchangeRates.csv",header=TRUE,sep=",")
usd <- dat$FXRUSD
cny <- dat$FXRCR
jpy <- dat$FXRJY
L <- length(usd)
ldrusd <- log(usd[2:L]/usd[1:(L-1)])
ldrcny <- log(cny[2:L]/cny[1:(L-1)])
ldrjpy <- log(jpy[2:L]/jpy[1:(L-1)])

par(mfrow=c(3,1))
plot(ldrusd, type='l', main="USD", xlab="Time", ylab="ldr")
```

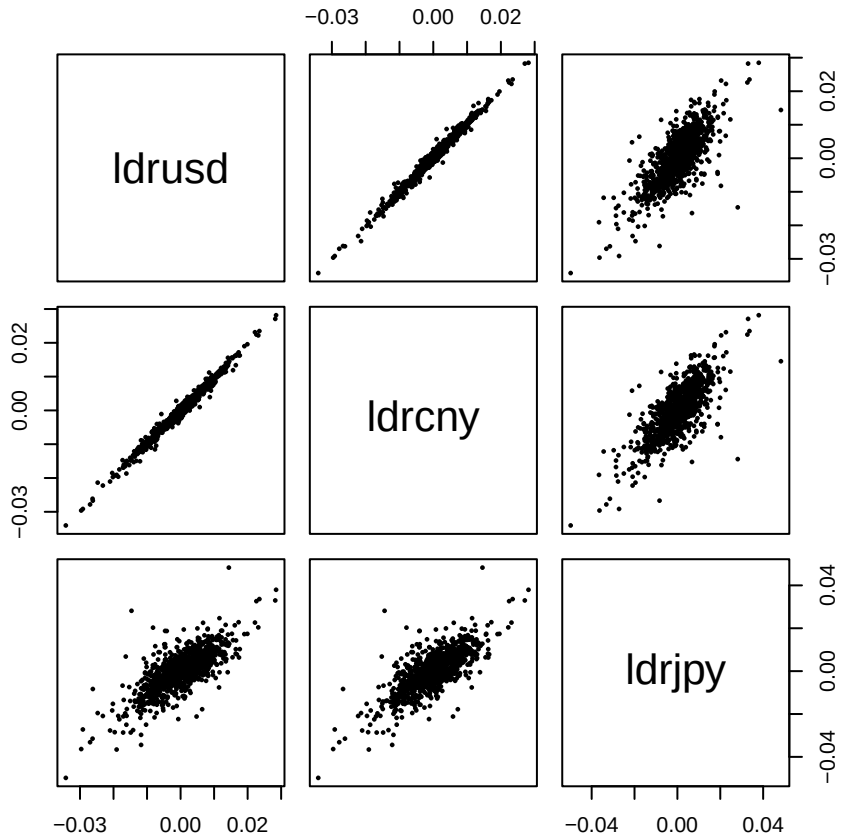
```
plot(ldrcny, type='l', main="CNY", xlab="Time", ylab="ldr")
plot(ldrjpy, type='l', main="JPY", xlab="Time", ylab="ldr")
```




```
pairs(cbind(ldrusd,ldrcny,ldrjpy), pch=16, cex=0.5)
```




The USD and CNY exchange rates are very highly dependent, and the dependence between all other currently pairs is positive and reasonably strong.

b) We previously considered that blocks of size $n = 21$ (monthly observations) were sufficiently large for the GEV limit for block maxima to approximately hold. Construct the block maxima and again produce time series and scatter plots of the data. Has the structure changed?

**Solution:**

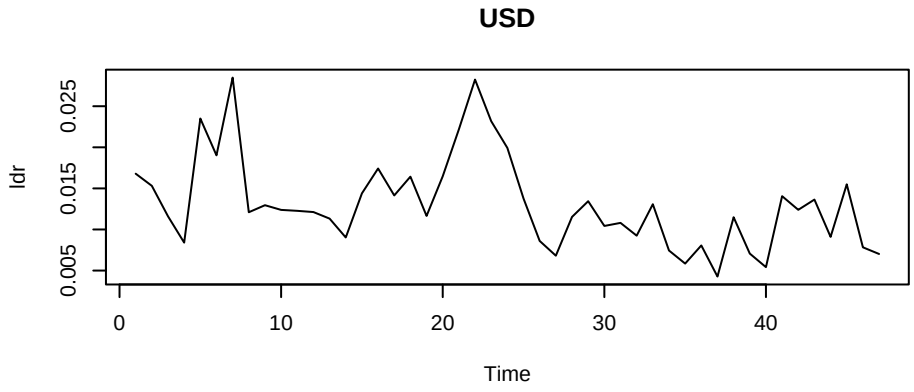
```
nn=21 # approximate no. of data per (monthly) block
temp=matrix(ldrusd,ncol=nn,byrow=T)
temp=temp[-dim(temp)[1],] # need to remove last row
bm.usd=apply(as.matrix(temp),1,max)

