Solution 1

(a) The case $\lambda \to \infty$ is straightforward. If $\lambda \to 0$ and $z_2 > z_1$, then $\lambda^{-1}\log(z_2/z_1) \to \infty$ and $\lambda^{-1}\log(z_1/z_2) \to -\infty$, so $V(z_1, z_2) \to 1/z_1 = 1/\min(z_1, z_2)$. Likewise, the limit is $1/z_2 = 1/\min(z_1, z_2)$ if $z_1 > z_2$.

Here $\theta = V(1, 1) = 2\Phi(\lambda/2) \in [1, 2)$ tends to 2 when $\lambda \to \infty$ and equals 1 when $\lambda = 0$.

(b) Suppose without loss of generality that $z_1 > z_2$. Then

$$(z_1^\alpha + z_2^\alpha)^{-1/\alpha} = z_1^{-1}\{1 + (z_2/z_1)^\alpha\}^{-1/\alpha},$$

where $1 + (z_2/z_1)^\alpha > 1$, so $\{1 + (z_2/z_1)^\alpha\}^{-1/\alpha} \to 0$ as $\alpha \to 0$. Hence $V(z_1, z_2) \to 1/z_1 + 1/z_2$, corresponding to independence.

Now suppose that $\alpha \to \infty$, and write

$$(z_1^\alpha + z_2^\alpha)^{-1/\alpha} = z_2^{-1}\{1 + (z_1/z_2)^\alpha\}^{-1/\alpha} \to z_2^{-1}\{(z_1/z_2)^\alpha\}^{-1/\alpha} = z_1^{-1}.$$

Hence $V(z_1, z_2) \to 1/z_2 = 1/\min(z_1, z_2)$, as required.

Here $\theta = V(1, 1) = 2 - 2^{-1/\alpha} \in (1, 2)$ tends to 2 when $\alpha \to 0$ and tends to 1 when $\alpha \to \infty$.

Solution 2 Independence of Z_1 and Z_2 gives

$$\begin{aligned}
\mathrm{P}(X_1 \le z_1, X_2 \le z_2) &= \mathrm{P}\{Z_1 \le z_1, \max(aZ_1, Z_2)/(a+1) \le z_2\} \\
&= \mathrm{P}\left[Z_1 \le \min(z_1, (a+1)z_2/a\}, Z_2 \le (a+1)z_2\right] \\
&= \mathrm{P}\left[Z_1 \le \min\{z_1, (a+1)z_2/a\}\right] \mathrm{P}\left\{Z_2 \le (a+1)z_2\right\} \\
&= \exp\left[-\max\{1/z_1, a/\{a+1)z_2\}\right] \exp\left[-1/\{(a+1)z_2\}\right],
\end{aligned}$$

resulting in the stated formula for V.

To compute the corresponding angular distribution we note that the Pickands function is

$$A(w) = V\{1/w, 1/(1-w)\} = \max\left\{w, \frac{a(1-w)}{a+1}\right\} + \frac{1-w}{a+1}, \quad 0 \le w \le 1,$$

and then we use the formulae from the slide on the Pickands function (slide 191) to check that

$$\nu[0, w] = \begin{cases} 0, & w < a/(2a+1), \\ \dfrac{2a+1}{2(a+1)}, & a/(2a+1) \le w < 1, \\ 1, & w = 1, \end{cases}$$

corresponding to the stated distribution; this is discrete and puts mass at $w = a/(2a+1)$ and at $w = 1$. Plots of the corresponding values of (z_1, z_2) lie on two rays out from the origin.

Solution 3 Notation is simplified by writing $V = V(z_1, \ldots, z_D)$, $V_1 = \partial V(z_1, \ldots, z_D)/\partial z_1$, $V_{12} = \partial^2 V(z_1, \ldots, z_D)/\partial z_1 \partial z_2$ and so on, and $z_1' = \partial z_1(y_1)/\partial y_1$ and so forth.

(a) The first derivative is $\partial \mathrm{P}(Y_1 \le y_1, \ldots, Y_D \le y_D)/\partial y_1 = -z_1' V_1 e^{-V}$, and the second is therefore

$$\frac{\partial \mathrm{P}(Y_1 \le y_1, \ldots, Y_D \le y_D)}{\partial y_1 \partial y_2} = -z_1' \times z_2' V_{12} \times e^{-V} - z_1' V_1 \times -z_2' V_2 e^{-V} = z_1' z_2' (V_1 V_2 - V_{12}) e^{-V},$$

as given in the notes.

