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

Dr. Boris Béranger

Week 5 - Lab Solutions

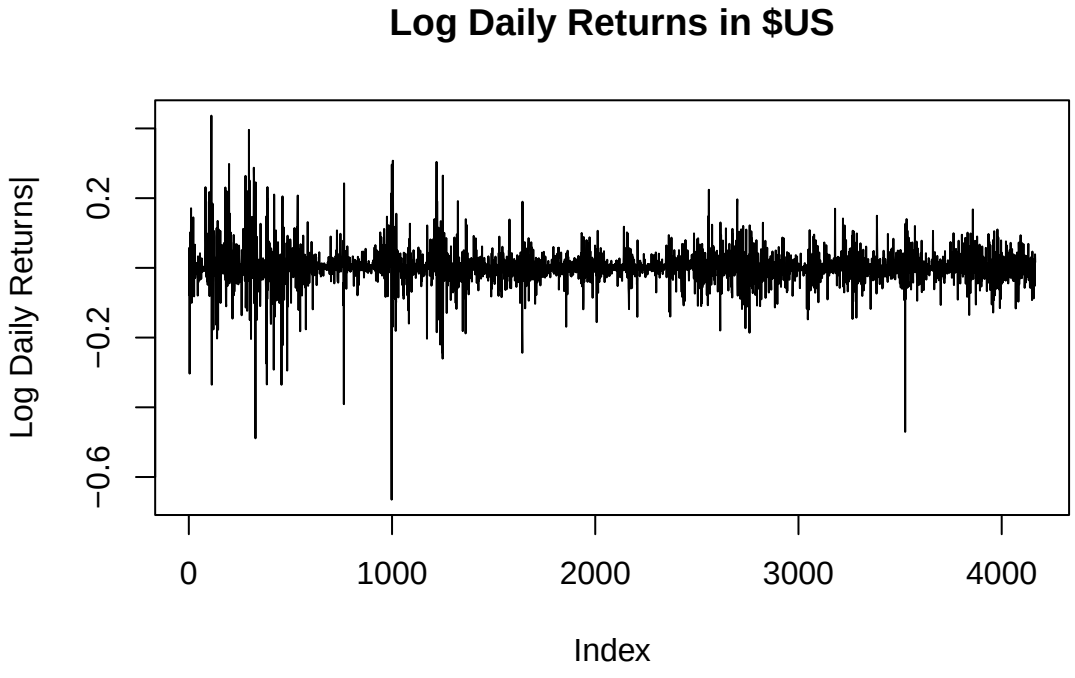
Exercise 1

The file `bitcoin.csv` contains 4,727 closing prices in US$ of bitcoin between 3rd January 2009 to 12th December 2021 downloaded from the website `https://coinmetrics.io`.

a) Plot the log daily returns of the bitcoin price. Is the series stationary? Do the positive and negative tails seem to have the same behaviour?

Solution:

```
bit <- read.table("../data/bitcoin.csv", header=T, sep=",")
L <- length(bit$priceUSD)
ldr <- log(bit$priceUSD[2:L]/bit$priceUSD[1:(L-1)])
ind <- is.na(ldr)
ldrok <- ldr[!ind]
plot(ldrok, ylab="Log Daily Returns|", main="Log Daily Returns in $US", type='l')
```



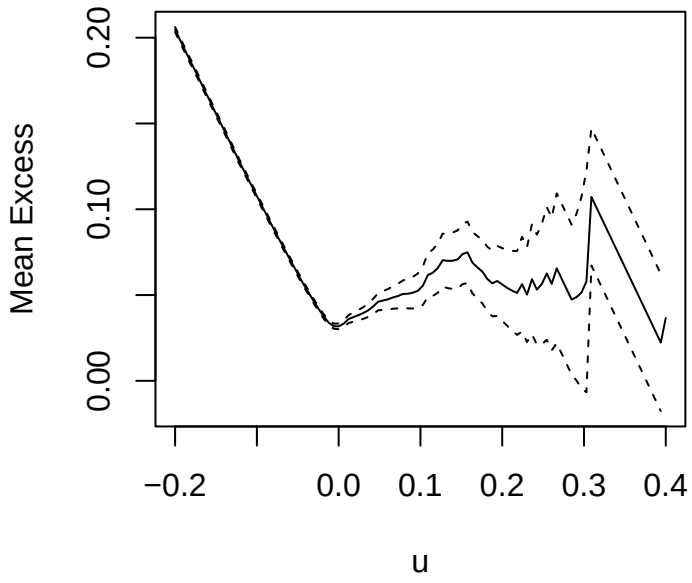
It seems as though there may have been more extreme values, positive and negative, earlier in the series compared to later on (with the exception of 1 extreme negative value towards the end of the series). It is not clear that the positive tail behaves particularly differently to the negative tail, although perhaps there are a few more "larger" values in the negative tail than the positive tail.

b) Suppose we are interested in the positive extremes. Produce a mean excess plot and identify a suitable threshold for a point process model for threshold exceedances. For convenience, let's work on a per-observation parameterisation by specifying `time.units="1/year"` within `fevd`. Verify your choice through goodness of fit diagnostics.

Note: For point process models, the diagnostics via e.g. `plot(fit)`, where `fit` is a `fevd` object, also include a Z-plot. This is a qq-plot of the distances between successive threshold exceedance events which, according to a *Poisson* point process model, should be exponentially distributed. For the (Poisson) point process model to be valid, we need convergence of the distribution of threshold exceedances (as for the GPD) as $u \to \infty$ in addition to the distribution of distances between successive threshold exceedances (as shown by the Z-plot). For more information on the Z-plot, see pp.25–26 of Gilleland and Katz (2016).

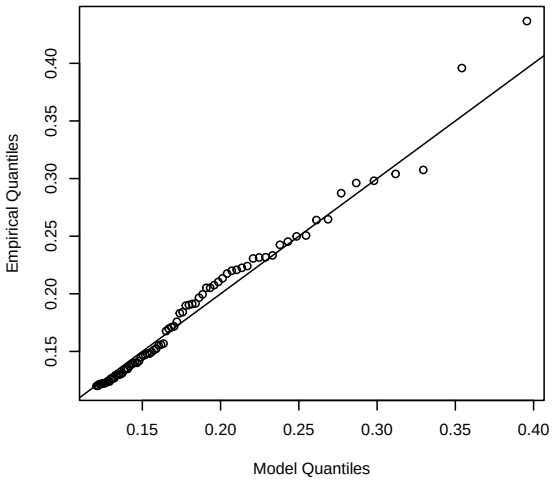
Solution:

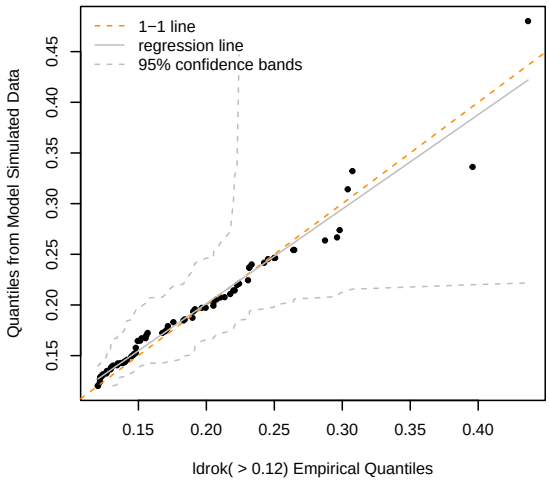
```
library(ismev)
mrl.plot(ldrok, umin=-0.2, umax=0.4) # ismev function
```

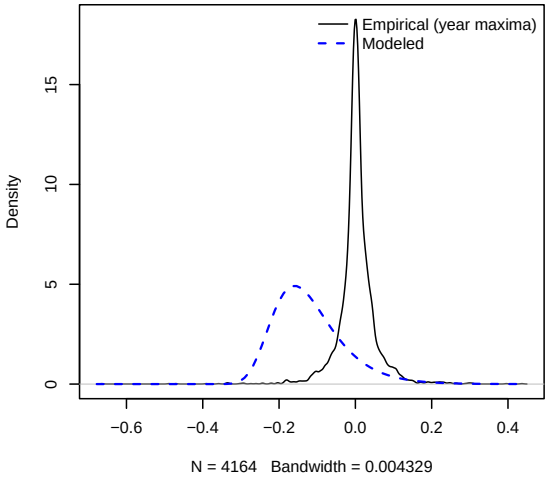


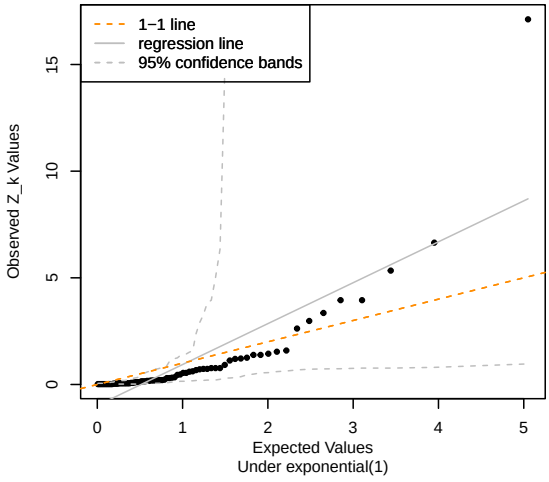

```
library(extRemes)
u <- 0.12 # choose threshold here
fit <- fevd(ldrok, threshold=u, type="PP", time.units="1/year")
plot(fit)
```

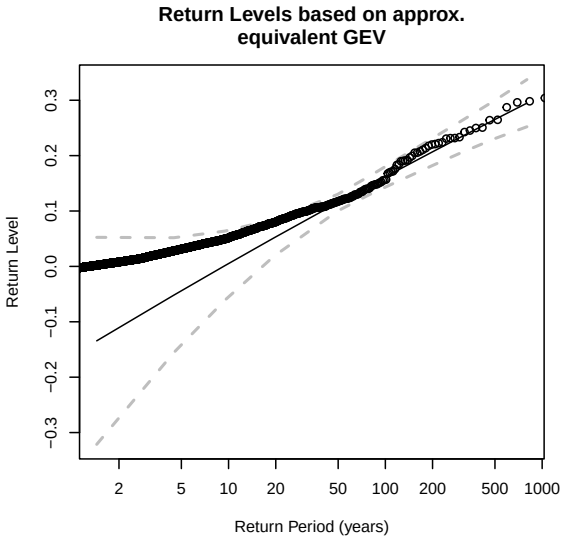
fevd(x = ldrok, threshold = u, type = "PP", time.units = "1/year")











The mean excess plot seems to be piecewise linear with an abrupt change at 0 and a less severe change around 1 to 1.5. Experimenting with a few potential thresholds in this range, and producing Poisson point process fit diagnostics, $u = 0.12$ (shown) produces a reasonable GPD threshold exceedance fit, and an acceptable exponential distribution of successive threshold exceedance distances, though this is less than perfect for the largest 8 or so distances.

c) For the chosen threshold, explore whether any of the point process model parameters, $\mu(t), \sigma(t), \xi(t)$, can be expressed as a function of time e.g., linear or quadratic. As part of this, examine whether regression modelling of $\log \sigma(t)$ might be more preferable to modelling $\sigma(t)$ directly.

Solution:

```
X <- data.frame(time=bit$index[!ind], time2=bit$index[!ind]^2)
fit.1 <- fevd(ldrok, data=X, threshold=u, location.fun=~time,type="PP",
              time.units="1/year")
lr.test(fit, fit.1) # overwhelming support
```

```
##
##  Likelihood-ratio Test
##
## data:  ldrokldrok
## Likelihood-ratio = 58.31, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 2.24e-14
## alternative hypothesis: greater
```

```
fit.2 <- fevd(ldrok, data=X, threshold=u, scale.fun=~time, type="PP",
              time.units="1/year")
lr.test(fit, fit.2) # overwhelming support
```

```
##
##  Likelihood-ratio Test
##
## data:  ldrokldrok
## Likelihood-ratio = 58.006, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 2.614e-14
## alternative hypothesis: greater
```

```
fit.3 <- fevd(ldrok, data=X, threshold=u, scale.fun=~time, use.phi=T,
              type="PP", time.units="1/year")
lr.test(fit, fit.3) # strongest overwhelming support
```

```
##
##  Likelihood-ratio Test
##
## data:  ldrokldrok
## Likelihood-ratio = 60.979, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 5.769e-15
## alternative hypothesis: greater
```

```
fit.4 <- fevd(ldrok, data=X, threshold=u, shape.fun=~time, type="PP",
              time.units="1/year")
lr.test(fit, fit.4) # overwhelming support
```

```
##
##  Likelihood-ratio Test
##
## data:  ldrokldrok
## Likelihood-ratio = 58.188, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 2.382e-14
## alternative hypothesis: greater
```

```
fit.5 <- fevd(ldrok, data=X, threshold=u, location.fun=~time, scale.fun=~time,
              use.phi=T, type="PP", time.units="1/year")
lr.test(fit.3, fit.5) # support
```

```
##
##  Likelihood-ratio Test
##
## data:  ldrokldrok
## Likelihood-ratio = 8.3338, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 0.003891
## alternative hypothesis: greater
```

```
fit.6 <- fevd(ldrok, data=X, threshold=u, scale.fun=~time, use.phi=T,
              shape.fun=~time,type="PP", time.units="1/year")
lr.test(fit.3, fit.6) # no support
```

```
##
##  Likelihood-ratio Test
##
## data:  ldrokldrok
## Likelihood-ratio = 0.0013628, chi-square critical value = 3.8415, alpha
## = 0.0500, Degrees of Freedom = 1.0000, p-value = 0.9706
## alternative hypothesis: greater
```

```
fit.7 <- fevd(ldrok, data=X, threshold=u, location.fun=~time+time2,
              scale.fun=~time, use.phi=T, type="PP", time.units="1/year")
lr.test(fit.5, fit.7) # no support
```