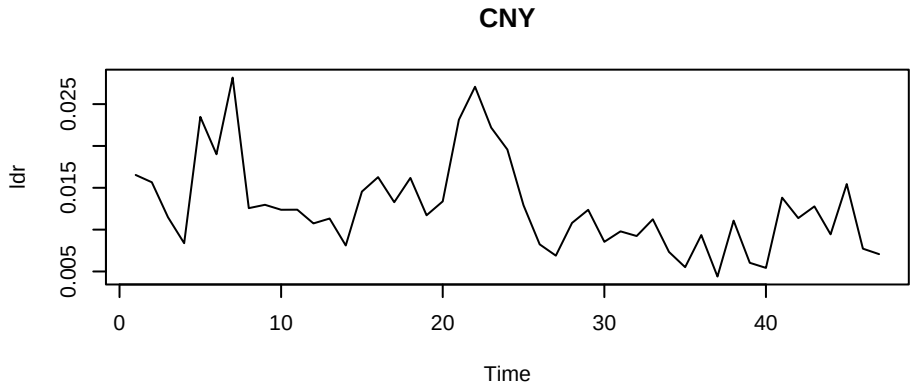
temp=matrix(ldrcny,ncol=nn,byrow=T)
temp=temp[-dim(temp)[1],] # need to remove last row
bm.cny=apply(as.matrix(temp),1,max)

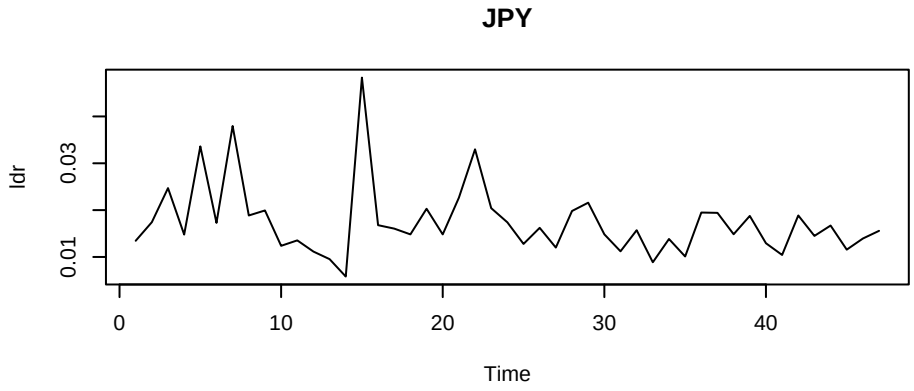
temp=matrix(ldrjpy,ncol=nn,byrow=T)
temp=temp[-dim(temp)[1],] # need to remove last row
bm.jpy=apply(as.matrix(temp),1,max)