(b) The third derivative is

$$z_1'z_2'(z_3'V_{13} \times V_2 + V_1 \times z_3'V_{23} - z_3'V_{123})e^{-V} + z_1'z_2'(V_1V_2 - V_{12}) \times -z_3'V_3e^{-V},$$

and reduces to

$$z_1'z_2'z_3'\{-V_1V_2V_3 + (V_{12}V_3 + V_{13}V_2 + V_{23}V_1) - V_{123}\}e^{-V}.$$

Recall that the set $\{1, 2, 3\}$ can be partitioned in five different ways, as

$$123, \quad 12 \mid 3, \quad 13 \mid 2, \quad 23 \mid 1, \quad 1 \mid 2 \mid 3$$

where the successive partitions have 1, 2, 2, 2, 3 blocks. Thus the term in braces in the third derivative is the sum over all possible partitions of $\{1, 2, 3\}$, with terms having sign $(-1)^b$, where b is the corresponding number of blocks. The same is true in general.

The number of terms in the sum increases very rapidly; there are over 10^5 terms when $D = 10$, so computation of the likelihood (a product of such terms!) quickly becomes infeasible.

Solution 4

(a) We start modelling extremal dependence by fitting the logistic model using the code

```
(ModelD1<-fbvevd(cbind(NO2, PM10), model="log"))
Call: fbvevd(x = cbind(NO2, PM10), model = "log")
Deviance: 9167.191
AIC: 9181.191
Dependence: 0.4369681

Estimates
   loc1   scale1   shape1     loc2   scale2   shape2      dep
32.4008  10.1389  -0.0286  29.8827  12.4332   0.2925   0.6443

Standard Errors
   loc1   scale1   shape1     loc2   scale2   shape2      dep
0.45986  0.31545  0.02202  0.57992  0.48408  0.03487  0.02499
```

We initially inspect the marginal fit of the model using

```
# marginal fits
par(mfrow=c(2,4)); plot(ModelD1, mar=1)
plot(ModelD1, mar=2)
```

Figure 1 shows reasonable fits for both air pollutants; however the fitted margins tend to respectively underestimate and overestimate the high quantiles of NO2 and PM10.

We now check the bivariate fit using

```
# bivariate fit
par(mfrow=c(2,3)); plot(ModelD1)
```

The fit of the logistic model in Figure 2 seems reasonable. The plot illustrating the Pickands dependence function reveals certain discrepancies between the empirical and the fitted dependence function, which may arise from the presence of asymmetry in extremal dependence.

(b) We now model the bivariate extremal dependence via the asymmetric logistic model using the code

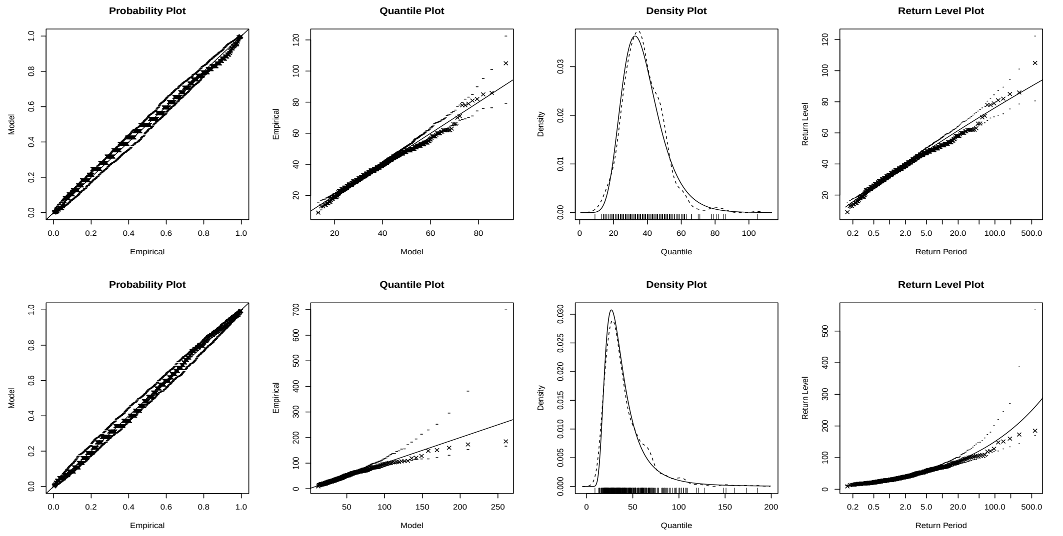

Figure 1: Diagnostic plots of marginal fits for NO2 (top panel) and PM10 (bottom panel).

```
 (ModelD2<-fbvevd(cbind(NO2, PM10), model="alog"))
 Call: fbvevd(x = cbind(NO2, PM10), model = "alog")
Deviance: 9165.467
AIC: 9183.467
Dependence: 0.4418221

Estimates
    loc1    scale1    shape1      loc2    scale2    shape2      asy1      asy2
32.32013  10.12205  -0.03254  30.02608  12.57328   0.28636   0.99990   0.89178
     dep
 0.61477

Standard Errors
     loc1     scale1     shape1       loc2     scale2     shape2       asy1
4.603e-01  3.152e-01  2.127e-02  5.931e-01  4.948e-01  3.606e-02  2.005e-06
     asy2        dep
8.331e-02  3.081e-02
```