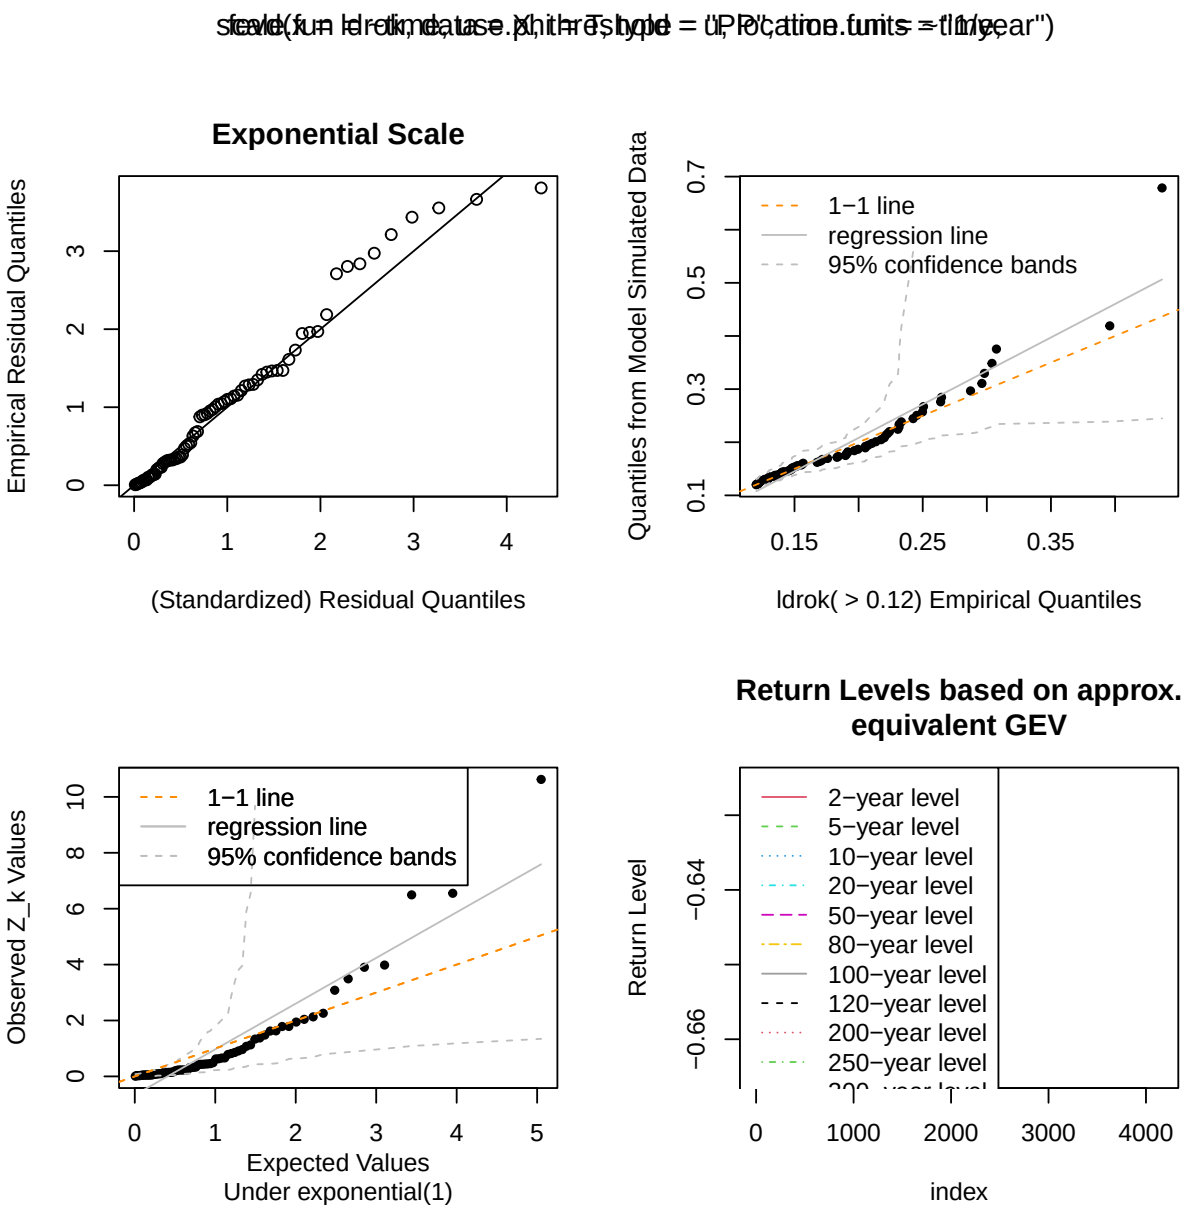
```
##
##  Likelihood-ratio Test
##
## data:  ldrokldrok
## Likelihood-ratio = 1.3296, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 0.2489
## alternative hypothesis: greater
```

```
fit.8 <- fevd(ldrok, data=X, threshold=u, location.fun=~time,
              scale.fun=~time+time2, use.phi=T, type="PP", time.units="1/year")
```

```
lr.test(fit.5, fit.8) # no support
```

```
##
##  Likelihood-ratio Test
##
## data:  ldrokldrok
## Likelihood-ratio = 2.799, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 0.09432
## alternative hypothesis: greater
```

```
plot(fit.5) # chosen model
```



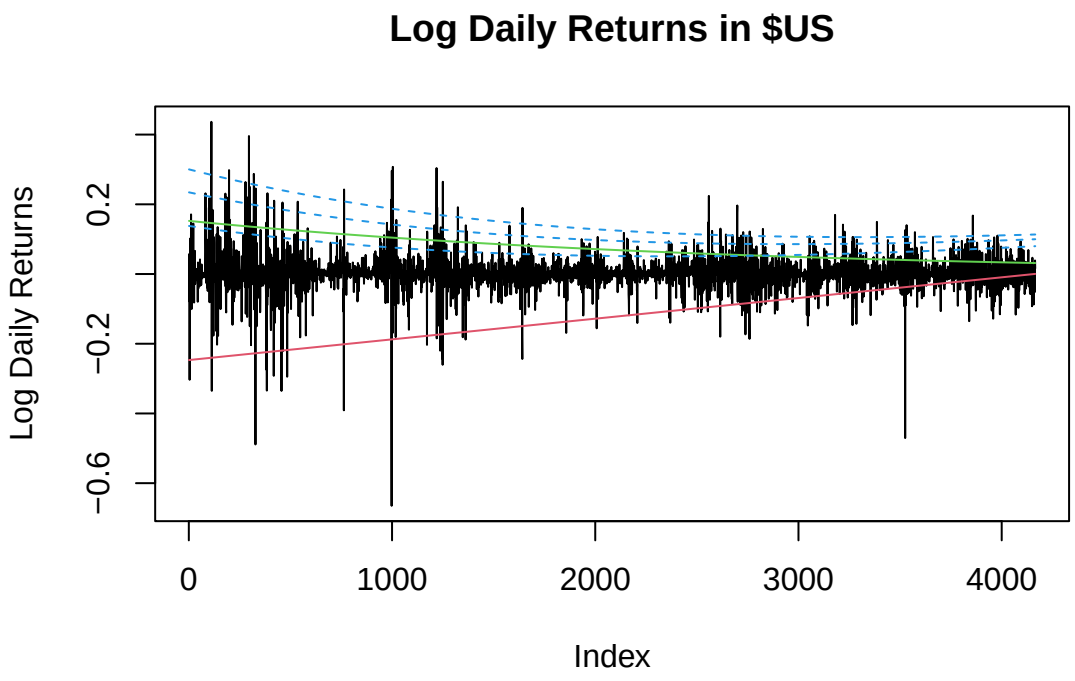
Because of the decrease in upper tail extremes over time, allowing any one parameter to change with time produces a significant improvement in model fit over a model with constant

parameters. Out of all possibilities modelling $\log(\sigma(t)) = \phi_0 + \phi_1 t$ gives the most significant improvement. Using this as a base, allowing $\mu(t) = \mu_0 + \mu_1 t$ to be linear in time also produces a significant contribution. No other predictors significantly improve the model fit beyond these. The resulting model fit diagnostics seem acceptable.

d) On a plot of the log daily returns, superimpose the time varying model parameters. Does the manner in which they vary make sense? Also superimpose 20-, 50- and 100-observation return levels (use the `return.level function`) that vary with the parameter values. Approximately 5%, 2% and 1% of observations should exceed these thresholds – do they? Use this information to comment on the suitability of this model for these data.

Solution:

```
p=fit.5$results$par
plot(ldrok, ylab="Log Daily Returns", main="Log Daily Returns in $US",
     type='l')
lines(p[1]+p[2]*X$time, col=2)
lines(exp(p[3]+p[4]*X$time), col=3)
v <- make.qcov(fit.5, vals=list(mu1=X$time, phi1=X$time))
a20 <- return.level(fit.5, return.period=c(20), qcov=v)
a50 <- return.level(fit.5, return.period=c(50), qcov=v)
a100 <- return.level(fit.5, return.period=c(100), qcov=v)
lines(a20, col=4, lty=2)
lines(a50, col=4, lty=2)
lines(a100, col=4, lty=2)
```




```
sum(ldrok>a20)/length(ldrok) # about 20% of observations?
```

```
## [1] 0.06314526
```

```r
sum(ldrok>a50)/length(ldrok) # about 2% of observations?
```

```
## [1] 0.02160864
```

```r
sum(ldrok>a100)/length(ldrok) # about 1% of observations?
```

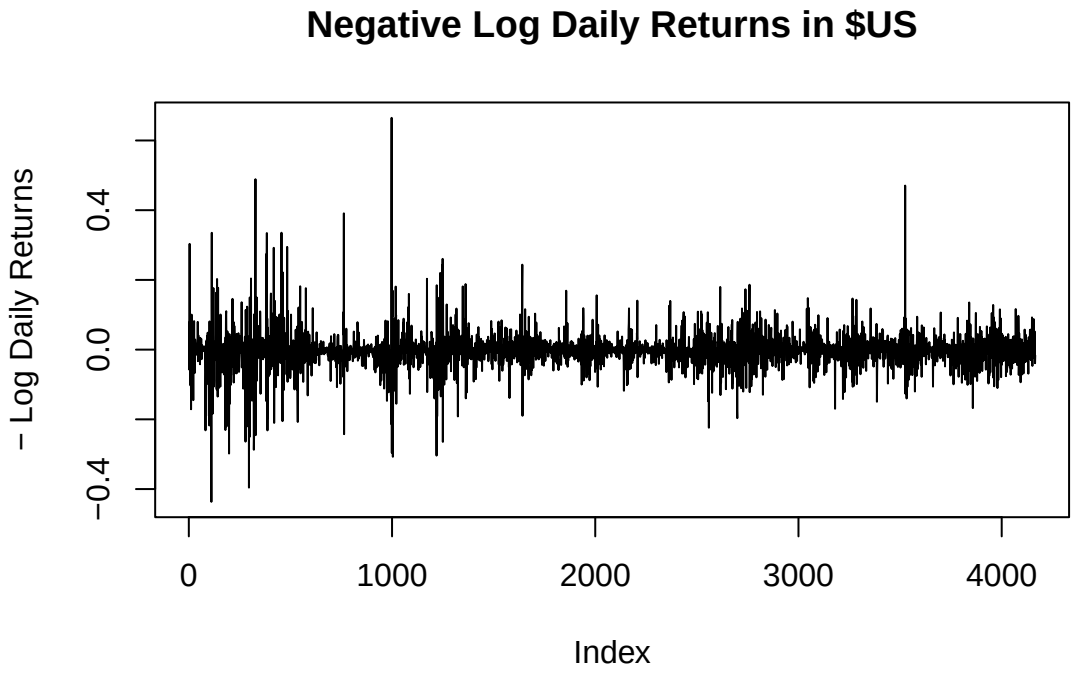
```
## [1] 0.009363745
```

$\mu(t)$ – the red line – shows an increase with time, as does $\log(\sigma(t))$ – the green line – which is accordingly an exponential decay. Both of these make sense in terms of the observed data showing reduced extreme tail behaviour as time goes on. Whether these models are useful for extrapolation, however, is the usual statistical question. Accordingly the return level estimates (blue dashed lines) also reduce over time, which makes sense. Supporting this is the empirical return level exceedance rates of 6.3%, 2.2% and 0.9% are relatively close to the model based rates of 5%, 2% and 1% respectively.

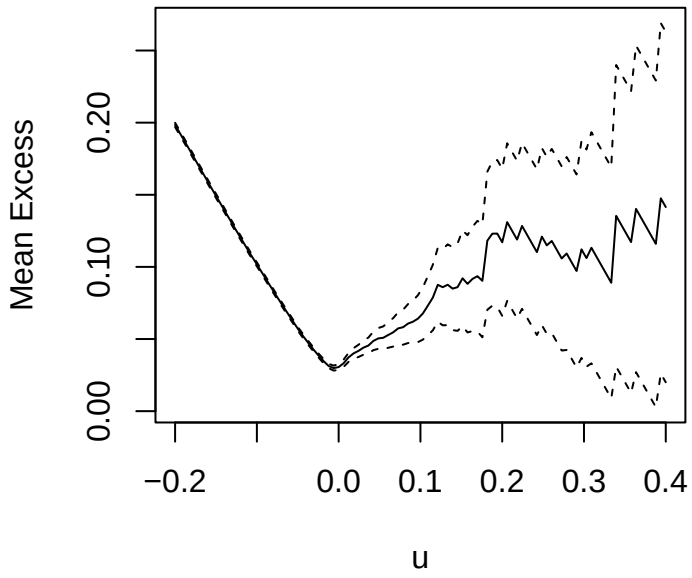
e) Repeat this process (threshold identification, model selection, return level estimation, model fit assessment) for the negative tail. Do both tails behave similarly?

