```
res1=fit.garch11.normal(d)       # Fit GARCH(1,1) assuming Z(t)~N(0,1)
                                 # This is the same as the quasi maximum likelihood
                                 # which assumes the Z(t) are normal
res2=fit.garch11.t(d)            # Fit GARCH(1,1); Z(t)~t_nu with E[Z]=0, Var[Z]=1
res3=fit.garch11.nct(d)          # Fit GARCH(1,1); Z(t)~NCT with E[Z]=0, Var[Z]=1
```

Type res1$par, res2$par, res3$par in sequence to compare the parameter estimates:

```
> res1$par
   sigma(1)       alpha0       alpha1        beta1
9.708470e-03 1.072343e-06 7.262521e-02 9.200175e-01
          mu
5.186208e-04
> res2$par
   sigma(1)       alpha0       alpha1        beta1
2.311459e-02 1.335119e-06 6.676990e-02 9.184381e-01
          mu           DF
5.606958e-04 6.633855e+00
> res3$par
    sigma(1)        alpha0        alpha1         beta1
 7.105520e-03  9.443817e-07  5.886972e-02  9.330010e-01
           mu            DF           ncp
 4.452859e-04  6.413187e+00 -2.852327e-01
```

The sigma1 parameter estimates seem to vary quite a bit. However, this can be considered to be a form of nuisance parameter as it has little influence over forecasts after time *n*. The other parameter estimates alpha0, alpha1 and beta1 have some similarities, but there is some variation between them, suggesting that full maximum likelihood is best when the distribution of Z(t) has been established.

```
ZQ=res1$Z       # QML volatility standardised residuals
                # Also the true volatility standardised residuals if the
                # Normal hypothesis is correct
```

####### Exercise 1 ######

Now test for the following questions:
Plot the ZQ vector
Are the ZQ i.i.d.?

Plot the ACF's of the ZQ's and the ZQ-squared (see below). The acf of ZQ shows a slight positive autocorrelation at lag 1 which does not cause me huge concern. The ACF of ZQ^2 is consistent with the data being i.i.d.

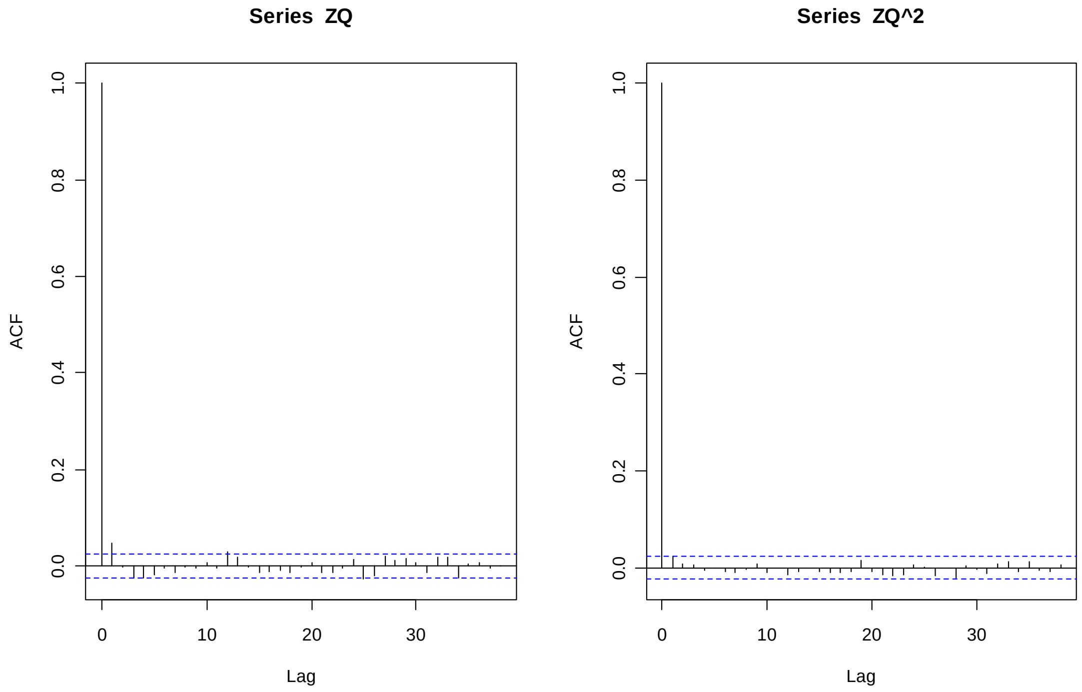

Now calculate the interexceedance times (a) for ZQ greater than its 95% quantile, (b) |ZQ| greater than its 95% quantile:

```
> tv=1:length(ZQ)
> T1=tv[ZQ > quantile(ZQ,0.95)]
> S1=diff(T1)
> qqgeometric(S1)
> abline(0,1)
> T2=tv[abs(ZQ) > quantile(abs(ZQ),0.95)]
> S2=diff(T2)
> qqgeometric(S2)
> abline(0,1)
```