The strength of extremal dependence is similar to that for the logistic model. Figure 3 shows a reasonable fit, although differences between the empirical and fitted Pickands dependence functions subsist.

Since the logistic model is a particular case of the asymmetric logistic, we rely on a likelihood ratio test and compare between the models using the code

```
 cp<-ModelD1$dev-ModelD2$dev
 1-pchisq(cp, df=2)
 0.422309
```

This is larger than 0.05, so it appears that the dependence can be modelled by a logistic dependence structure; there is little evidence for asymmetry,

(c) We only proceed here with the Hüsler–Reiss model as the fits from other models differ only slightly and do not provide any improvements. We use again the code

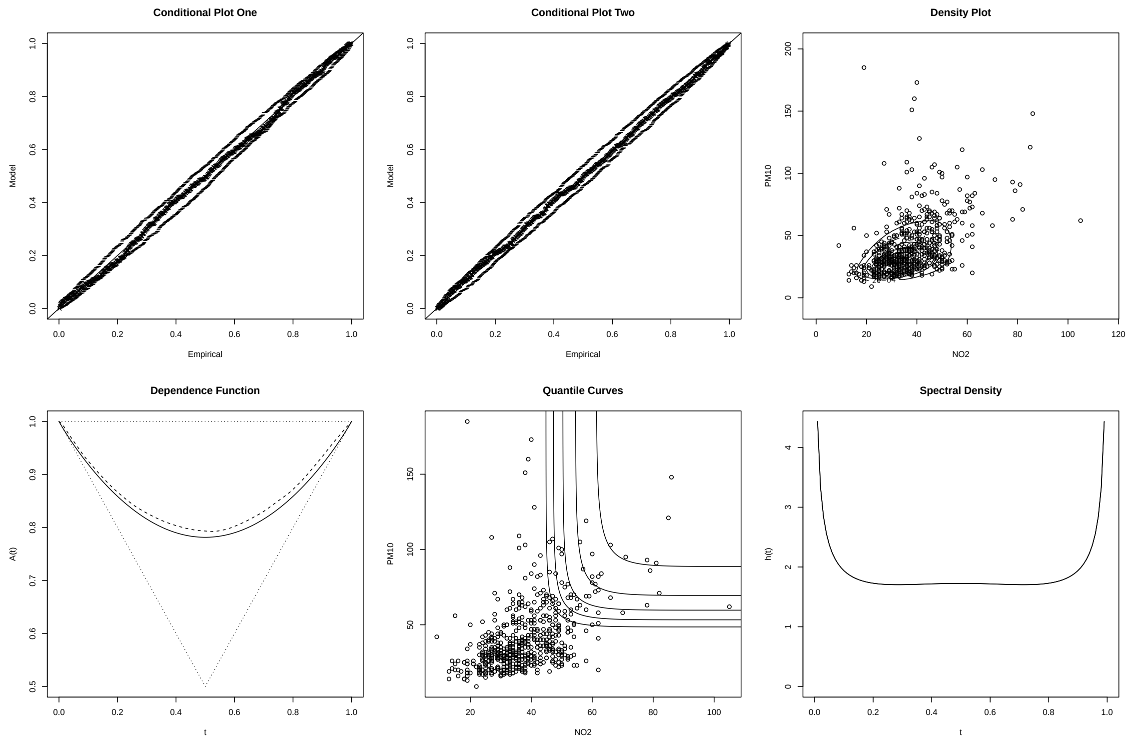

Figure 2: Diagnostic plots of the fit of the bivariate logistic model for daily maxima.

```
 (ModelD3<-fbvevd(cbind(NO2, PM10), model="hr"))

 Call: fbvevd(x = cbind(NO2, PM10), model = "hr")
Deviance: 9179.674
AIC: 9193.674
Dependence: 0.4168728

Estimates
    loc1    scale1    shape1      loc2    scale2    shape2       dep
32.35193  10.12590  -0.02004  29.80346  12.43712   0.31536   1.23174

Standard Errors
   loc1   scale1   shape1     loc2   scale2   shape2      dep
0.45893  0.31604  0.02024  0.57933  0.48909  0.03347  0.07125
```

Figure 4 shows a reasonable fit, but both the deviance and AIC for this fit are appreciably larger than those in (a) and (b), indicating that those may be better models.