Solution:

```r
plot(-ldrok, ylab="- Log Daily Returns",
     main="Negative Log Daily Returns in $US",type='l')
```

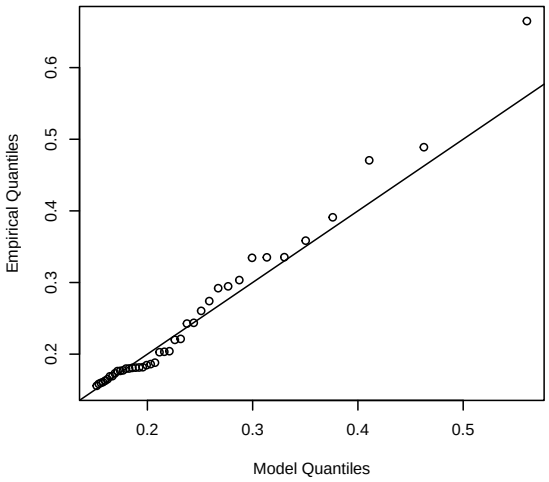


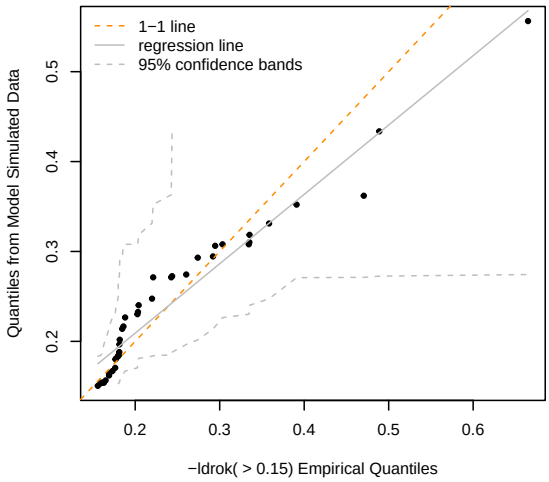

```r
mrl.plot(-ldrok, umin=-0.2, umax=0.4) # ismev function
```

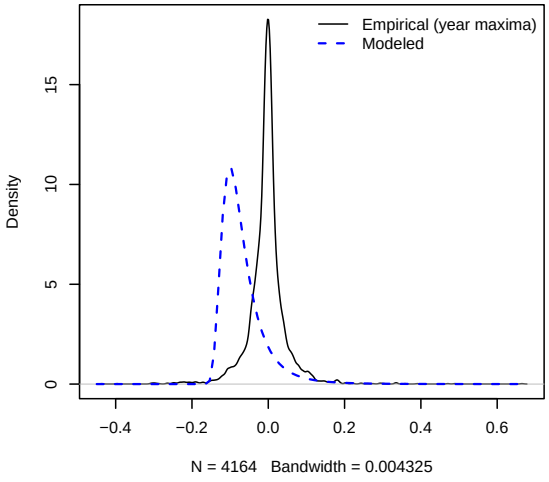


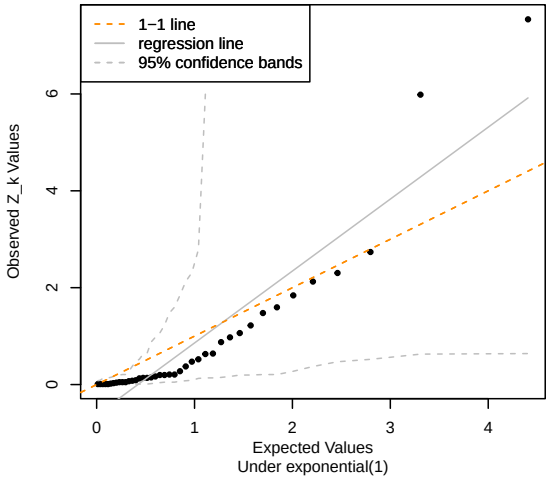

```
u <- 0.15 # choose threshold here
fit <- fevd(-ldrok,threshold=u,type="PP",time.units="1/year")
plot(fit)
```

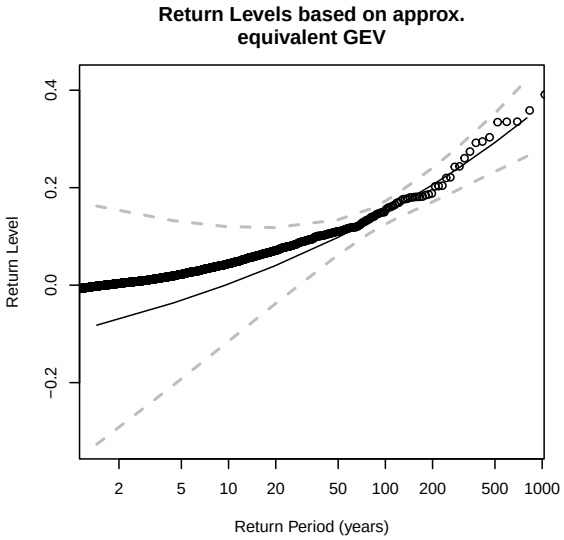
fevd(x = −ldrok, threshold = u, type = "PP", time.units = "1/year")











The features of the negative tail are similar to the positive tail. A threshold of $u = 0.15$ seems sensible from the mean excess plot in addition to having reasonable point process model diagnostics.

```
fit.1<- fevd(-ldrok, data=X, threshold=u, location.fun=~time, type="PP",
             time.units="1/year")
lr.test(fit, fit.1) # overwhelming support
```

```
##
##  Likelihood-ratio Test
##
## data:  -ldrok-ldrok
## Likelihood-ratio = 36.53, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 1.504e-09
## alternative hypothesis: greater
```

```
fit.2 <- fevd(-ldrok,data=X, threshold=u, scale.fun=~time, type="PP",
              time.units="1/year")
lr.test(fit, fit.2) # overwhelming support
```

```
##
##  Likelihood-ratio Test
##
## data:  -ldrok-ldrok
## Likelihood-ratio = 37.493, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 9.174e-10
## alternative hypothesis: greater
```

```
fit.3 <- fevd(-ldrok, data=X, threshold=u, scale.fun=~time, use.phi=T,
              type="PP", time.units="1/year")
lr.test(fit, fit.3) # strongest overwhelming support
```

```
##
##  Likelihood-ratio Test
##
## data:  -ldrok-ldrok
## Likelihood-ratio = 38.003, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 7.062e-10
## alternative hypothesis: greater
```

```
fit.4 <- fevd(-ldrok, data=X, threshold=u, shape.fun=~time, type="PP",
              time.units="1/year")
lr.test(fit, fit.4) # overwhelming support
```

```
##
##  Likelihood-ratio Test
##
## data:  -ldrok-ldrok
## Likelihood-ratio = 32.266, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 1.344e-08
## alternative hypothesis: greater
```

```
fit.5 <- fevd(-ldrok, data=X, threshold=u, location.fun=~time,
              scale.fun=~time, use.phi=T, type="PP", time.units="1/year")
lr.test(fit.3, fit.5) # no support
```

```
##
##  Likelihood-ratio Test
##
## data:  -ldrok-ldrok
## Likelihood-ratio = 0.018857, chi-square critical value = 3.8415, alpha
## = 0.0500, Degrees of Freedom = 1.0000, p-value = 0.8908
## alternative hypothesis: greater
```

```
fit.6 <- fevd(-ldrok, data=X, threshold=u, scale.fun=~time, shape.fun=~time,
              use.phi=T, type="PP", time.units="1/year")
lr.test(fit.3, fit.6) # no support
```

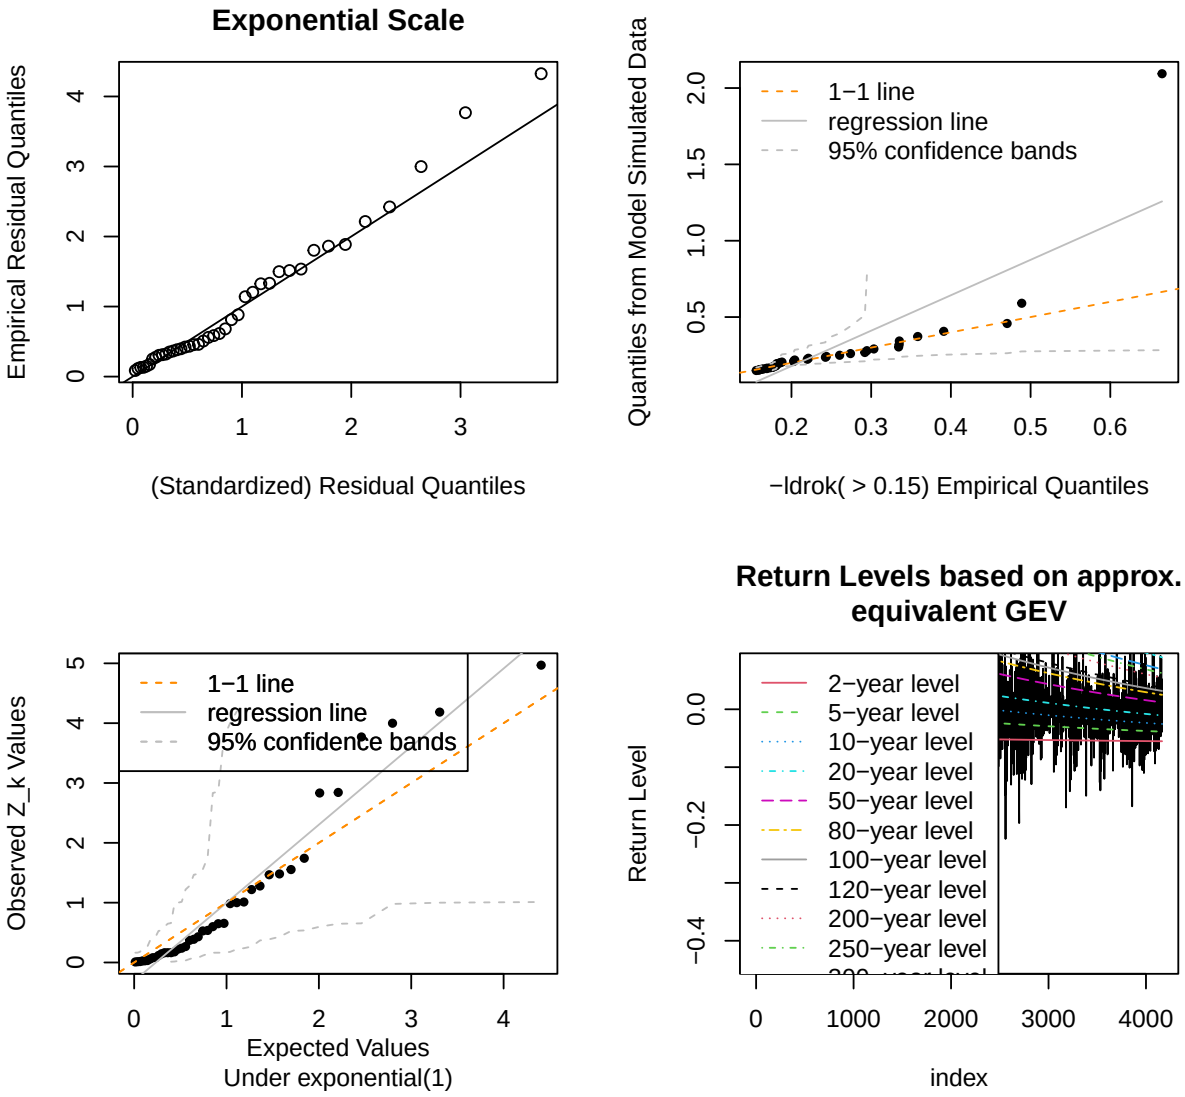
```
##
##  Likelihood-ratio Test
##
## data:  -ldrok-ldrok
## Likelihood-ratio = 0.13601, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 0.7123
## alternative hypothesis: greater
```

```
fit.7 <- fevd(-ldrok, data=X, threshold=u, scale.fun=~time+time2, use.phi=T,
              type="PP", time.units="1/year")
lr.test(fit.3, fit.7) # no support
```

```
##
##  Likelihood-ratio Test
##
## data:  -ldrok-ldrok
## Likelihood-ratio = 0.060735, chi-square critical value = 3.8415, alpha
## = 0.0500, Degrees of Freedom = 1.0000, p-value = 0.8053
## alternative hypothesis: greater
```

```
plot(fit.3) # Final model
```

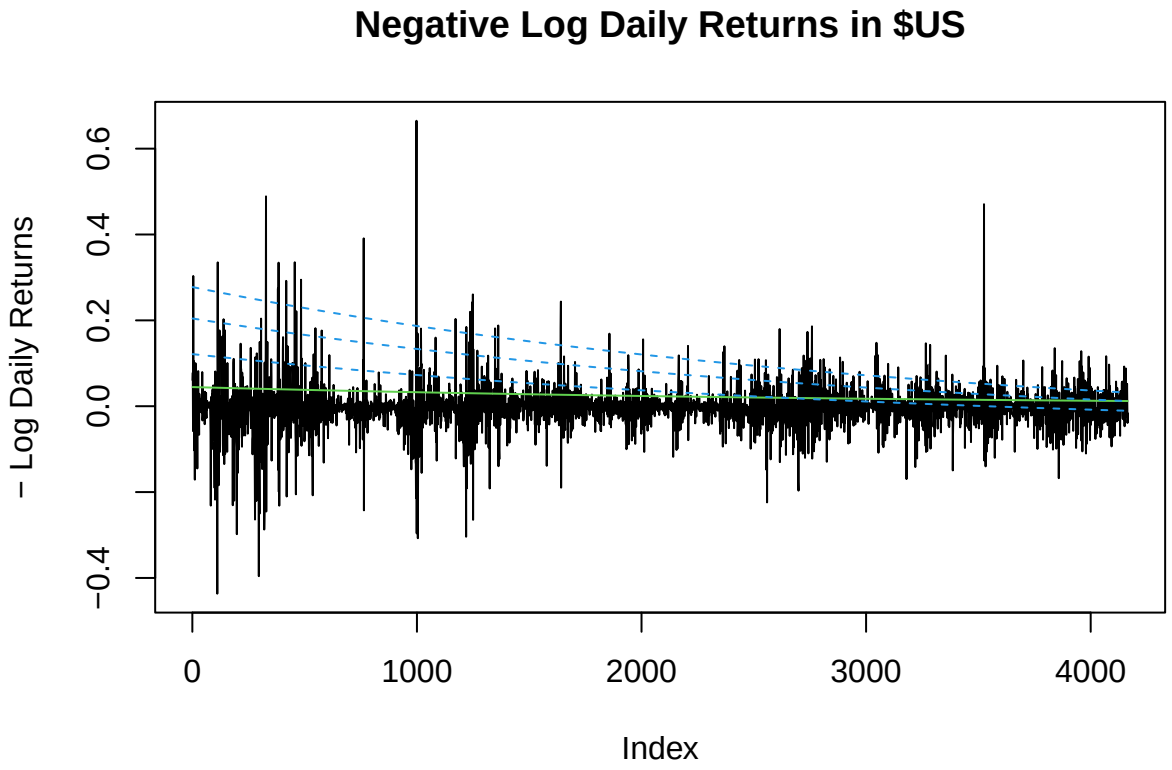
fevd(x = -ldrok, data = X, threshold = 0.15, time = time, time.units = "1/year")