par(mfrow=c(3,1))
plot(bm.usd,type='l',main="USD",xlab="Time",ylab="ldr")
```

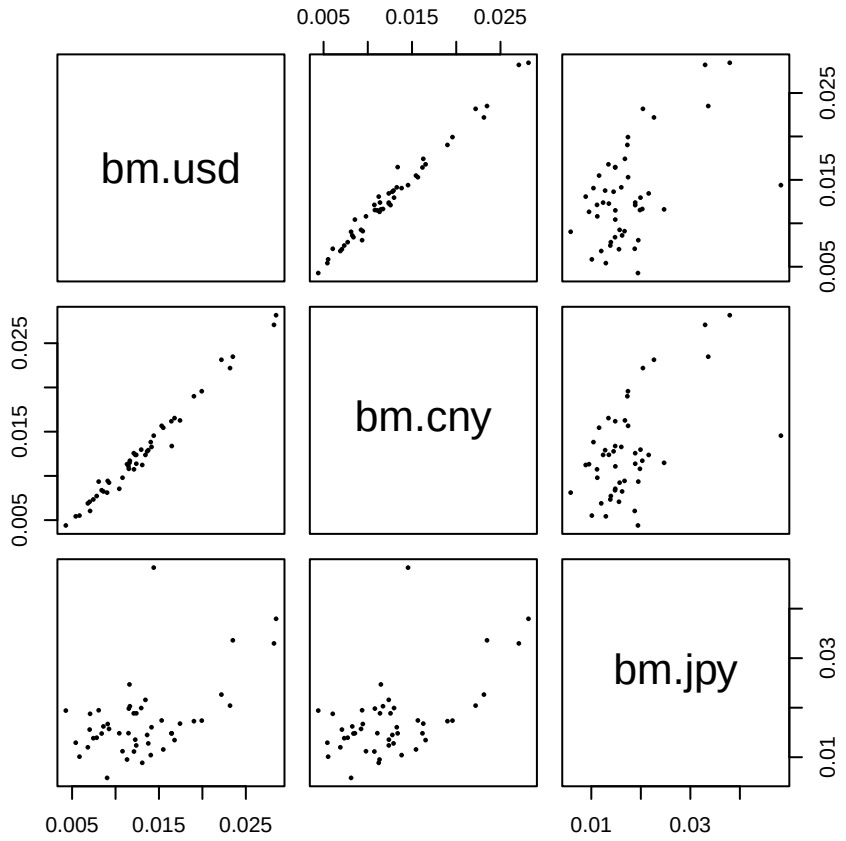
```
plot(bm.cny,type='l',main="CNY",xlab="Time",ylab="ldr")
plot(bm.jpy,type='l',main="JPY",xlab="Time",ylab="ldr")
```




```
bm.all=cbind(bm.usd,bm.cny,bm.jpy)
pairs(bm.all,pch=16,cex=0.5)
```




Obviously there has been a large reduction in the amount of data, but we can be reasonably confident that an extreme value distribution can be used to model these block maxima. The very strong dependence between USD and CNY is still present, and there is a reasonable chance that some of the (USD,CNY) block maxima were actually observed at the same time point in the original series. The dependence between the other currency pairs perhaps looks a little weaker, though its not clear due to the small amount of data. There are likely very few block maxima pairs that correspond to observations in the original series.

c) Based on marginal GEV fits, transform the data to unit Fréchet margins, and view scatterplots again.

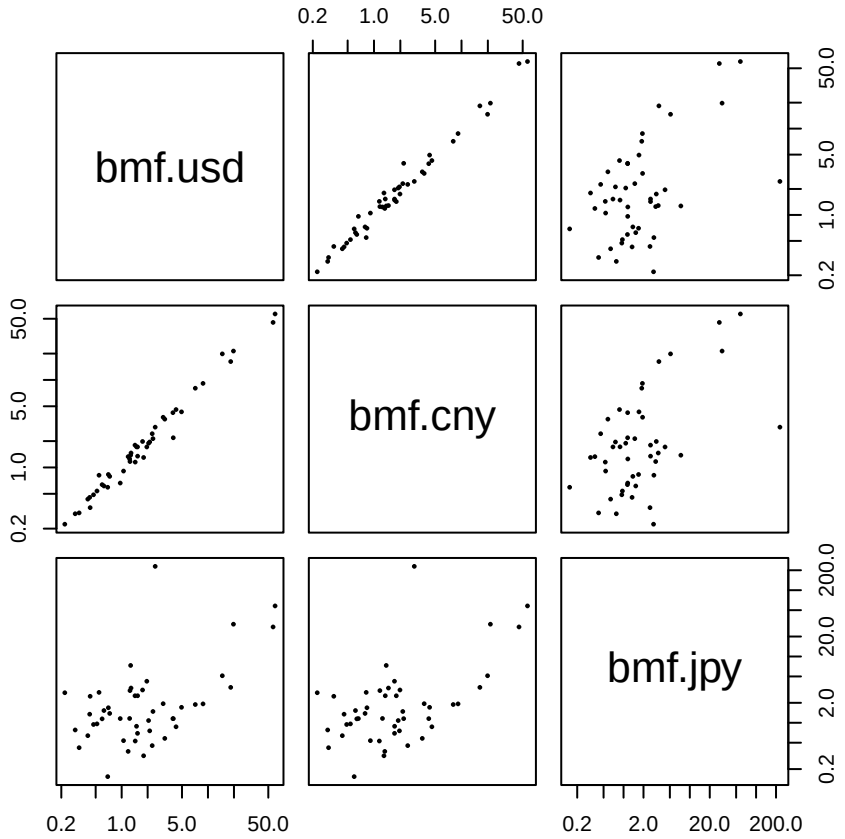
**Solution:**