(d) In this part we work with weekly maxima of NO2 and PM10, obtained using

```
#  We want to split 578 observations into blocks of 7,
#  so we initially add zero observations at the end of the vectors
NO2[578:581]<-0; PM10[578:581]<-0;

NO2_w<-apply(matrix(NO2, ncol=7, byrow=F), 1, function(x)max(x))
PM10_w<-apply(matrix(PM10, ncol=7, byrow=F), 1, function(x)max(x))
```

We model extremal dependence via the logistic dependence structure and obtain the following fit:

```
(ModelW1<-fbvevd(cbind(NO2_w, PM10_w), model="log"))
```

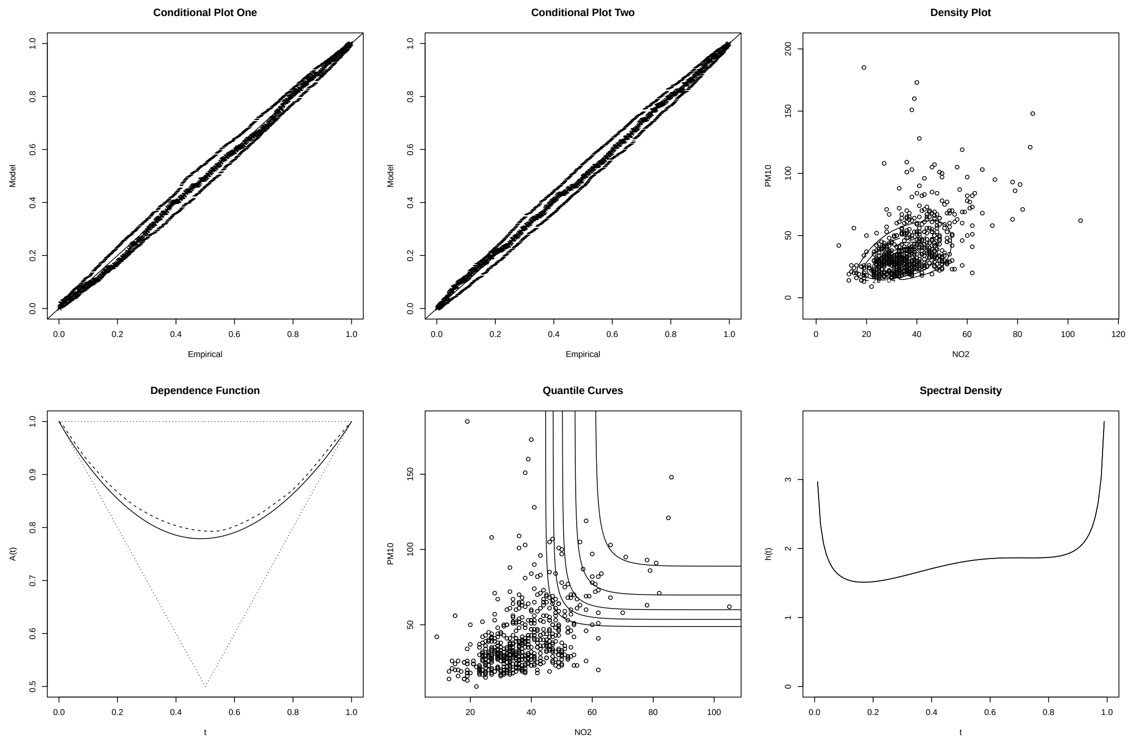

Figure 3: Diagnostic plots of the fit of the bivariate asymmetric logistic model for daily maxima.

```
Call: fbvevd(x = cbind(NO2_w, PM10_w), model = "log")
Deviance: 1395.31
AIC: 1409.31
Dependence: 0.4604533

Estimates
   loc1   scale1   shape1     loc2   scale2   shape2      dep
45.5834   9.7346   0.1019  51.4028  20.2904   0.2448   0.6225

Standard Errors
   loc1   scale1   shape1     loc2   scale2   shape2      dep
1.20519  0.88295  0.08197  2.57124  2.07558  0.10188  0.06447
```

The strength of extremal dependence is similar to that for the daily data.

Figure 5 shows a reasonable fit for the dependence structure, although we see again similar problems to the daily data with the Pickands dependence function.