Following the same process as before, allowing any of the model parameters to vary with time sees an overwhelmingly significant improvement to the model fit. The largest improvement came from expressing $\log(\sigma(t)) = \phi_0 + \phi_1 t$ as linear in time. Adopting this into the model, no further improvements were made. The final model fit seems fairly reasonable.

```
p <- fit.3$results$par
plot(-ldrok, ylab="- Log Daily Returns",
     main="Negative Log Daily Returns in $US", type='l')
lines(exp(p[2]+p[3]*X$time), col=3)
v <- make.qcov(fit.3, vals=list(phi1=X$time))
a20 <- return.level(fit.3, return.period=c(20), qcov=v)
a50 <- return.level(fit.3, return.period=c(50), qcov=v)
a100 <- return.level(fit.3, return.period=c(100), qcov=v)
lines(a20,col=4,lty=2)
```

```
lines(a50,col=4,lty=2)
lines(a100,col=4,lty=2)
```




```
sum(ldrok>a20)/length(ldrok) # about 5% of observations
```

```
## [1] 0.2388956
```

```
sum(ldrok>a50)/length(ldrok) # about 2% of observations
```

```
## [1] 0.08115246
```

```
sum(ldrok>a100)/length(ldrok) # about 1% of observations
```

```
## [1] 0.0362545
```

The location parameter, $\mu(t)$, is estimated to be constant over time. The scale parameter is reducing log-linearly with time, which is sensible, and as a result the return levels also reduce over time, which is in line with the observed data. However, while this reduction seems sensible for most of the dataset, towards the end, the 20-observation return level is substantially below most of the data, suggesting a serious failure of the model to capture the correct tail behaviour at this point. This indicates that further investigation into suitable forms for the model parameters is required before we can feel satisfied that we have an appropriate model. Supporting this is the empirical return level exceedance rates of 24%, 8% and 4% being somewhat in excess of the model based rates of 5%, 2% and 1% respectively. More work is required here!

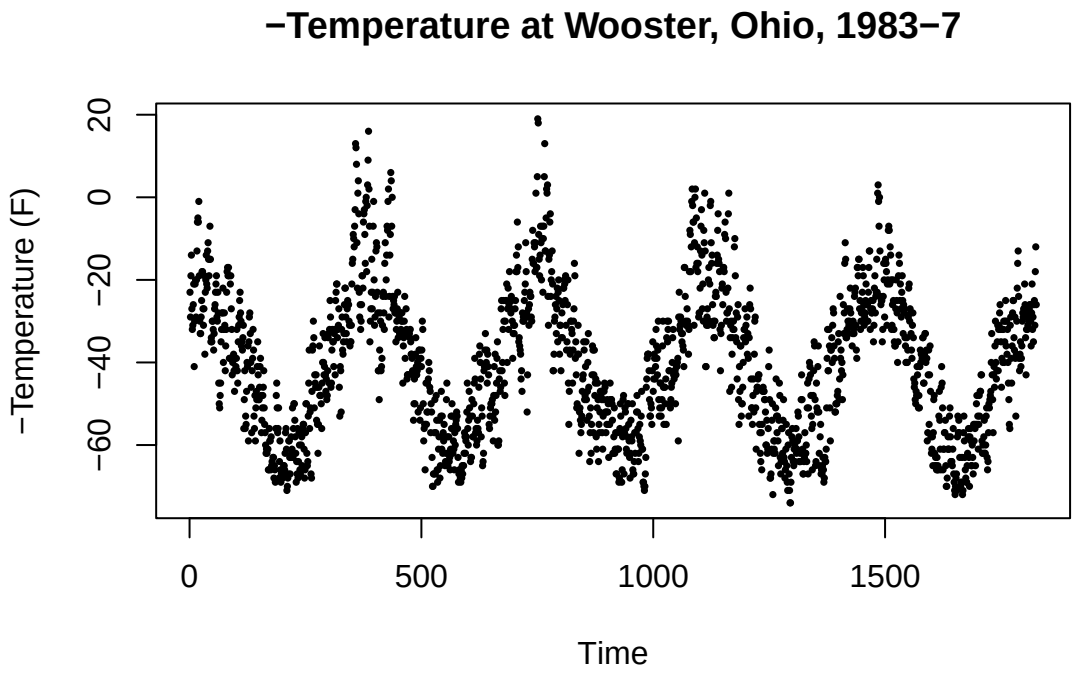
Exercise 2

The dataset `wooster` in the `R` package `ismev` contains 1826 daily minimum temperatures, in degrees Fahrenheit, at Wooster, Ohio, over the 5-year period from 1st January 1983 to 31st December 1987 (1984 is a leap year). We are interested in the extremely low temperatures.

a) Plot the (negative) temperature data, and observe the obvious seasonality. Is there any trend?

Solution:

```
library(ismev)
data(wooster)
w <- -wooster
plot(w, pch=16, cex=0.5, xlab="Time", ylab="-Temperature (F)",
     main="-Temperature at Wooster, Ohio, 1983-7")
```

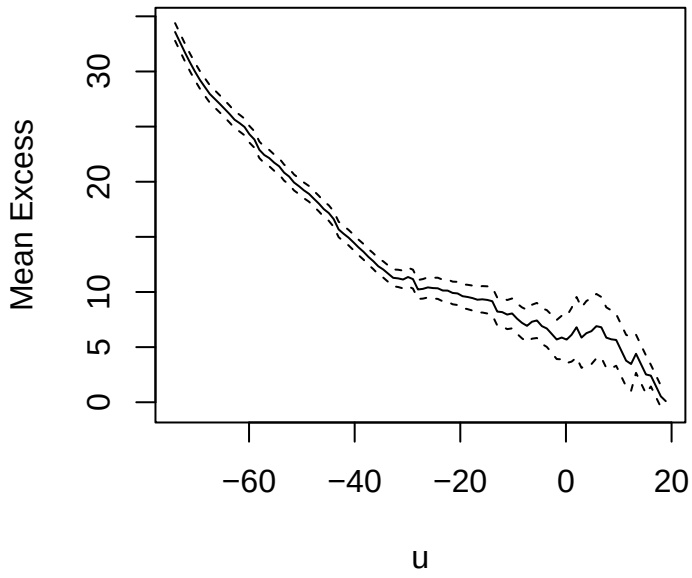


The obvious cyclic annual trend shows colder temperatures in the winter months and warmer in the summer months. There may(?) be a slight downwards trend over time (so in general, slightly warmer temperatures).

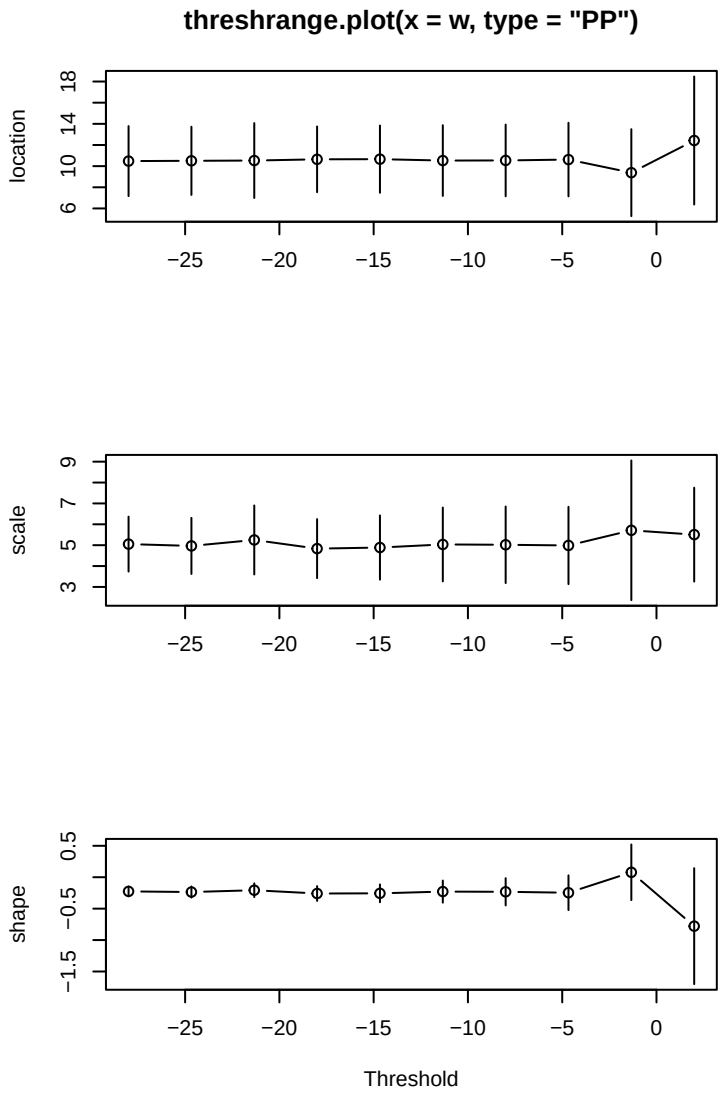
b) Use mean excess (mean residual life) plots and goodness-of-fit diagnostics to identify the most suitable (constant) threshold u for a point process threshold-exceedance model with constant parameters.

Solution:

```
mrl.plot(w)
```

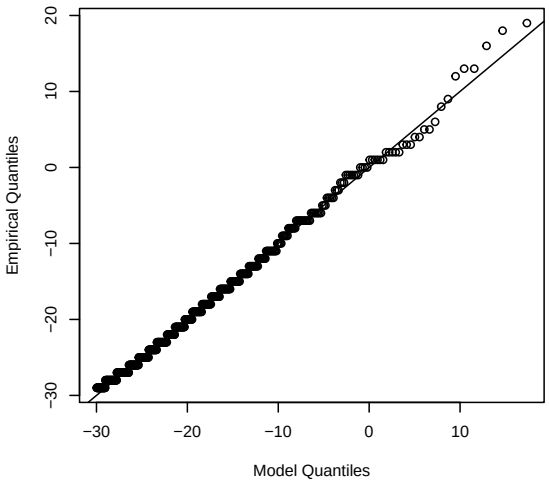


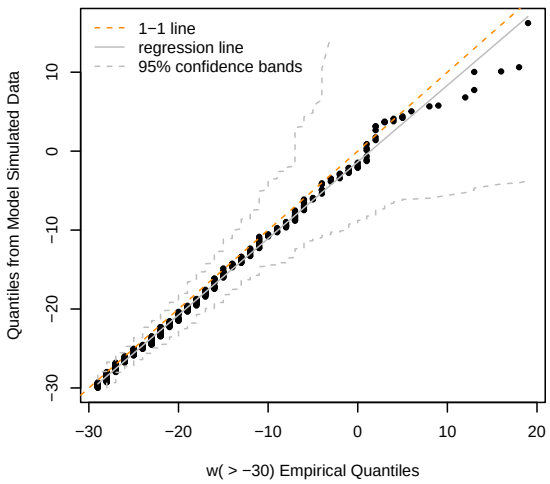

```
threshrange.plot(w,type="PP")
```

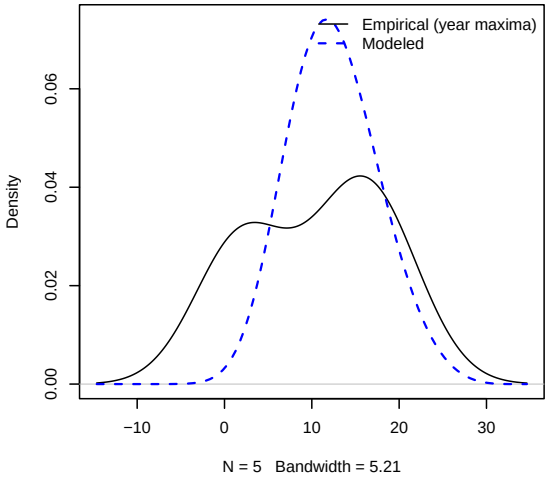


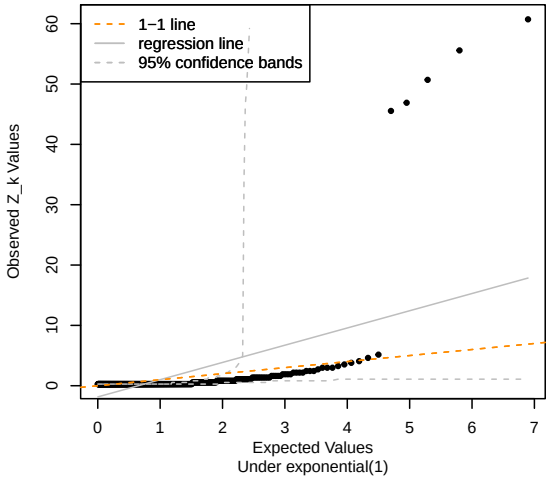

```
u <- -30 # Choose candidate threshold value here
fit <- fevd(w, threshold=u, type="PP", time.units="365.25/year")
plot(fit)
```

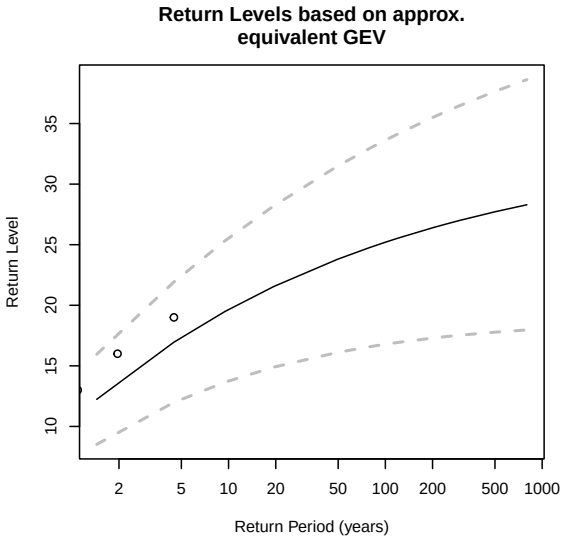
fevd(x = w, threshold = u, type = "PP", time.units = "365.25/year")











The mean excess plot suggests linearity above 5, which only gives 9 data points, so this is unreliable. Potentially the mean excess is linear above $u = -30$ or $u = -20$. The parameter estimates for the point process model fitted at different thresholds give no clues in this direction. Looking at the model diagnostics for the fitted model at $u = -30$, the qq-plot fit seems reasonable (plots for other thresholds not shown). Note that in the Z-plot (which tests how reasonably the distances between successive exceedances follow an exponential distribution), while the majority of distances are on the $y = x$ line, there are five very large outlier values, indicating that there is something wrong with our threshold-exceedance assumption. Obviously these five values are due to the large gap of exceedances between years when assuming a threshold of (say) $u = -30$. This is a natural consequence of the obvious seasonality in these data (which we have ignored, currently), and so increasing the threshold u will not improve the Poisson process fit.