```
library(ExtremalDep)
library(extRemes)

fit <- fevd(bm.usd, type="GEV")
bmf.usd <- trans2UFrechet(data=bm.usd, pars=fit$results$par, type="GEV")

fit <- fevd(bm.cny, type="GEV")
bmf.cny <- trans2UFrechet(data=bm.cny, pars=fit$results$par, type="GEV")
fit <- fevd(bm.jpy, type="GEV")
bmf.jpy <- trans2UFrechet(data=bm.jpy, pars=fit$results$par, type="GEV")
```

```
bmf.all <- cbind(bmf.usd, bmf.cny, bmf.jpy)
pairs(bmf.all,pch=16,cex=0.5,log="xy")
```




On the log scale, while the scales have changed, the data don't look too different from the original block maxima.

d) Fit $d = 3$-dimensional Hüsler-Reiss and Extremal-$t$ models to these block maxima (use the command `fExtDep` with argument `method="Composite"` in the `ExtremalDep` package). Produce plots of Pickands dependence function for each pair of dimensions for each model (see Week 7 Lab). Do the fitted models appropriately capture the dependence in the data?

**Solution:**

```
p0 <- c(rep(0.5,3),1)
fit.et <- fExtDep(method="Composite", data=bmf.all, model="ET", par.start=p0,trace=2)
```

```
## initial  value 564.100065
## iter  10 value 506.420025
## iter  20 value 504.402750
## final  value 504.237277
## converged
```

```
fit.et
```

```
## $par
## [1] 0.996 0.581 0.616 3.357
```

```
##
## $LL
## [1] -504.237
##
## $SE
## [1] 0.001 0.144 0.126 4.491
##
## $TIC
## [1] 1041.218
```

```
p1 <- rep(0.5,3)
fit.hr <- fExtDep(method="Composite", data=bmf.all, model="HR", par.start=p1,trace=2)
```

```
## initial  value 578.139462
## iter  10 value 505.207946
## iter  10 value 505.207941
## final  value 505.207941
## converged
```

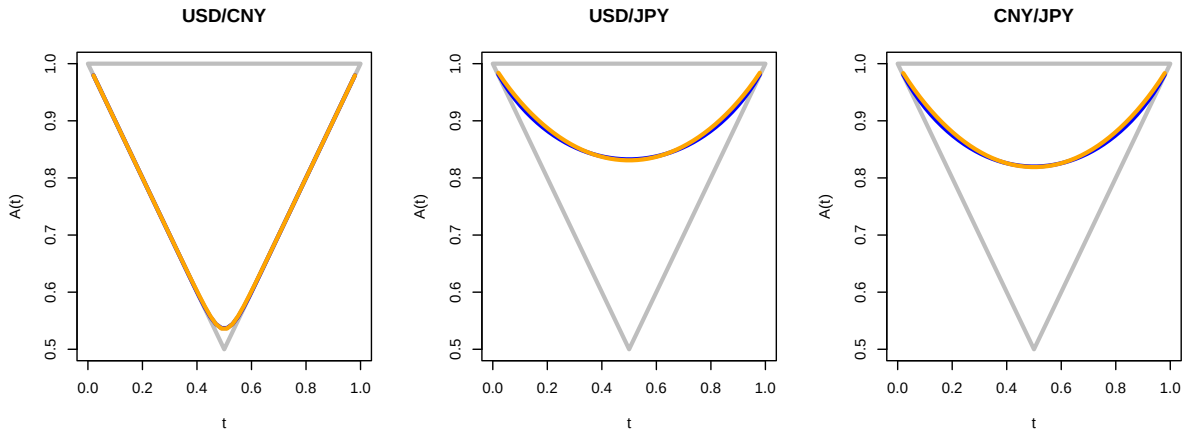
```
fit.hr
```

```
## $par
## [1] 0.091 0.964 0.915
##
## $LL
## [1] -505.208
##
## $SE
## [1] 0.013 0.136 0.118
##
## $TIC
## [1] 1018.377
```