The QQ plots below (ZQ on the left, |ZQ| on the right) look reasonably linear. However, the left-hand plot suggests some structure. However, if you try

```
> qqgeometric(rgeometric(325,0.95))
> abline(0,1)
```

and repeat a few times, you can see that the structure in the left hand plot might be caused by randomness in the data. So these plus the ACF's give reasonable but not 100% support for the i.i.d. hypothesis.

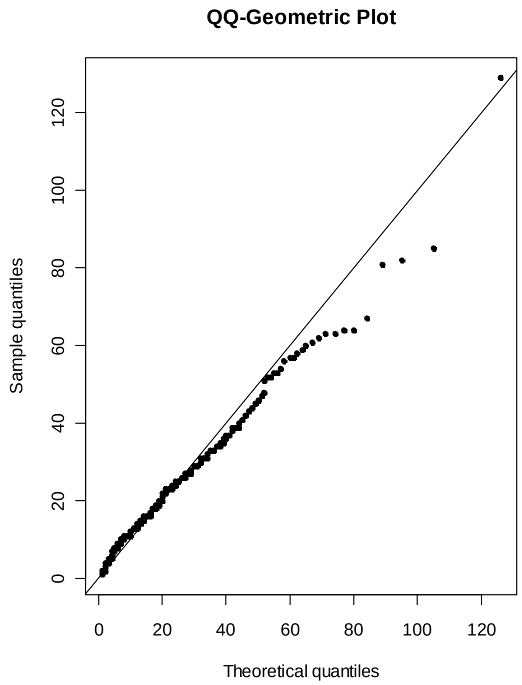

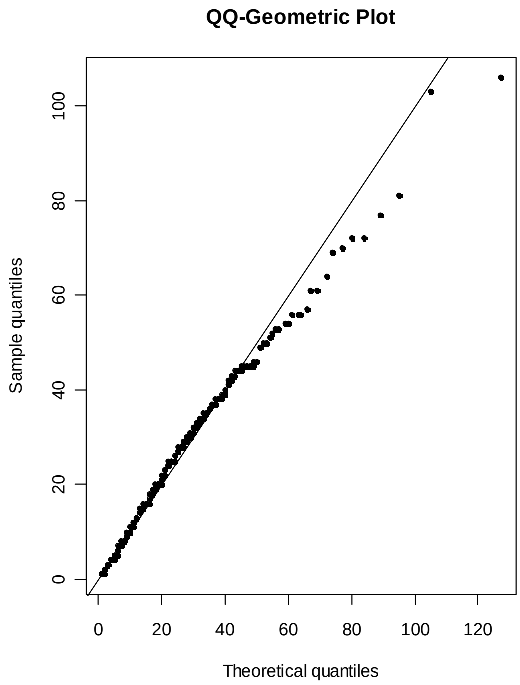

Are the ZQ Normally distributed?

Do a QQ-plot against the theoretical normal. See below. Clearly non linear. But the QQ plot also reveals that the underlying distribution is probably negatively skewed with a fatter left hand tail than the right hand tail.