For now we note that, apart from this obvious failure of the assumptions of the (Poisson) point process model, the rest of the model fit diagnostics look ok, and so we select $u = -30$ for now.

c) To what extent is this point process model, a good model for these data?

Solution:
There are a number of assumption violations in this model: threshold exceedances are not constant over time, and the data exhibit clear within-year seasonality which is not being accounted for. The analyst would not use this model with any confidence!

d) Retaining the constant threshold u, develop time varying point process model parameters, $\mu(t), \sigma(t), \xi(t)$, that better describe the non-stationarity that can be observed in the data. E.g., Consider adding the time-varying predictors

$$\cos\left(2\pi \times \text{doy}_t/365.25\right), \qquad \sin\left(2\pi \times \text{doy}_t/365.25\right),$$

for the seasonality, and t (time) for the potential trend, where doy_t is the day of the year (note: only 1984 is a leap year). Also consider whether it is better to model $\sigma(t)$ or $\log(\sigma(t))$ as a linear function of predictors.

Solution:

```
doy <- c(1:365,1:366,1:365,1:365,1:365)
X <- data.frame(time=seq(-1,1,length=length(w)), ccos=cos(2*pi*doy/365.25),
                ssin=sin(2*pi*doy/365.25))
fit.2 <- fevd(w, data=X, threshold=u, type="PP", location.fun=~ccos+ssin,
              time.units="365.25/year")
lr.test(fit, fit.2) # strong support p-value = 2.2e-16
```

```
##
##  Likelihood-ratio Test
##
## data:  ww
## Likelihood-ratio = 791.38, chi-square critical value = 5.9915, alpha =
## 0.0500, Degrees of Freedom = 2.0000, p-value < 2.2e-16
## alternative hypothesis: greater
```

```r
fit.3 <- fevd(w, data=X, threshold=u, type="PP", scale.fun=~ccos+ssin,
              time.units="365.25/year")
lr.test(fit, fit.3) # strong support p-value = 2.2e-16
```

```
##
##  Likelihood-ratio Test
##
## data:  ww
## Likelihood-ratio = 611.34, chi-square critical value = 5.9915, alpha =
## 0.0500, Degrees of Freedom = 2.0000, p-value < 2.2e-16
## alternative hypothesis: greater
```

```r
fit.4 <- fevd(w, data=X, threshold=u, type="PP", shape.fun=~ccos+ssin,
              time.units="365.25/year")
lr.test(fit, fit.4) # strong support p-value = 2.2e-16
```

```
##
##  Likelihood-ratio Test
##
## data:  ww
## Likelihood-ratio = 556.35, chi-square critical value = 5.9915, alpha =
## 0.0500, Degrees of Freedom = 2.0000, p-value < 2.2e-16
## alternative hypothesis: greater
```

```r
# fit.2 has the largest likelihood ratio, so take this.
fit.5 <- fevd(w, data=X, threshold=u, type="PP", location.fun=~ccos+ssin,
              scale.fun=~ccos+ssin, time.units="365.25/year")
lr.test(fit.2, fit.5) # no support
```

```
##
##  Likelihood-ratio Test
##
## data:  ww
## Likelihood-ratio = -18.619, chi-square critical value = 5.9915, alpha =
## 0.0500, Degrees of Freedom = 2.0000, p-value = 1
## alternative hypothesis: greater
```

```r
fit.6 <- fevd(w, data=X, threshold=u, type="PP", location.fun=~ccos+ssin,
              shape.fun=~ccos+ssin, time.units="365.25/year")
lr.test(fit.2, fit.6) # no support (p-value =1?)
```

```
##
##  Likelihood-ratio Test
##
## data:  ww
## Likelihood-ratio = 43.877, chi-square critical value = 5.9915, alpha =
## 0.0500, Degrees of Freedom = 2.0000, p-value = 2.966e-10
## alternative hypothesis: greater
```

```
fit.7 <- fevd(w, data=X, threshold=u, type="PP", location.fun=~ccos+ssin+time,
              time.units="365.25/year")
lr.test(fit.2, fit.7) # no support p=0.2
```

```
##
##  Likelihood-ratio Test
##
## data:  ww
## Likelihood-ratio = -0.11839, chi-square critical value = 3.8415, alpha
## = 0.0500, Degrees of Freedom = 1.0000, p-value = 1
## alternative hypothesis: greater
```

```
fit.8 <- fevd(w, data=X, threshold=u, type="PP", location.fun=~ccos+ssin,
              scale.fun=~time, time.units="365.25/year")
lr.test(fit.2, fit.8) # no support p=0.2
```

```
##
##  Likelihood-ratio Test
##
## data:  ww
## Likelihood-ratio = 1.4437, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 0.2295
## alternative hypothesis: greater
```

```
fit.9 <- fevd(w, data=X, threshold=u, type="PP", location.fun=~ccos+ssin,
              scale.fun=~time, use.phi=T, time.units="365.25/year")
lr.test(fit.2, fit.9) # moderate support p=0.09
```

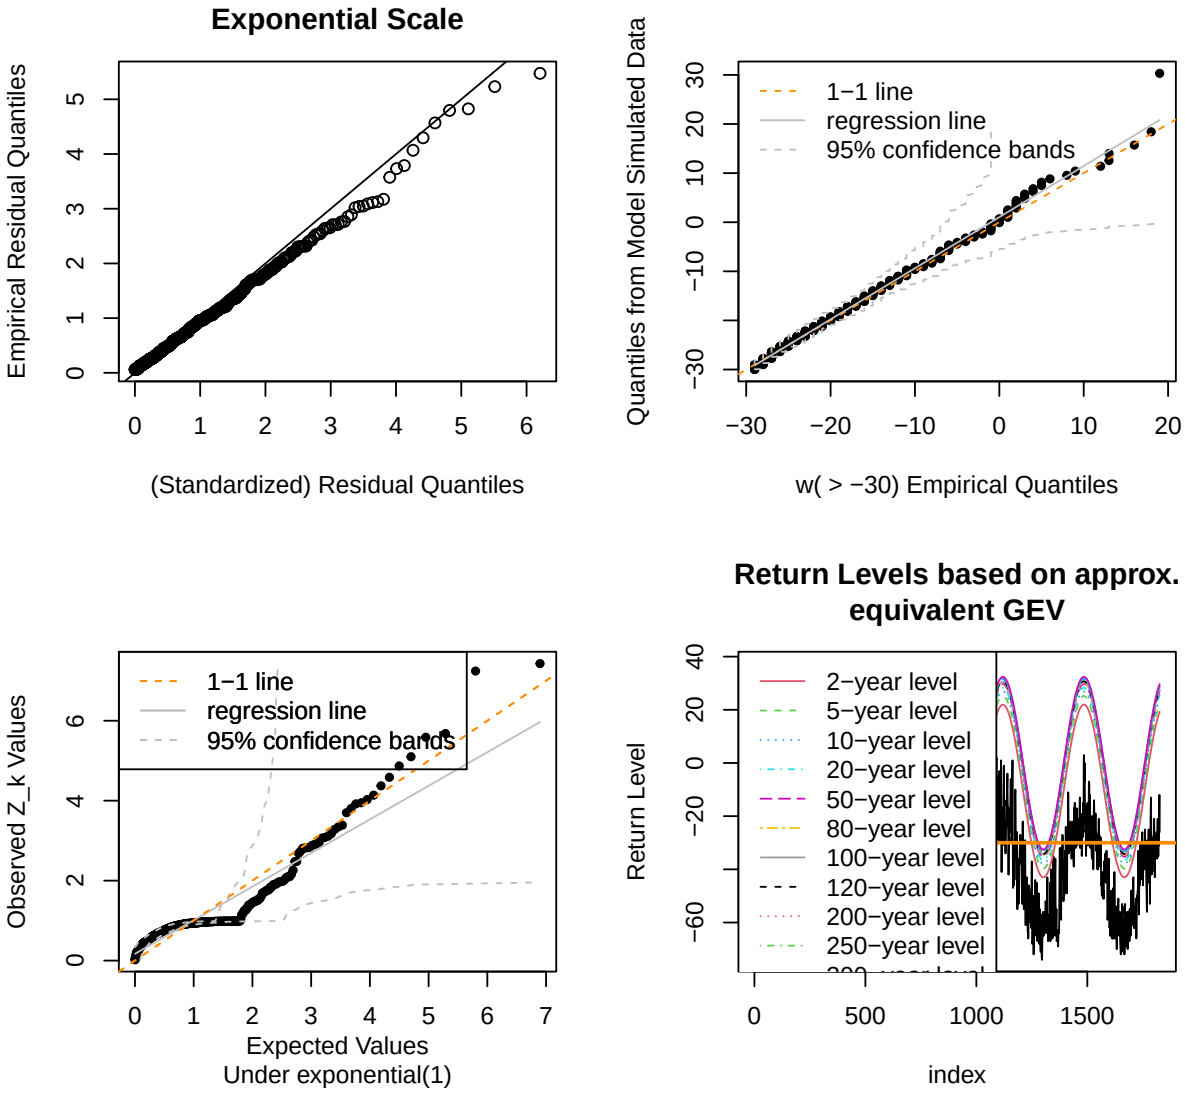
```
##
##  Likelihood-ratio Test
##
## data:  ww
## Likelihood-ratio = -21.705, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 1
## alternative hypothesis: greater
```

```
fit.10 <- fevd(w, data=X, threshold=u, type="PP", location.fun=~ccos+ssin,
               shape.fun=~time, time.units="365.25/year")
lr.test(fit.2, fit.10) # moderate support p=0.1
```

```
##
##  Likelihood-ratio Test
##
## data:  ww
## Likelihood-ratio = 5.8713, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 0.01539
## alternative hypothesis: greater
```

```
# So some potential evidence for changes over time, but not strong
plot(fit.2)
```

fevd(x = w, data = X, threshold = u, scale = sc, shape = xi, location = loc + cos + ssin,

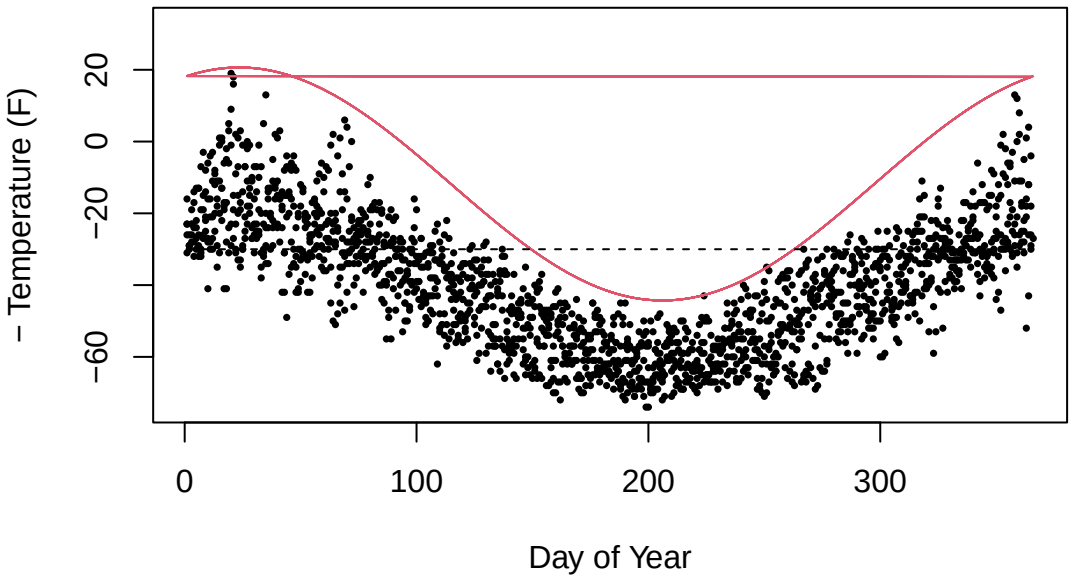


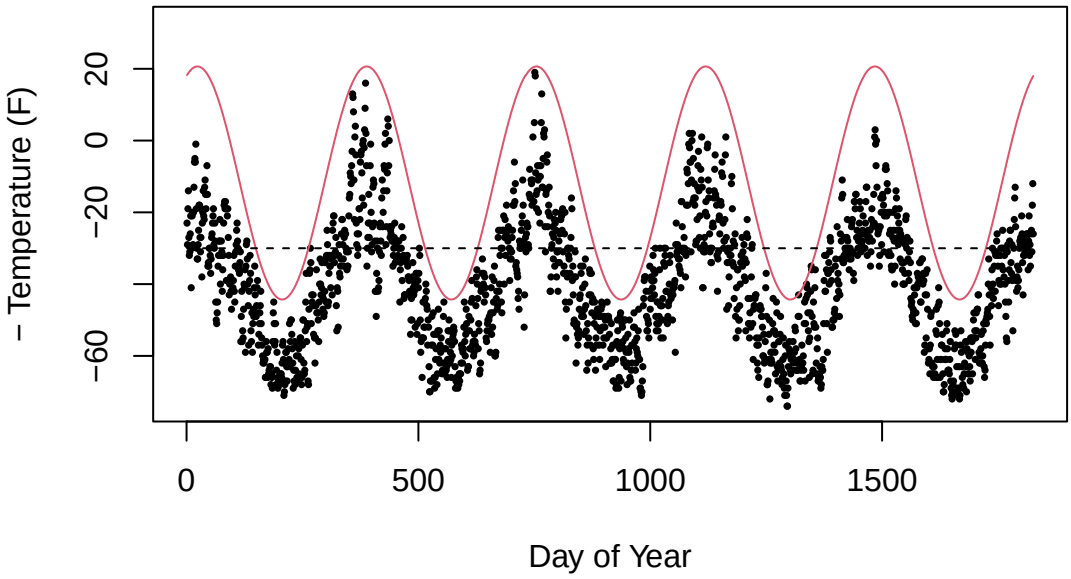
The qq-plot seems just as good as before. The Z-plot shows improved fit for the 5-largest distances between threshold exceedances, though the full distribution looks less than a perfect fit, suggesting that we are still some way from a Poisson point process representation. While any of the three point process parameters could significantly vary with time of year, the location was the strongest. Once this was added, there was no sufficiently strong (p-values between 0.1 and 0.2) for any other predictors, including linear time trends.

e) The qq-plot seems just as good as before. The Z-plot shows improved fit for the 5-largest distances between threshold exceedances, though the full distribution looks less than a perfect fit, suggesting that we are still some way from a Poisson point process representation. While any of the three point process parameters could significantly vary with time of year, the location was the strongest. Once this was added, there was no sufficiently strong (p-values between 0.1 and 0.2) for any other predictors, including linear time trends.