```
par(mfrow=c(1,3))
W <- simplex(d=2)
plot(NA, NA, xlim=c(0,1), ylim=c(0.5, 1), ylab="A(t)", xlab="t",
     main="USD/CNY")
polygon(c(0, 0.5, 1), c(1, 0.5, 1), lwd=3, border = 'grey')
lines(W[,1], apply(W, 1, index.ExtDep, object="pickands", model="HR",
                   par=fit.hr$par[1]), col="blue", lwd=3)
lines(W[,1], apply(W, 1, index.ExtDep, object="pickands", model="ET",
                   par=fit.et$par[c(1,4)]), col="orange", lwd=3)

plot(NA, NA, xlim=c(0,1), ylim=c(0.5, 1), ylab="A(t)", xlab="t",
     main="USD/JPY")
polygon(c(0, 0.5, 1), c(1, 0.5, 1), lwd=3, border = 'grey')
lines(W[,1], apply(W, 1, index.ExtDep, object="pickands", model="HR",
                   par=fit.hr$par[2]), col="blue", lwd=3)
lines(W[,1], apply(W, 1, index.ExtDep, object="pickands", model="ET",
                   par=fit.et$par[c(2,4)]), col="orange", lwd=3)
```

```
plot(NA, NA, xlim=c(0,1), ylim=c(0.5, 1), ylab="A(t)", xlab="t",
     main="CNY/JPY")
polygon(c(0, 0.5, 1), c(1, 0.5, 1), lwd=3, border = 'grey')
lines(W[,1], apply(W, 1, index.ExtDep, object="pickands", model="HR",
                   par=fit.hr$par[3]), col="blue", lwd=3)
lines(W[,1], apply(W, 1, index.ExtDep, object="pickands", model="ET",
                   par=fit.et$par[c(3,4)]), col="orange", lwd=3)
```



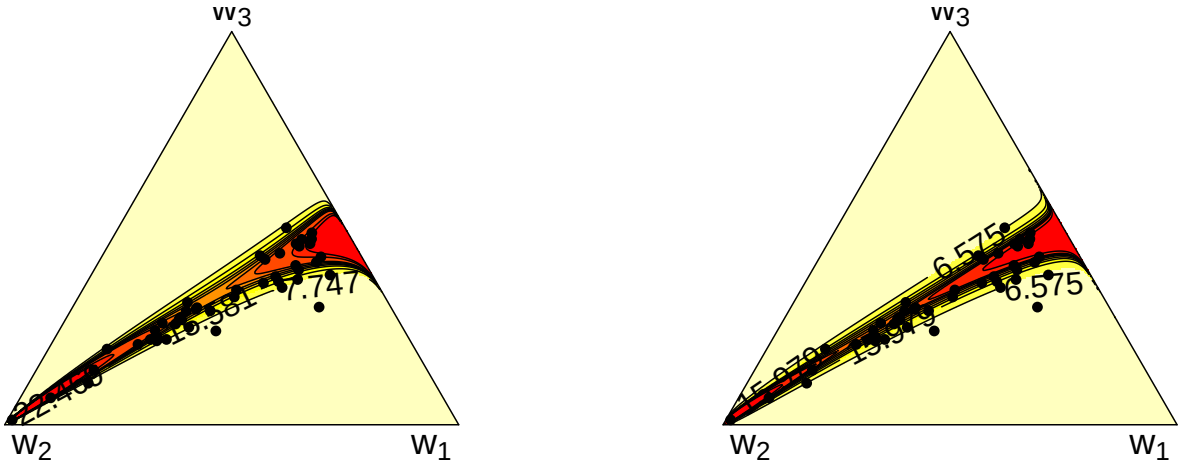

Both models appear to describe the dependence between the currencies similarly. There is moderate dependence between USD/JPY and CNY/JPY and the expected very strong dependence between USD/CNY. The fitted dependencies are very clear: these currencies are not independent.

Note that we have taken the parameters from a 3D model and used them as proxies for the parameter values in 2D model fits. Of course the parameters are not directly transferable in this way, so these plots are only approximate. To make more suitable 2D comparisons we could either integrate out 1 dimension from the 3D model (perhaps the most appropriate), or fit multiple 2D models (which will not correspond to the fitted 3D model).

e) Plot the fitted angular densities $h(\omega)$ for each model. Discuss their features.

**Solution:**