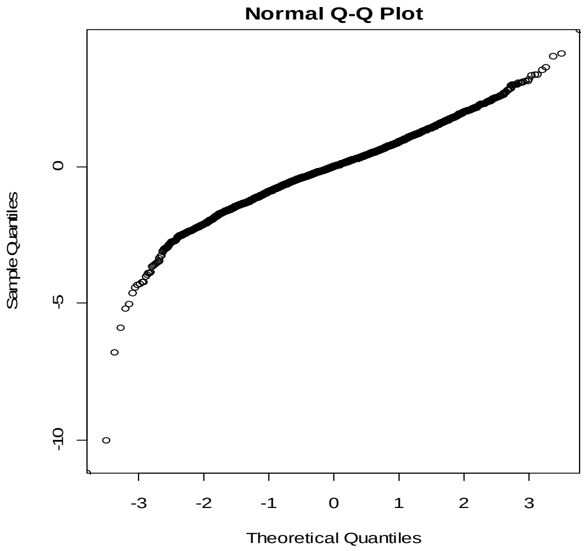


```
# Work out how to simulate the next day's daily log return as follows:
# Suppose that you are using the GARCH(1,1)-Normal model
# First work out sigma(n+1) given the values of sigma(n), alpha0,
# alpha1, beta1, and Z(n) from the res1 output.

> alpha0=res1$par[2]
> alpha1=res1$par[3]
> beta1=res1$par[4]
> n0=length(ZQ)
> sigma0=res1$sigma[n0] # the value of sigme for time n
> Z0=ZQ[n0] # the final estimated value of Z(n)
> sigma1.squared=alpha0+(alpha1*Z0^2+beta1)*sigma0^2
> sigma1=sqrt(sigma1.squared) # the estimate of sigma[n+1]
> sigma1
     alpha0
0.007417527
> attributes(sigma1)=NULL
> sigma1
[1] 0.007417527
> res1$sigma[n0]
[1] 0.00734467
```

The last line provides a check that sigma(n+1) is close to the final sigma(n).

```
> update.sigma.R(res1)
[1] 0.007417527

# Then note that Z(n+1) is N(0,1).
# Hence delta(n+1), conditional of the history up to time n,
# is normal with mean mu and variance sigma(n+1)^2.

> xv=((-500):500)/10000
> setaxes(-0.05,0.05,0,60)
```

```
> lines(xv,dnorm(xv,mu,sigma1))
> abline(v=0,lwd=0.5)
> title(xlab="delta(n+1)",ylab="density")
```