We also check the asymmetric logistic model and obtain

```
(ModelW2<-fbvevd(cbind(NO2_w, PM10_w), model="alog"))

Call: fbvevd(x = cbind(NO2_w, PM10_w), model = "alog")
Deviance: 1395.579
AIC: 1413.579
Dependence: 0.4389836

Estimates
    loc1    scale1    shape1      loc2    scale2    shape2      asy1      asy2
45.84077   9.75481   0.09113  51.95000  20.55651   0.22959   0.99944   0.93725
     dep
```

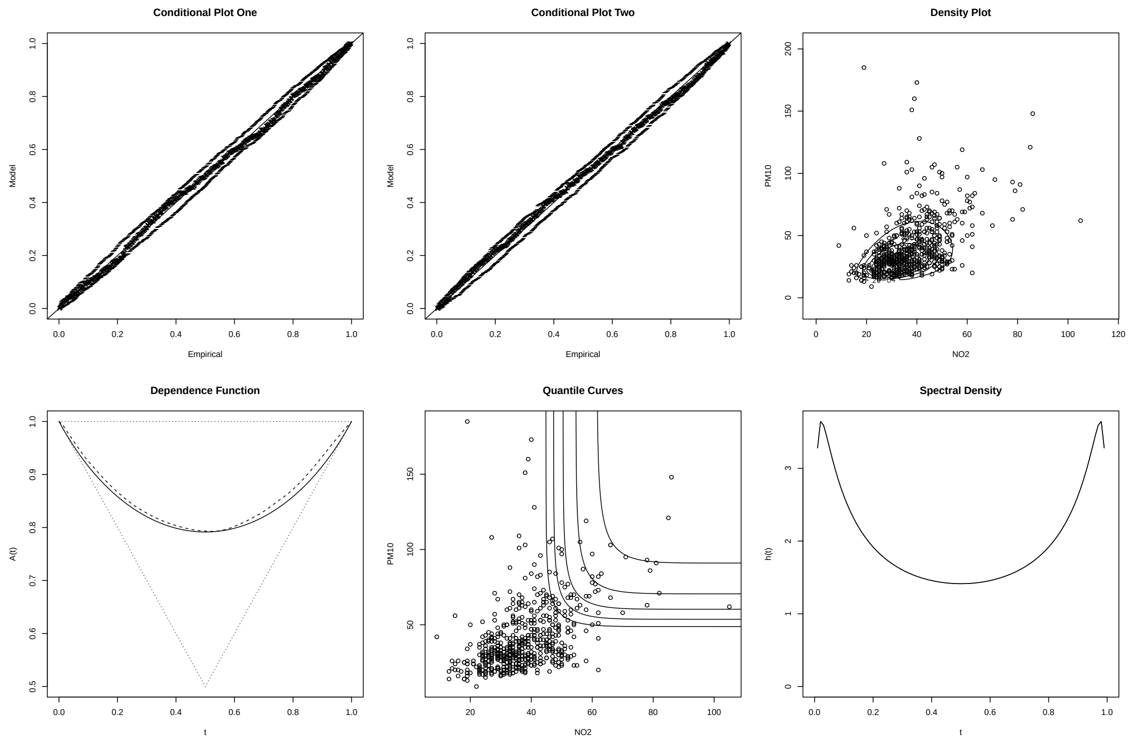

Figure 4: Diagnostic plots of the fit of the bivariate Hüsler–Reiss model for daily maxima.

```
 0.62872

Standard Errors
     loc1     scale1     shape1       loc2     scale2     shape2       asy1
1.215e+00  8.841e-01  8.084e-02  2.701e+00  2.116e+00  1.085e-01  2.001e-06
     asy2        dep
3.930e-01  9.031e-02
```

The fitted Pickands dependence function in Figure 6 fits its empirical counterpart better than for the logistic model; nevertheless the deviances are very similar, so the logistic model still seems reasonable.

Solution 5

(a) We first fit the logistic dependence structure to the exceedances of NO2 and PM10 using the code

```
airpollutants<-cbind(NO2, PM10)
(POT1 <- evd::fbvpot(airpollutants, apply(airpollutants, 2, quantile, 0.95),
 model="log"))

 Call: evd::fbvpot(x = airpollutants, threshold = apply(airpollutants, 2, quantile,
  0.95), model = "log")
Likelihood: censored
Deviance: 851.0401
AIC: 861.0401
Dependence: 0.2399396

Threshold: 58 84.15
Marginal Number Above: 26 29
Marginal Proportion Above: 0.045 0.0502
```