```
bmf.all.ang <- bmf.all / rowSums(bmf.all)
par(mfrow=c(1,2))
plot_ExtDep(model='HR', par=fit.hr$par, log=FALSE,
            cex.lab=1.5, cex.cont=1.3, data=bmf.all.ang)
plot_ExtDep(model='ET', par=fit.et$par, log=FALSE,
            cex.lab=1.5, cex.cont=1.3, data=bmf.all.ang)
```




Both angular densities capture the dependence in the data similarly, although the extremal-$t$ model (right plot) has a narrower (higher) degree of dependence between $(\omega_1, \omega_3)$ (presumably USD/CNY). Apart from this, the density for $(\omega_1, \omega_3)$ is strongly centred around 0.5, although the fit looks slightly asymmetric for the Husler-Reiss (left plot) model. The bivariate angular densities for the other currency pairs have a clear $\cup$-shape, though with some density for $\omega$ around 0.5.

f) Compute the extremal coefficient $1 \leq \theta \leq 3$. (Hint: use the function `index.ExtDep`). What does this say about the dependence in the model?

**Solution:**

```
theta.hr <- index.ExtDep(object="extremal", model="HR",par=fit.hr$par)
theta.hr
```

```
## [1] 1.432671
```

```
theta.et <- index.ExtDep(object="extremal", model="ET",par=fit.et$par)
theta.et
```

```
## [1] 1.703463
```

Both extremal coefficients ($\theta = V(1,1,1)$) are between 1.5 and 1.7. This can be interpreted as the effective number of independent components out of the 3 considered. We know that there is very strong dependence between USD and CNY so there is credibly only a little over 1 independent component there, and we know that there is reasonable dependence between the other pairs, so 1.5-1.7 effective independent components seems credible.

g) The upper 95% quantile of the unit Fréchet is approximately $z = 20$. Using the extremal-$t$ distribution (via the function `pExtDep`), calculate the probability that (on the Fréchet transformed scale):

- $\mathbb{P}(\text{JPY} > 20)$
- $\mathbb{P}(\text{JPY} > 20 \mid \text{CNY} > 20)$
- $\mathbb{P}(\text{JPY} > 20 \mid \text{CNY} > 20, \text{USD} > 20)$

and compare these to the available empirical estimates. Comment on your findings.

**Solution:**

```
z <- rep(20,3)
prob1 <- 1-pevd(z[3], loc=1, scale=1, shape=1, type="Frechet")
prob1
```

```
## [1] 0.04877058
```

```
# For simplicity, fit bivariate model
fit.et2 <- fExtDep(method="Composite", data=bmf.all[,2:3], model="ET",
                   par.start=c(0.5,1), trace=2)
```

```
## initial  value 199.567415
## iter  10 value 195.740426
## iter  20 value 195.584311
## iter  30 value 195.579104
## iter  40 value 195.577191
## final  value 195.576015
## converged
```

```
prob2 <- pExtDep(q=z[2:3], type="upper", model="ET", par=fit.et2$par)/prob1
prob2
```

```
## [1] 0.3780266
```

```
prob3 <- pExtDep(q=z, type="upper", model="ET", par=fit.et$par)/(prob2*prob1)
prob3
```

```
## [1] 0.9489331
```

```
ind <- bmf.all>20
emp1 <- mean(ind[,3])
emp1
```

```
## [1] 0.08510638
```

```
emp2 <- sum(ind[,2]*ind[,3])/sum(ind[,2])
emp2
```

```
## [1] 1
```

```
emp3 <- sum(ind[,1]*ind[,2]*ind[,3]) / sum(ind[,2]*ind[,3])
emp3
```

```
## [1] 0.6666667
```

The first probability of JPY being large is around 5% by construction. If we know that CNY is large, then due to the positive dependence, the probability of JPY also being large increases to around 0.39. If we know that both CNY and USD are large, then the probability of JPY also being large increases to 0.92, due to the strong dependence between the currencies. The empirical estimates of the first probability are acceptable, though with accuracy limited by the number of data points. Estimates of the second two probabilities are wildly variable and unreliable due to the low amount of data in the region of interest. Working with the model is

more reliable here, and is, of course, the only option once we look into the tails beyond the observed data.
Comment: Here we approximated the joint distribution of (CNY,JPY) with a 2D extreme value model, rather than integrating USD out of the 3D (USD,CNY,JPY) model fit.

Query: This code doesn't seem to work if we replace JPY with (say) USD, in that prob3<prob2 which doesn't make sense as all variables are positively dependent. Similarly, repeating the calculations for the Husler-Reiss model produces negative probabilities. Suspect my code is wrong. Any ideas?