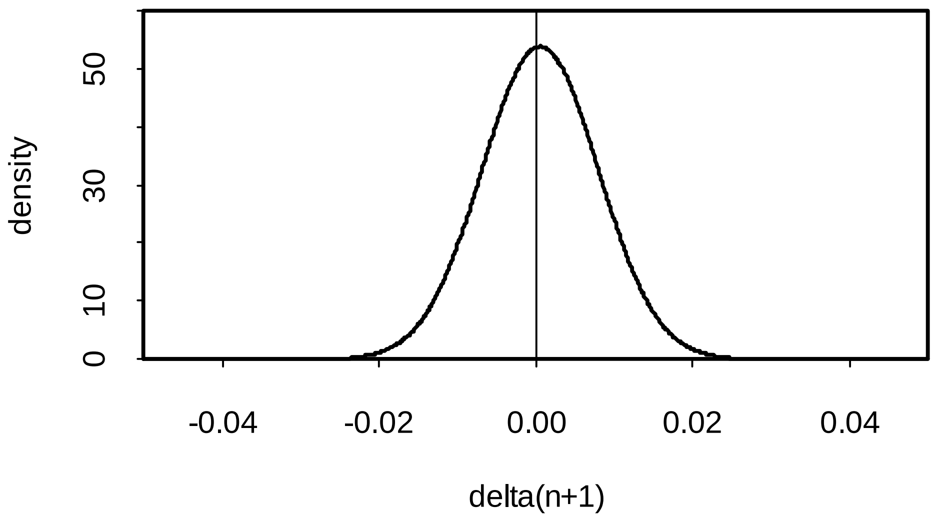

The plot above shows the predicted normal distribution for delta(n+1) given delta(1),…,delta(n) and the GARCH(1,1)-normal model fitted to that historical data. You can see that the impact of mu>0 is barely noticeable in relation to the uncertainty.

```
# Next do the same for the t distribution. But you will need to
# work out the number of degrees of freedom, the location and the
# scaling parameter for the three parameter t distribution
# for the conditional distribution of delta(n+1).
# Finally, do the same for the NCT. This is much more difficult!
# You will have to first work out the location and the scaling parameter
# parameters are for the NCT given that the Z(t) have mean 0 and
# variance 1. You will need to (a) look up Chapter 4
# to get the formulae for E[X] and E[X^2], where X has a standard
# NCT distribution given the degrees of freedom and the non-centrality
# parameter NCP. Then work out constants a=a(nu,NCP) and b=b(nu,NCP)
# that result if E[a+bX]=0 and Var(a+bX)=1.
# Define Z(t)=a+b.X(t) where the X(t) are i.i.d. standard NCT(nu,NCP)'s.
# Finally, delta(n+1) = mu + sigma(n+1) Z(n+1)
#              = mu + sigma(n+1) [ a + b.Z(n+1) ]
#              = [mu + sigma(n+1).a] + [sigma(n+1).b].X(n+1)
# i.e. NCT with nu degrees of freedom, non-centrality parameter NCP
# location parameter [mu + sigma(n+1).a] and scaling parameter
# [sigma(n+1).b]
```

```
> alpha0=res3$par[2]
> alpha1=res3$par[3]
> beta1=res3$par[4]
> Z0=res3$Z[n0]
```

```
> sigma0=res3$sigma[n0]
> sigma1.squared=alpha0+(alpha1*Z0^2+beta1)*sigma0^2
> sigma1=sqrt(sigma1.squared) # the estimate of sigma[n+1]
> sigma1
0.007247292
> sigma0
[1] 0.007183485
> mu=res3$par[5]
> df=res3$par[6]
> ncp=res3$par[7]
> # Now let X be a standard NCT with parameters df and ncp
> muX=sqrt(df/2)*gamma(0.5*(df-1))/gamma(0.5*df)*ncp      # E[X]
> muX2=df/(df-2)*(1+ncp*ncp)                               # E[X^2]
> sdX=sqrt(muX2-muX^2)                                     # S.D.[X]
> # Hence Z(n+1) = (X – muX)/sdX has mean 0 and variance 1
> location=mu-sigma1*muX/sdX
> scale=sigma1/sdX
```

Therefore the paramaters of the NCT distribution are as follows

```
> location
0.002391104
> scale
0.005986117
> df
6.413187
> ncp
-0.2852327
```

Now plot the density of the NCT for delta(n+1) and compare it with the previously fitted Normal.

```
> setaxes(-0.05,0.05,0,80)
> lines(xv,dnorm(xv,mu,sigma1))
> abline(v=0,lwd=0.5)
> lines(xv,1/scale*dt((xv-location)/scale,df=df,ncp=ncp),col=2)
```

You can see a clear difference between the two. The NCT is more narrow peaked, it is slightly negatively skewed and it has a fatter left hand tail than the normal.

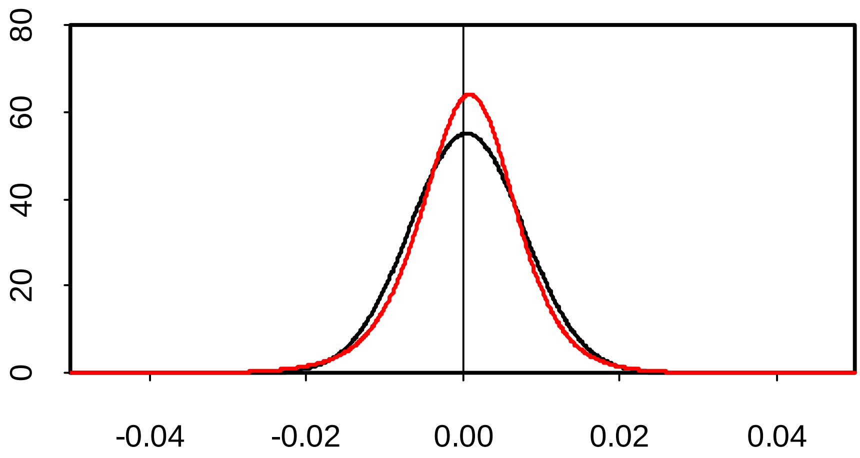

Exercise 2

```
# Now overlay res1$sigma, res2$sigma, res3$sigma using different colours
# Are they quite similar? Where do the differences occur?

> setaxes(1980,2006,-0.08,0.08)
> points(tv,d,pch=20,cex=0.6)
> abline(h=0,col=rgb(1,1,1),lwd=4)
> abline(h=0,lwd=1)
> lines(tv,res1$sigma,col=2,lwd=1)
> lines(tv,res2$sigma,col=3,lwd=1)
> lines(tv,res3$sigma,col=4,lwd=1)
```

In the graph below we can barely detect any difference between the blue (on top) red and green (underneath) lines for sigma(t) under the three distributional assumptions for Z(t).