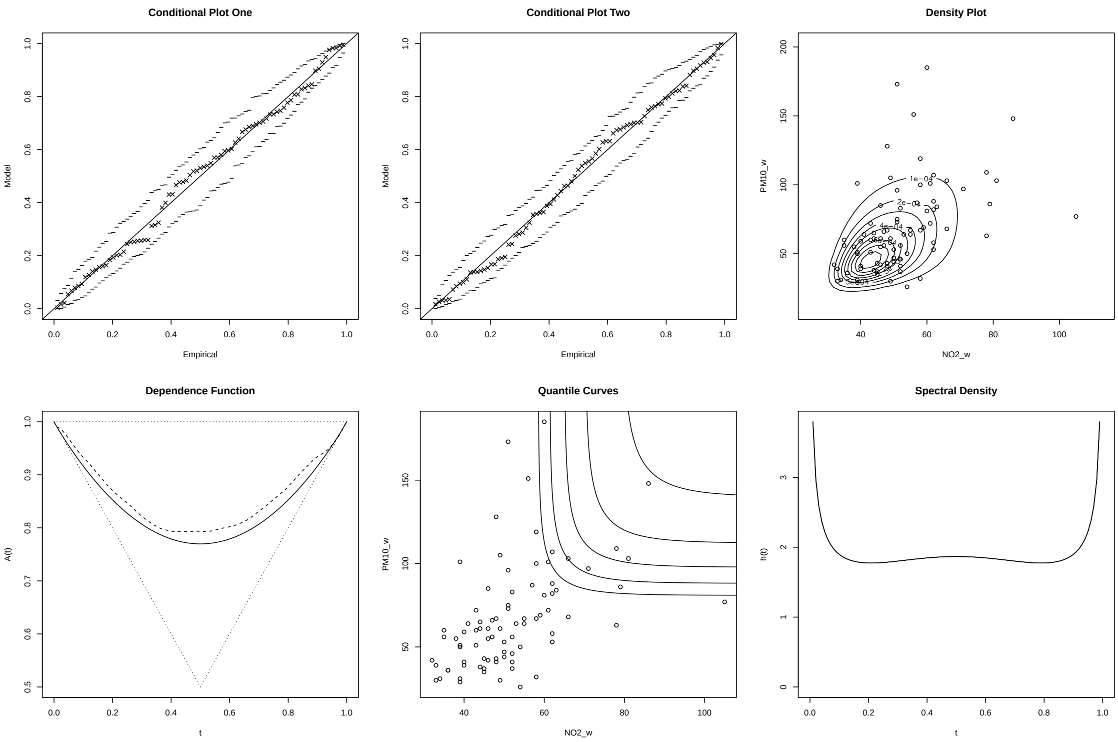

Figure 5: Diagnostic plots of the fit of the bivariate logistic model for weekly maxima.

```
Number Above: 8
Proportion Above: 0.0138

Estimates
  scale1    shape1    scale2    shape2       dep
11.79466   0.01502  26.85926   0.04004   0.81562

Standard Errors
 scale1   shape1   scale2   shape2      dep
3.63755  0.25698  7.86268  0.23691  0.06329
```

The estimated extremal dependence is weaker than for daily maxima, perhaps because selecting a high threshold means that we only consider large exceedances.

We inspect the marginal fits using

```
par(mfrow=c(2,4)); plot(POT1, mar=1)
plot(POT1, mar=2)
```

Figure 7 seems reasonable, in particular for PM10, but problems may arise from the presence of ties and from the small number of exceedances, as shown in the top panel for NO2.

We now check the bivariate fit using

```
par(mfrow=c(1,4)); plot(POT1)
```

Figure 8 shows a reasonable fit for the logistic model, and we notice from the third plot that there are rather few extreme observations jointly in both pollutants.

(b) We repeat the same steps as in part (a) for the asymmetric logistic model using the code