Solution:

```
p <- fit.2$results$par
p1 <- p[1]+p[2]*X$ccos+p[3]*X$ssin
par(mfrow=c(2,1))
plot(doy, w, pch=16, cex=0.5, ylim=c(min(w),33), xlab="Day of Year",
     ylab="- Temperature (F)")
lines(doy, p1, col=2)
lines(c(1,366), c(u,u), lty=2)
tt <- 1:length(w)
plot(tt, w, pch=16, cex=0.5, ylim=c(min(w),33), xlab="Day of Year",
     ylab="- Temperature (F)")
lines(tt, p1, col=2)
lines(c(1,length(tt)), c(u,u), lty=2)
```





The location parameter (red) increases when the (negative) temperatures are largest, which is sensible as this mirrors the variations in the observed data. However having a parameter that varies in a particular way when there is no data there is somewhat worrying – these parameter values are unreliable in regions where there are no data.

f) Is the time varying parameter point process model an improvement over the constant parameter model? To what extent is the constant threshold u responsible for any modelling failures?

Solution:

Restricting ourselves to considering data only above the threshold, there is no data in the middle of the year, so why should we believe that the extremes have a lower location parameter

in this region, when in fact we have effectively zero probability of observing a threshold exceedance in this time period? Because there is no data in the middle of the year, there is nothing to penalise this poor choice of model.

The poor model selection here is primarily driven by the obvious seasonality in threshold exceedances. The lack of the data in some regions is not penalising obvious drawbacks in the fitted models. The Poisson process model also assumes an exponential distribution between successive threshold exceedances, and time-homogeniety in this distribution. These assumptions are clearly violated here. It seems that we are unable to produce a credible model using time varying parameters and a constant threshold for these data.

g) Given the strong seasonality in this data, it would seem to make sense to develop a point process model with a time varying threshold $u(t)$. For example, the distribution of temperatures doesn't seem to change too much beyond the obvious location shift, and so a time varying threshold could normalise these variations and leave a relatively simple model for the exceedances.

While there are no easy tools to identify such threshold in the `R` packages that we are using (thought question: how might you identify a time varying threshold as part of the inference process?), trial and error suggests that a possible threshold could be given by

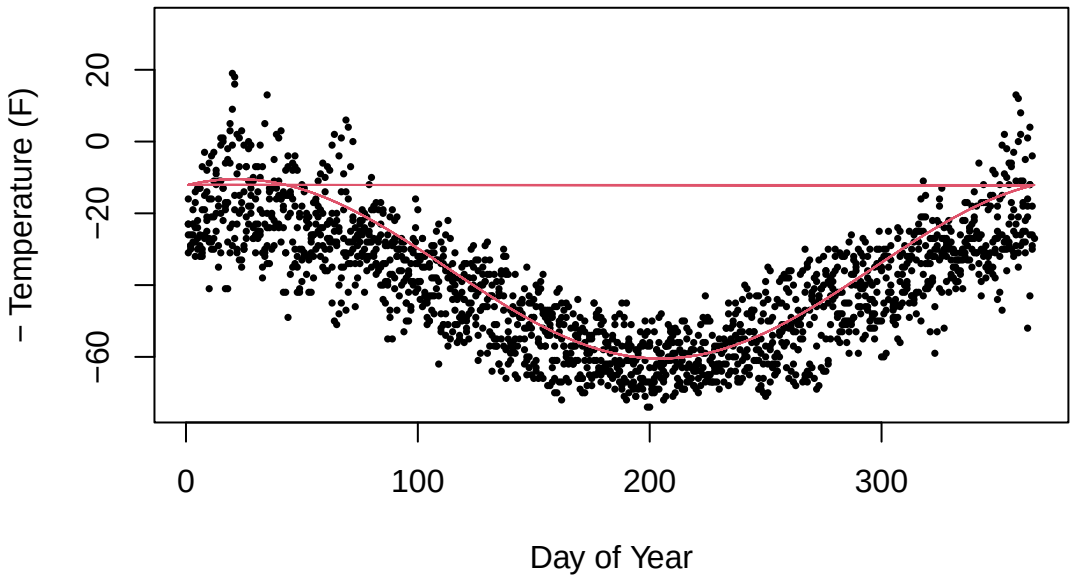
$$u(t) = 25\cos(2\pi(\text{doy}_t - 17)/366) + \bar{x}_t + 5$$

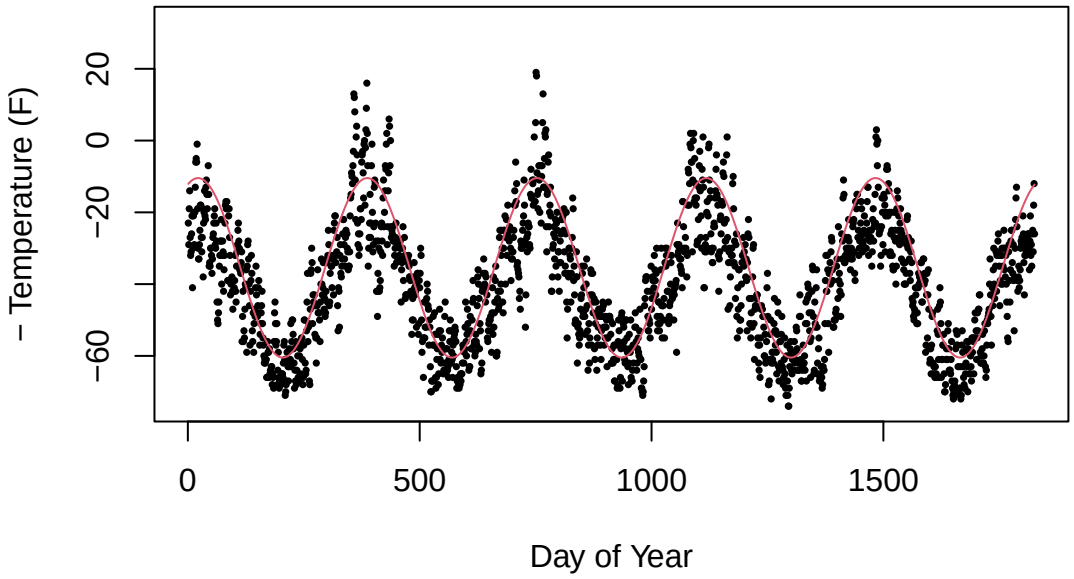
where doy_t is the day of year for observation x_t (note: `doy=c(1:365,1:366,1:365,1:365,1:365)` in `R`).

Examine plots of threshold exceedances over time and versus day of year. Does a variable threshold model seem more likely to produce useful results? Is this threshold a suitable one?

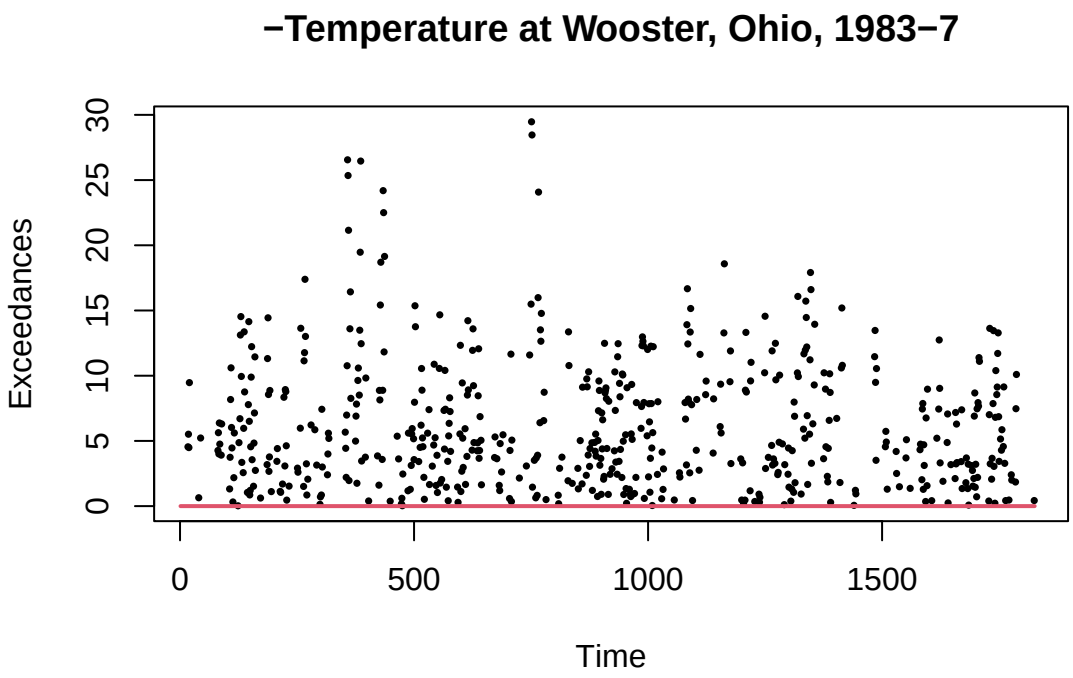
Solution:

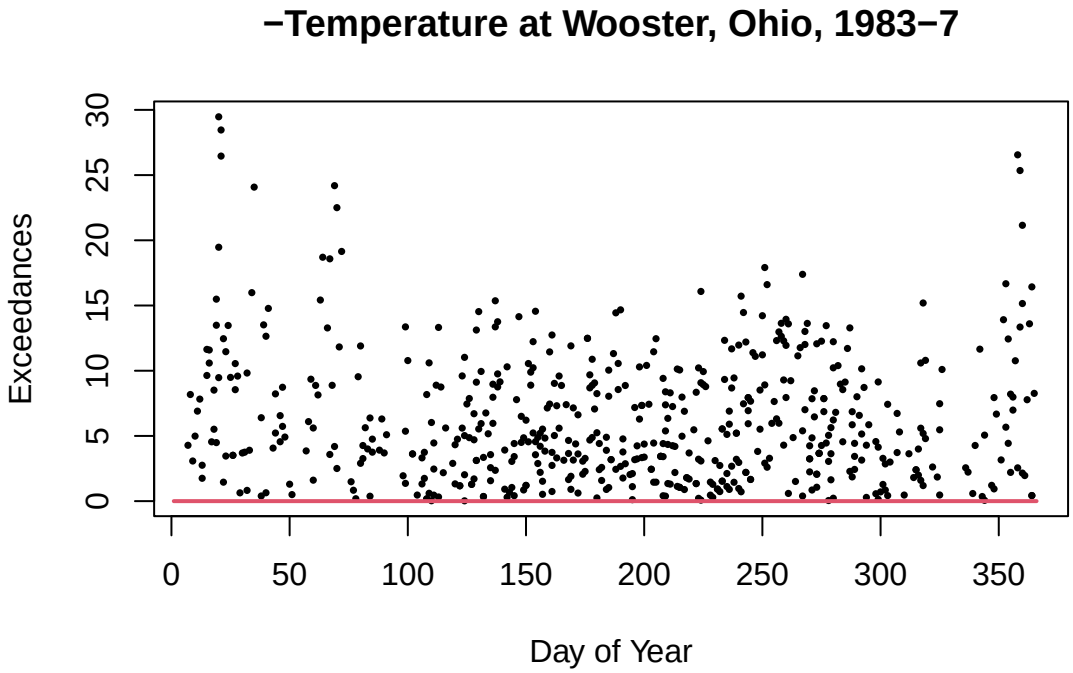
```
ut <- 25 * cos(2*pi*(doy-22)/365.25) + mean(w) + 5
par(mfrow=c(2,1))
plot(doy, w, pch=16, cex=0.5, ylim=c(min(w),33), xlab="Day of Year",
     ylab="- Temperature (F)")
lines(doy, ut, col=2)
plot(tt, w, pch=16, cex=0.5, ylim=c(min(w),33), xlab="Day of Year",
     ylab="- Temperature (F)")
lines(tt, ut, col=2)
```