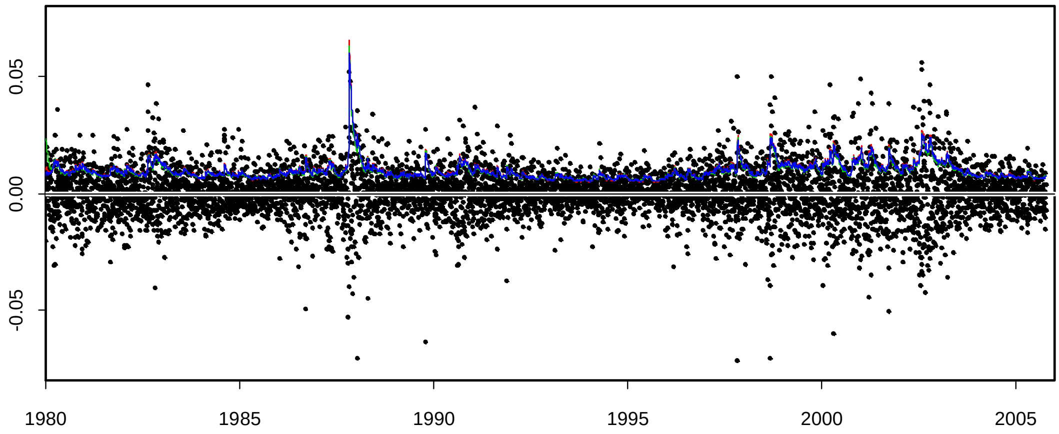

Also try:

```
plot(res1$sigma,res2$sigma,pch=20,cex=0.5)
abline(0,1,col=2)
plot(res1$sigma,res3$sigma,pch=20,cex=0.5)
abline(0,1,col=2)
plot(res2$sigma,res3$sigma,pch=20,cex=0.5)
abline(0,1,col=2)
```

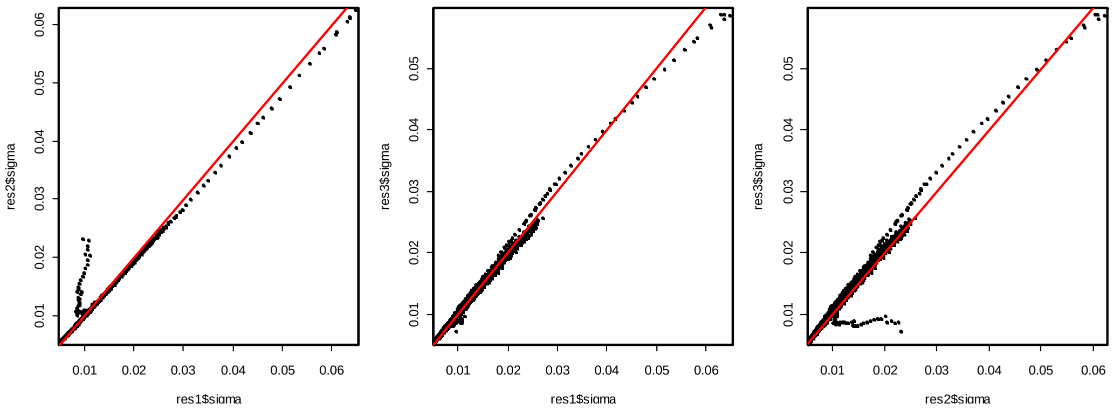

These plots reveal some small differences between the estimated sigma(t) for each time t, but essentially these confirm that the values are very close given the close alignment of the plots along the x=y diagonal . The exception is the sigma2(t) has some outliers, but these are linked to the estimated value of sigma(1) and this is very unimportant for forecasting after time n.

########## Exercise 2B ################

Now compare the Z(t)'s

```
cdfplot(res1$Z)
cdfplot(res2$Z,a=1,co=2)
cdfplot(res3$Z,a=1,co=3)
```