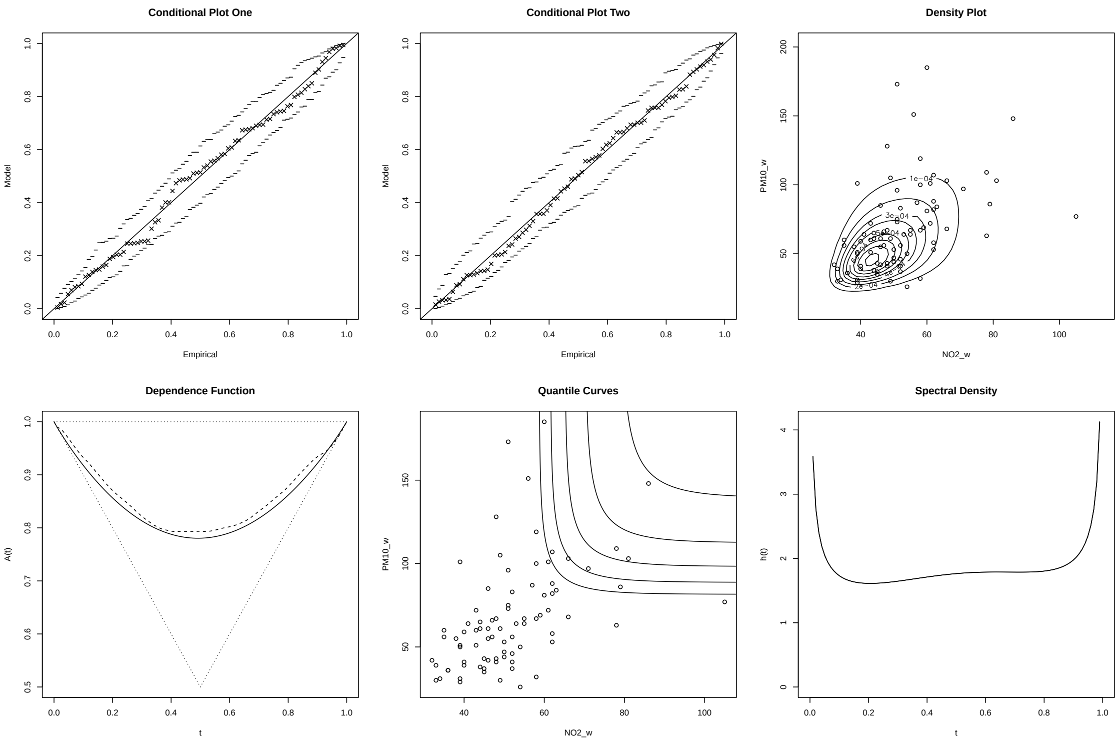

Figure 6: Diagnostic plots of the fit of the bivariate asymmetric logistic model for weekly maxima.

```
(POT2 <- evd::fbvpot(airpollutants, apply(airpollutants, 2, quantile, 0.95),
 model="alog"))

Call: evd::fbvpot(x = airpollutants, threshold = apply(airpollutants, 2, quantile,
 0.95), model = "alog")
Likelihood: censored
Deviance: 842.7556
AIC: 856.7556
Dependence: 0.2295116

Threshold: 58 84.15
Marginal Number Above: 26 29
Marginal Proportion Above: 0.045 0.0502
Number Above: 8
Proportion Above: 0.0138

Estimates
   scale1     shape1     scale2     shape2       asy1       asy2        dep
12.444427  -0.070632  28.987009   0.002408   0.735653   0.242493   0.367385

Standard Errors
 scale1   shape1   scale2   shape2     asy1     asy2      dep
3.82401  0.24100  8.66101  0.25969  0.18538  0.08982  0.11088
```

In contrast to the previous exercise, here we notice that the presence of stronger asymmetry in the fitted model, indicated by the large difference between the `asy1` and `asy2` parameters.

The plots of the bivariate fits in Figure 9 confirm that the fitted dependence structure is highly asymmetric, and the corresponding deviance is much lower than for the logistic model. The likelihood ratio test between the logistic and asymmetric logistic models gives