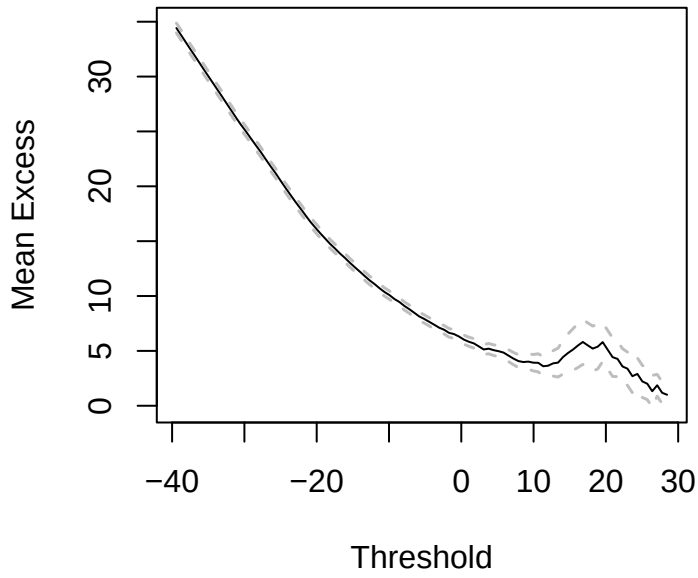



```
r <- w-ut # r is wut?
L <- length(r)
par(mfrow=c(2,1))
plot((1:L)[r>0], r[r>0], pch=16, cex=0.5, xlab="Time", ylab="Exceedances",
     main="-Temperature at Wooster, Ohio, 1983-7")
lines(c(1,length(r)), c(0,0), col=2, lwd=2)
plot(doy[r>0], r[r>0], pch=16, cex=0.5, xlab="Day of Year",
     ylab="Exceedances", main="-Temperature at Wooster, Ohio, 1983-7")
lines(c(1,366), c(0,0), col=2, lwd=2)
```



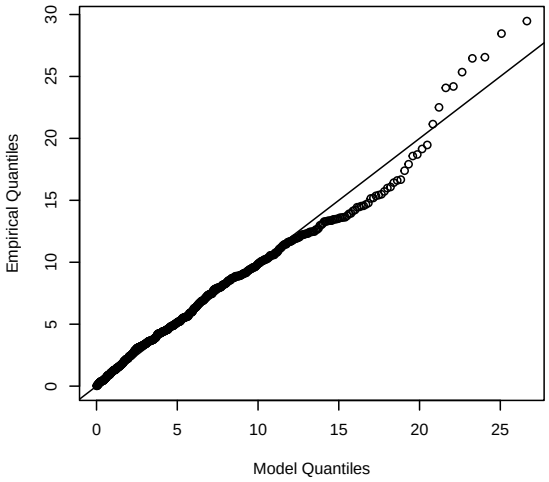


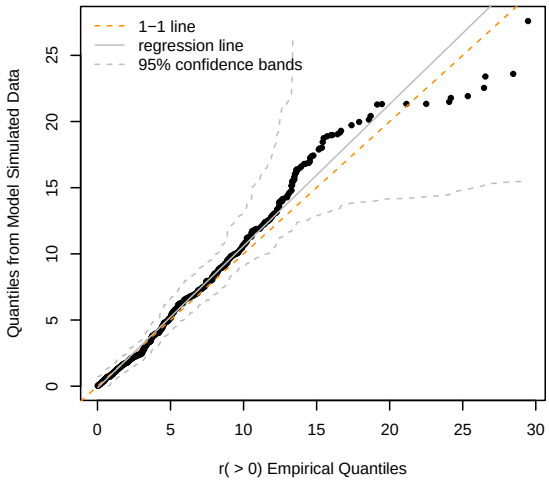

```
par(mfrow=c(1,1))
mrlplot(r)
```

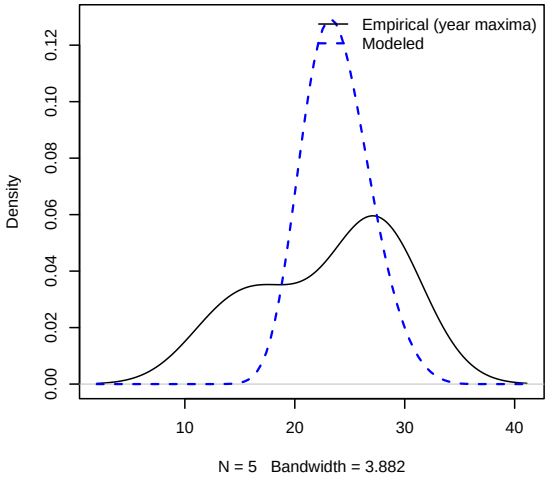


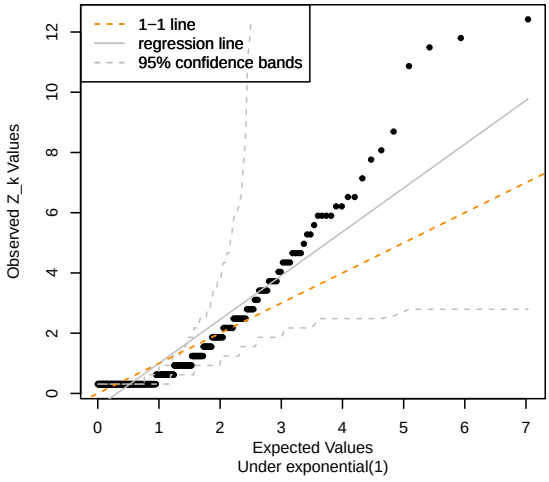

```
fit <- fevd(r, threshold=0, type="PP", time.units="365.25/year")
plot(fit)
```

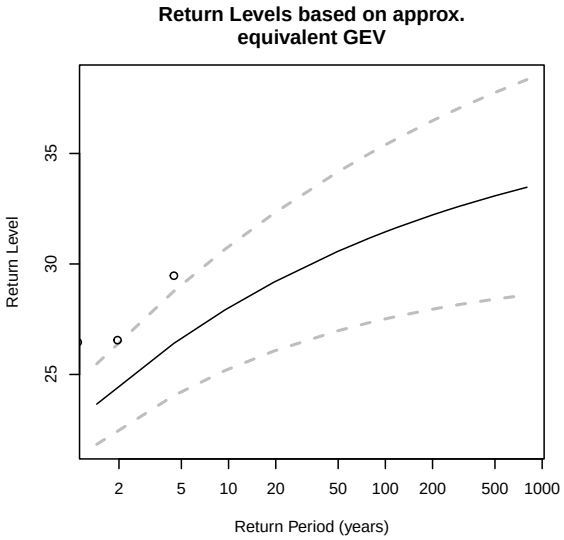
fevd(x = r, threshold = 0, type = "PP", time.units = "365.25/year")











Compared to plots of the original data over a constant threshold, these threshold exceedances have substantially less structure, which is good for the representation of the location of the threshold exceedances as a poisson process. In principle, if we can choose a threshold so that we eliminate most of the time-clustering of exceedances, we can then model the exceedances via a threshold exceedance model in the usual way.

Here the chosen threshold is borderline acceptable as there is still some clear time clustering in the exceedances over time. Ideally this would be removed by a better choice of threshold. Against day of year it is clear that exceedances are still more likely around the middle of the year compared to the endpoints. There is some indication that larger extremes are observable in the winter months (which makes sense), although the plot over time indicates that this is primarily driven by two cold years (out of 5), so it is not clear if this is a general quality of the data.

Overall, the chosen threshold is not great (the distribution of inter-exceedance times is not as exponential as we might like), but finding a better one will take a little more work.

h) Using this variable threshold identify and fit a suitable point-process threshold exceedance model, with time varying parameters.

Solution:

```
iind <- (doy<=90 | doy>=330) # two potential seasons?
X <- data.frame(time=1:length(r), iind=iind, ccos=cos(2*pi*doy/365.25),
                ssin=sin(2*pi*doy/365.25))
fit.1 <- fevd(r, data=X, threshold=0, location.fun=~iind, type="PP",
              time.units="365.25/year")
lr.test(fit, fit.1) # strong support
```

```
##
##  Likelihood-ratio Test
##
## data:  rr
## Likelihood-ratio = 21.789, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 3.043e-06
## alternative hypothesis: greater
```

```
fit.2 <- fevd(r, data=X, threshold=0, scale.fun=~iind, type="PP",
              time.units="365.25/year")
lr.test(fit, fit.2) # overwhelming support
```

```
##
##  Likelihood-ratio Test
##
## data:  rr
## Likelihood-ratio = 47.69, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 4.992e-12
## alternative hypothesis: greater
```

```
fit.3 <- fevd(r, data=X, threshold=0, scale.fun=~iind, use.phi=T, type="PP",
              time.units="365.25/year")
lr.test(fit, fit.3) # identical to fit.2
```

```
##
##  Likelihood-ratio Test
##
## data:  rr
## Likelihood-ratio = 47.69, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 4.992e-12
## alternative hypothesis: greater
```

```
fit.4 <- fevd(r, data=X, threshold=0, shape.fun=~iind, type="PP",
              time.units="365.25/year")
lr.test(fit, fit.4) # overwhelming support
```

```
##
##  Likelihood-ratio Test
##
## data:  rr
## Likelihood-ratio = 49.07, chi-square critical value = 3.8415, alpha =
## 0.0500, Degrees of Freedom = 1.0000, p-value = 2.47e-12
## alternative hypothesis: greater
```

```
fit.5 <- fevd(r, data=X, threshold=0, location.fun=~ccos+ssin, type="PP",
              time.units="365.25/year")
lr.test(fit, fit.5) # strong support
```

```
##
##  Likelihood-ratio Test
##
## data:  rr
## Likelihood-ratio = 38.509, chi-square critical value = 5.9915, alpha =
## 0.0500, Degrees of Freedom = 2.0000, p-value = 4.343e-09
## alternative hypothesis: greater
```

```
fit.6 <- fevd(r, data=X, threshold=0, scale.fun=~ccos+ssin, type="PP",
              time.units="365.25/year")
lr.test(fit, fit.6) # strongest overwhelming support
```

```
##
##  Likelihood-ratio Test
##
## data:  rr
## Likelihood-ratio = 65.362, chi-square critical value = 5.9915, alpha =
## 0.0500, Degrees of Freedom = 2.0000, p-value = 6.408e-15
## alternative hypothesis: greater
```

```
fit.7 <- fevd(r, data=X, threshold=0, shape.fun=~ccos+ssin, type="PP",
              time.units="365.25/year")
```

```r
lr.test(fit, fit.7) # overwhelming support
```

```
##
##  Likelihood-ratio Test
##
## data:  rr
## Likelihood-ratio = 64.185, chi-square critical value = 5.9915, alpha =
## 0.0500, Degrees of Freedom = 2.0000, p-value = 1.154e-14
## alternative hypothesis: greater
```

```r
#
fit.8 <- fevd(r,data=X, threshold=0, location.fun=~ccos+ssin,
              scale.fun=~ccos+ssin, type="PP", time.units="365.25/year")
lr.test(fit.6, fit.8) # strong support
```

```
##
##  Likelihood-ratio Test
##
## data:  rr
## Likelihood-ratio = 27.031, chi-square critical value = 5.9915, alpha =
## 0.0500, Degrees of Freedom = 2.0000, p-value = 1.35e-06
## alternative hypothesis: greater
```

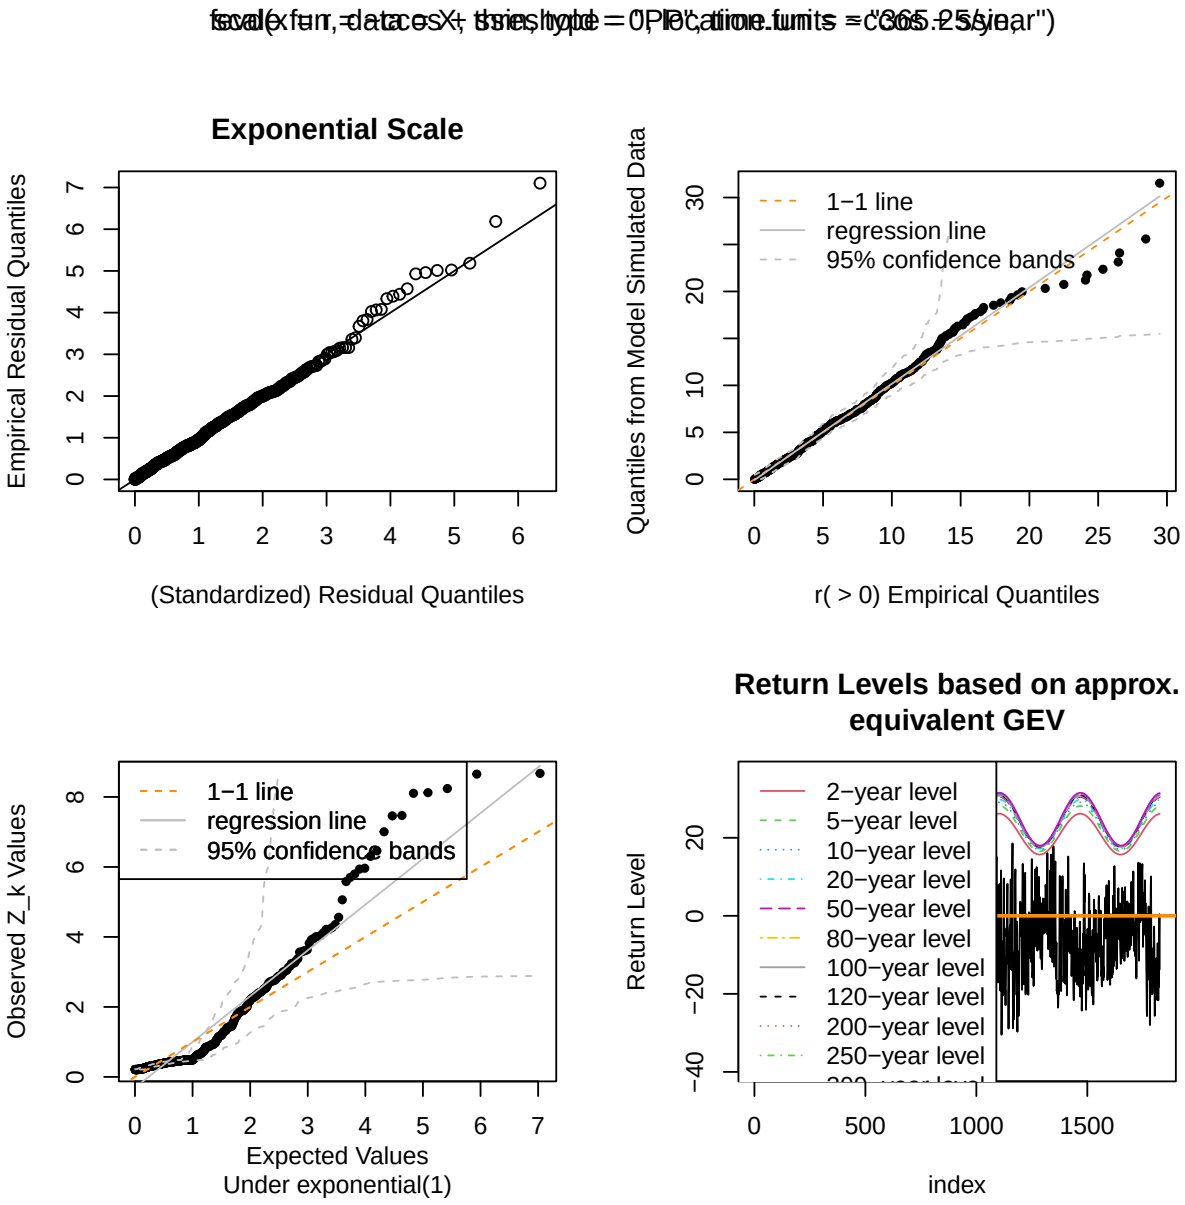
```r
fit.9 <- fevd(r, data=X, threshold=0, shape.fun=~ccos+ssin,
              scale.fun=~ccos+ssin, type="PP", time.units="365.25/year")
lr.test(fit.6, fit.9) # no support
```

```
##
##  Likelihood-ratio Test
##
## data:  rr
## Likelihood-ratio = 1.7349, chi-square critical value = 5.9915, alpha =
## 0.0500, Degrees of Freedom = 2.0000, p-value = 0.42
## alternative hypothesis: greater
```

```r
#
fit.10 <- fevd(r, data=X, threshold=0, location.fun=~ccos+ssin,
               scale.fun=~ccos+ssin, shape.fun=~ccos+ssin,type="PP",
               time.units="365.25/year")
lr.test(fit.8, fit.10) # no support
```

```
##
##  Likelihood-ratio Test
##
## data:  rr
## Likelihood-ratio = 0.48577, chi-square critical value = 5.9915, alpha =
## 0.0500, Degrees of Freedom = 2.0000, p-value = 0.7844
## alternative hypothesis: greater
```

```
#
plot(fit.8) # Final model
```