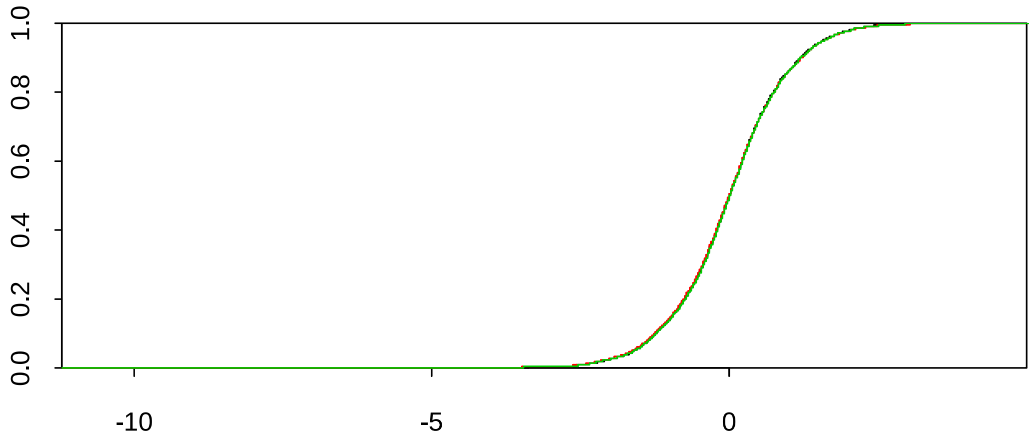

This plot supports the idea that the assumption of normal, t or NCT makes very little difference to the estimated Z(t)’s.

```
############ Exercise 2C #################

# Now compare the Z(t)'s by individual date

plot(res1$Z,res2$Z,pch=20,cex=0.5)
abline(0,1)

plot(res1$Z,res3$Z,pch=20,cex=0.5)
abline(0,1)

plot(res2$Z,res3$Z,pch=20,cex=0.5)
abline(0,1)
```

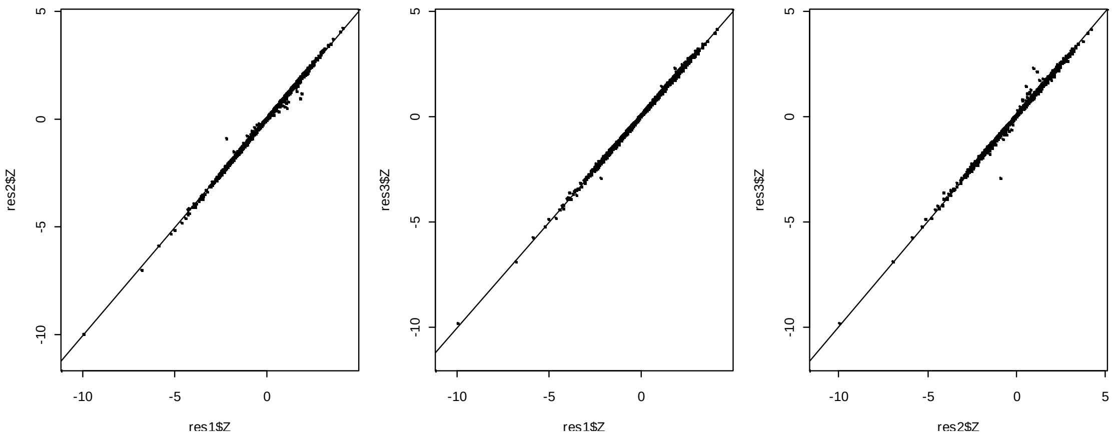

What do you conclude?

These plots reiterate the idea that the assumption of normal, t or NCT makes very little difference to the estimated Z(t)’s.

####### Exercise 4 ######

e.g. the NCT: Use the parameter values recorded in Exercise 1.

```
> D.n.plus.1=location+scale*rt(1000,df=df,ncp=ncp)
> setaxes(-0.05,0.05,0,1)
> lines(xv,pt((xv-location)/scale,df=df,ncp=ncp),col=2)
> cdfplot(D.n.plus.1,a=1)
```

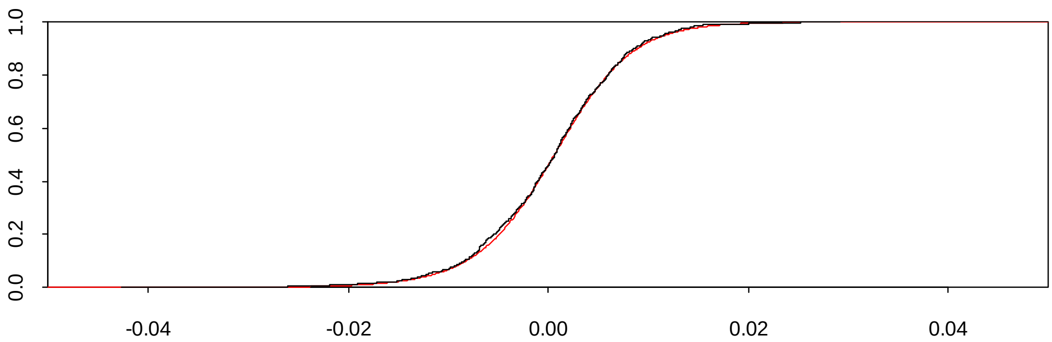

The plot superimposes the empirical CDF (wiggly black line) on to of the theoretical CDF (smooth red line). We can see that we get a good match as we would expect.