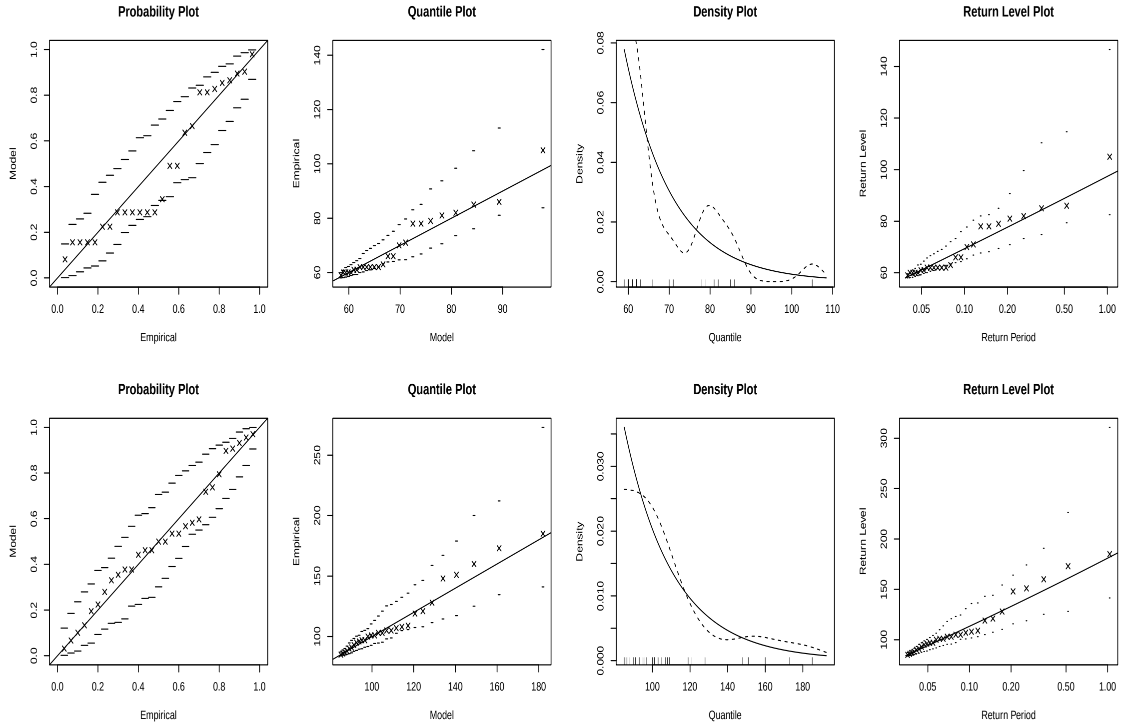

Figure 7: Diagnostic plots of marginal fits for NO2 (top panel) and PM10 (bottom panel).

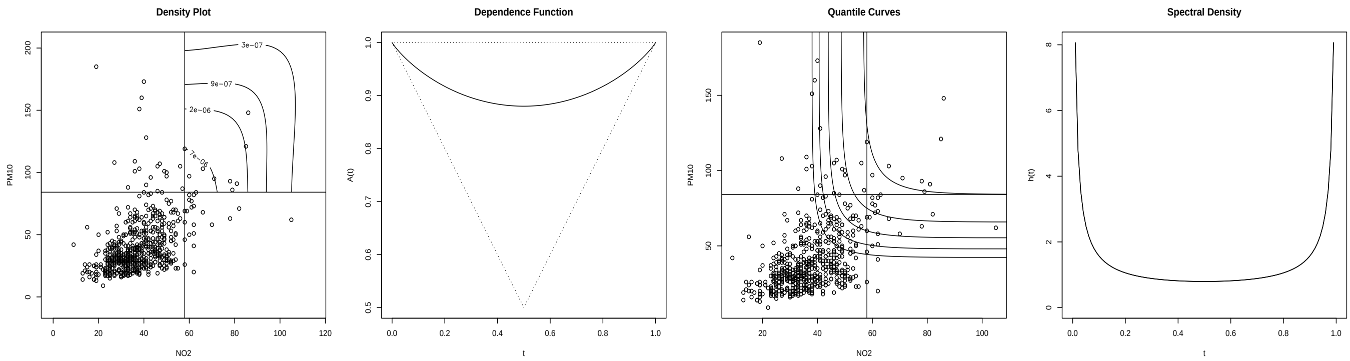

Figure 8: Diagnostic plots of the fit of the bivarite logistic model for exceedances of daily maxima.

```
cp2<-POT1$dev-POT2$dev
1-pchisq(cp2, df=2)
0.01588705
```

which suggests a fairly strong departure from symmetry.

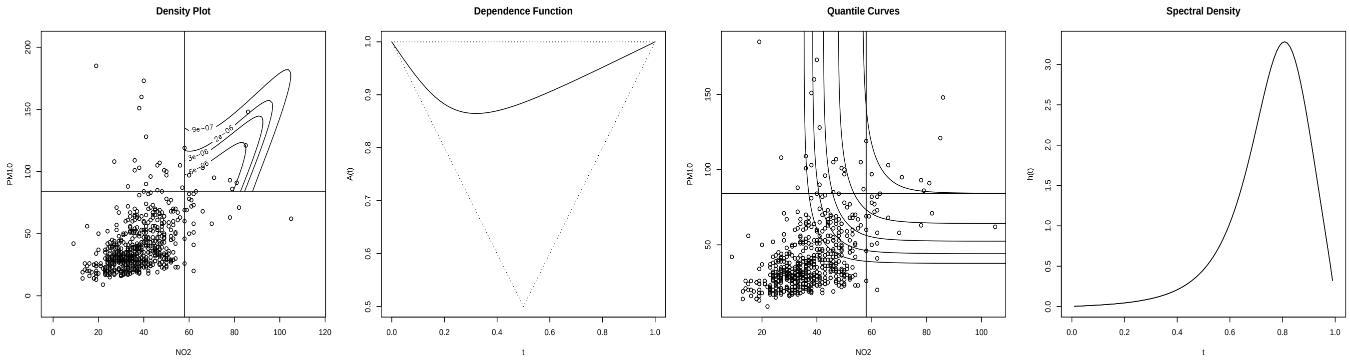

Figure 9: Diagnostic plots of the fit of the bivarite asymmetric logistic model for exceedances of daily maxima.