Motivated by the exceedance versus day of year plot we create an indicator variable (`iind`) for two different seasons, allowing for different (constant) parameter values in each season. We also include the usual sin and cos representation. Compared to the constant parameter model any predictor as part of any parameter makes a significant improvement. The strongest of these is writing $\sigma(t)$ as a function of `ccos` and `ssin` (modelling on the log scale does not make any difference). Further improvement is also obtained by writing $\mu(t)$ as a function of `ccos` and `ssin`.

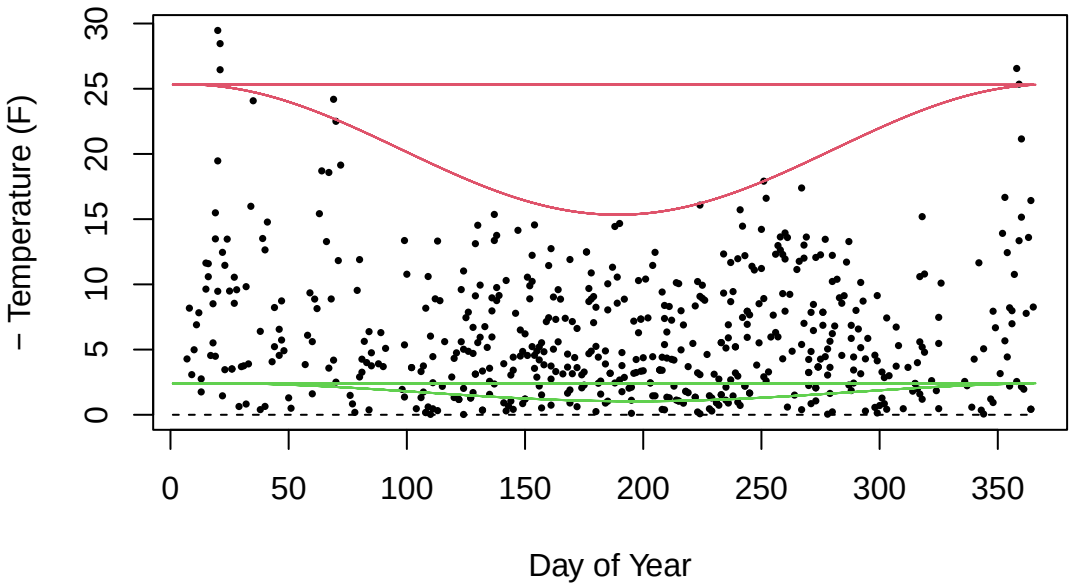
Compared to the constant parameter model, the threshold exceedances have an improved GPD fit (the distribution of threshold exceedance occurences is, of course, the same).

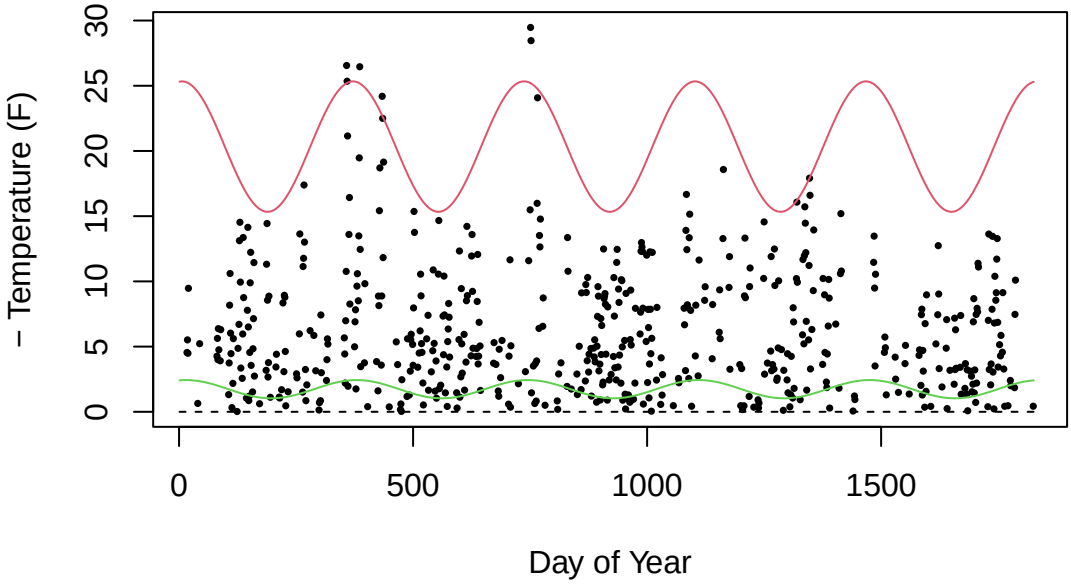
i) Superimpose the parameter variation on plots of the threshold exceedance data versus time

and day of year. Do these parameter variations explain these the data?

Solution:

```
p <- fit.8$results$par
p1 <- p[1] + p[2] * X$ccos + p[3] * X$ssin
p2 <- p[4] + p[5] * X$ccos + p[6] * X$ssin
par(mfrow=c(2,1))
plot(doy[r>0], r[r>0], pch=16, cex=0.5, xlab="Day of Year",
     ylab="- Temperature (F)")
lines(doy, p1, col=2)
lines(doy, p2, col=3)
lines(c(1,366), c(0,0), lty=2)
tt <- 1:length(r)
plot(tt[r>0], r[r>0], pch=16, cex=0.5, xlab="Day of Year",
     ylab="- Temperature (F)")
lines(tt, p1, col=2)
lines(tt, p2, col=3)
lines(c(1,length(tt)), c(0,0), lty=2)
```



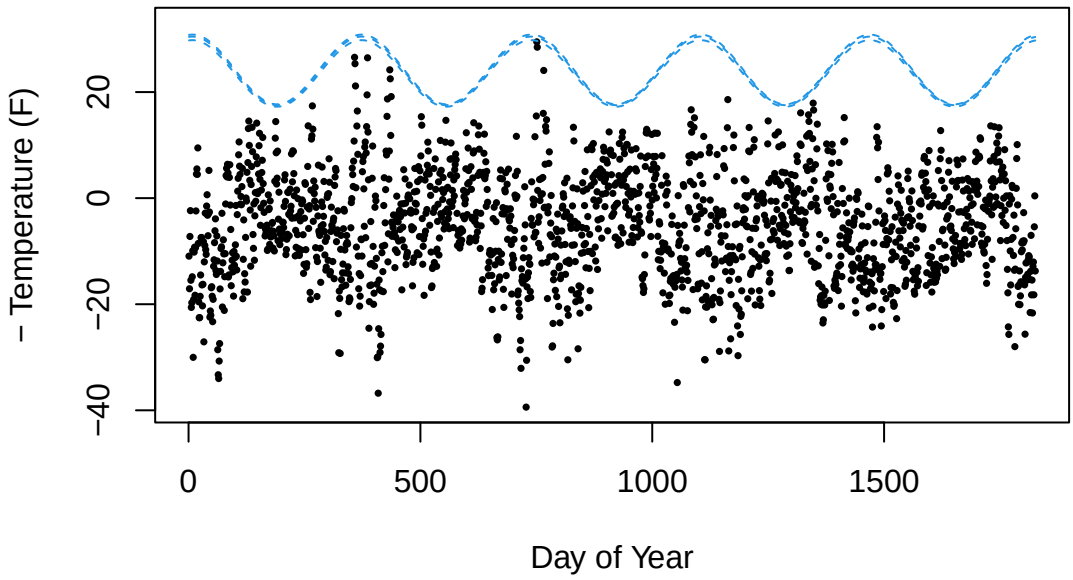


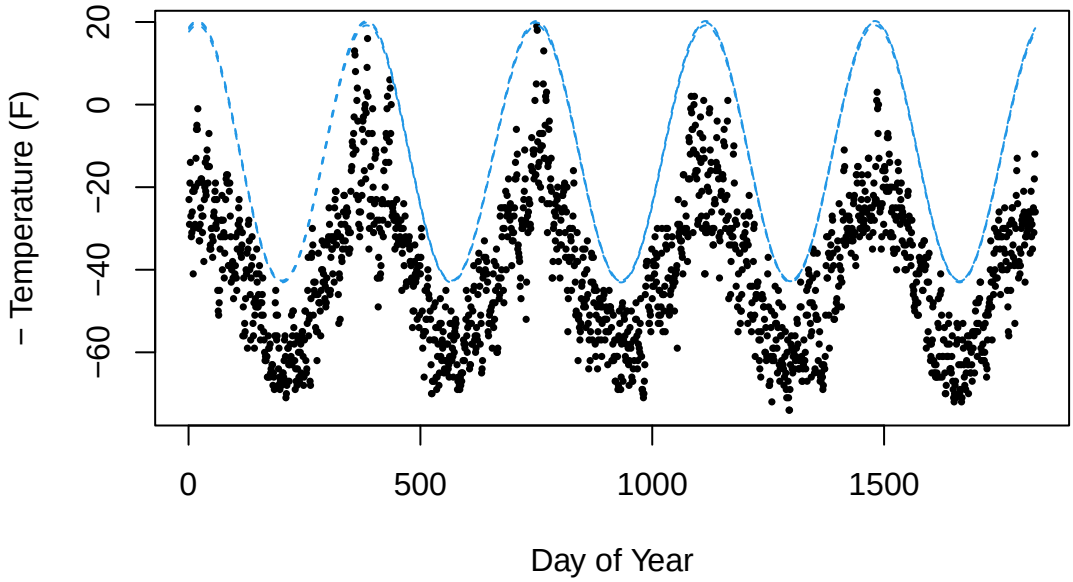
The larger extremes at the ends of each year are accounted for by an increase in both location and scale parameters, which seems entirely reasonable. (The shape parameter remains constant.)

j) Calculate indicative 20-year, 50-year and 100-year return levels over the range of the dataset, and superimpose on plots of both threshold exceedances and of the original data. Comment on their features. ('Indicative' here means to calculate an x-year return level given the parameter values present at each time point. The resulting x-year return level is the level we would expect to be exceeded on average once every x years if the parameter values were constant at the specified levels.)

Solution:

```
par(mfrow=c(2,1))
plot(tt, r, pch=16, cex=0.5, ylim=c(min(r),33), xlab="Day of Year",
     ylab="- Temperature (F)")
v <- make.qcov(fit.8, vals=list(mu1=X$ccos, mu2=X$ssin, sigma1=X$ccos,
               sigma2=X$ssin))
a20 <- return.level(fit.8, return.period=c(20), qcov=v)
a50 <- return.level(fit.8, return.period=c(50), qcov=v)
a100 <- return.level(fit.8, return.period=c(100), qcov=v)
lines(a20, col=4, lty=2)
lines(a50, col=4, lty=2)
lines(a100, col=4, lty=2)
plot(tt, w, pch=16, cex=0.5, xlab="Day of Year", ylab="- Temperature (F)")
lines(a20+ut, col=4, lty=2)
lines(a50+ut, col=4, lty=2)
lines(a100+ut, col=4, lty=2)
```





The return levels clearly increase and decrease appropriately depending on the time of year, so that what is considered an extremely low temperature in winter is vastly different to what is considered an extremely low temperature in summer.

Interestingly the residuals from the time varying threshold have disturbed the obvious discretisation in the original dataset.

Note that the (indicative) 20-year, 50-year and 100-year return levels are very close to eachother. This is due to the strongly negative estimate of the shape parameter $\hat{\xi} = -0.3554189$ indicating a finite upper endpoint of the temperature distribution ($\hat{\xi} > -0.5$ so the usual MLE asypmtotics exist). Physically this might be questionable – it would be more intuitive if the shape parameter

was ≥ 0. In combination with the choice of time-varying threshold, further investigation into suitable forms of regression equations for each of the model parameters is probably warranted before we can be confident that we have implemented a good analysis here. We might be suspicious that while the idea of a time varying threshold is a good one, perhaps our choice of regression models are poor in this case.