ORIE 678 — D. Ruppert
R2WinBUGS example: linear regression with subject-specific intercept: from hw #5

R Program:

```
library(R2WinBUGS)
data = read.csv("hw5_data.csv",header=TRUE)
ID = data[,1]
time = data[,2]
conc = data[,3]
N = length(ID)
data=list("ID","time","conc","N")
inits=function(){list(alpha_sub=rnorm(5,mean=0,sd=10),beta=rnorm(1,mean=0,sd=10),
alpha=rnorm(1,mean=0,sd=10),taualpha=rgamma(1,1,.1),taue=rgamma(1,1,.0))}

rand_int = bugs(data,inits,model.file="rand_int.bug",
parameters=c("alpha_sub","beta","alpha","taualpha","taue"),n.chains = 3,
n.iter=3100,n.burnin=100,n.thin=1,
bugs.directory="c:/Program Files/WinBUGS14/",codaPkg=FALSE)

pdf(file='plot1.pdf',height=9)
par(mfrow=c(5,2))
ts.plot(rand_int$sims.array[30*(1:100),,1])
ts.plot(rand_int$sims.array[30*(1:100),,2])
ts.plot(rand_int$sims.array[30*(1:100),,3])
ts.plot(rand_int$sims.array[30*(1:100),,4])
ts.plot(rand_int$sims.array[30*(1:100),,5])
ts.plot(rand_int$sims.array[30*(1:100),,6])
ts.plot(rand_int$sims.array[30*(1:100),,7])
ts.plot(rand_int$sims.array[30*(1:100),,8])
ts.plot(rand_int$sims.array[30*(1:100),,9])
ts.plot(rand_int$sims.array[30*(1:100),,10])

pdf(file='plot2.pdf',height=9)
par(mfrow=c(5,2))
acf(rand_int$sims.array[,1,1])
acf(rand_int$sims.array[,1,2])
acf(rand_int$sims.array[,1,3])
acf(rand_int$sims.array[,1,4])
acf(rand_int$sims.array[,1,5])
acf(rand_int$sims.array[,1,6])
acf(rand_int$sims.array[,1,7])
acf(rand_int$sims.array[,1,8])
acf(rand_int$sims.array[,1,9])
acf(rand_int$sims.array[,1,10])

graphics.off()
```

Bugs program:

```
model{
for(i in 1:N){
conc[i] ~ dnorm(muy[i],taue)
muy[i] <- alpha_sub[ID[i]] + beta*time[i]
}
for(i in 1:5)
{
alpha_sub[i] ~ dnorm(alpha,taualpha)
}
beta ~ dnorm(0.0,1.0E-6)
alpha ~ dnorm(0.0,1.0E-6)
taualpha ~dgamma(0.1,0.01)
taue ~dgamma(0.1,0.01)
}
```

Output:

```
> print(rand_int)
Inference for Bugs model at "rand_int.bug", fit using winbugs,
 3 chains, each with 3100 iterations (first 100 discarded)
 n.sims = 9000 iterations saved
             mean  sd 2.5%  25%  50%  75% 97.5% Rhat n.eff
alpha_sub[1]  3.4 0.1  3.1  3.3  3.4  3.4   3.6    1   520
alpha_sub[2]  0.5 0.1  0.3  0.4  0.5  0.6   0.8    1   480
alpha_sub[3] -0.5 0.1 -0.8 -0.6 -0.5 -0.4  -0.3    1   470
alpha_sub[4] -0.2 0.1 -0.4 -0.3 -0.2 -0.1   0.0    1   360
alpha_sub[5] -0.3 0.1 -0.5 -0.4 -0.3 -0.2   0.0    1   470
beta         -0.5 0.0 -0.5 -0.5 -0.5 -0.5  -0.5    1   290
alpha         0.6 0.9 -1.3  0.1  0.6  1.1   2.5    1  5400
taualpha      0.4 0.3  0.1  0.2  0.4  0.6   1.2    1  1400
taue         17.3 4.3 10.0 14.3 17.0 20.0  26.7    1  9000
deviance      0.4 4.3 -5.5 -2.7 -0.4  2.6  11.0    1  5100

For each parameter, n.eff is a crude measure of effective sample size,
and Rhat is the potential scale reduction factor (at convergence, Rhat=1).

pD = 7.2 and DIC = 7.6 (using the rule, pD = Dbar-Dhat)
DIC is an estimate of expected predictive error (lower deviance is better).
```

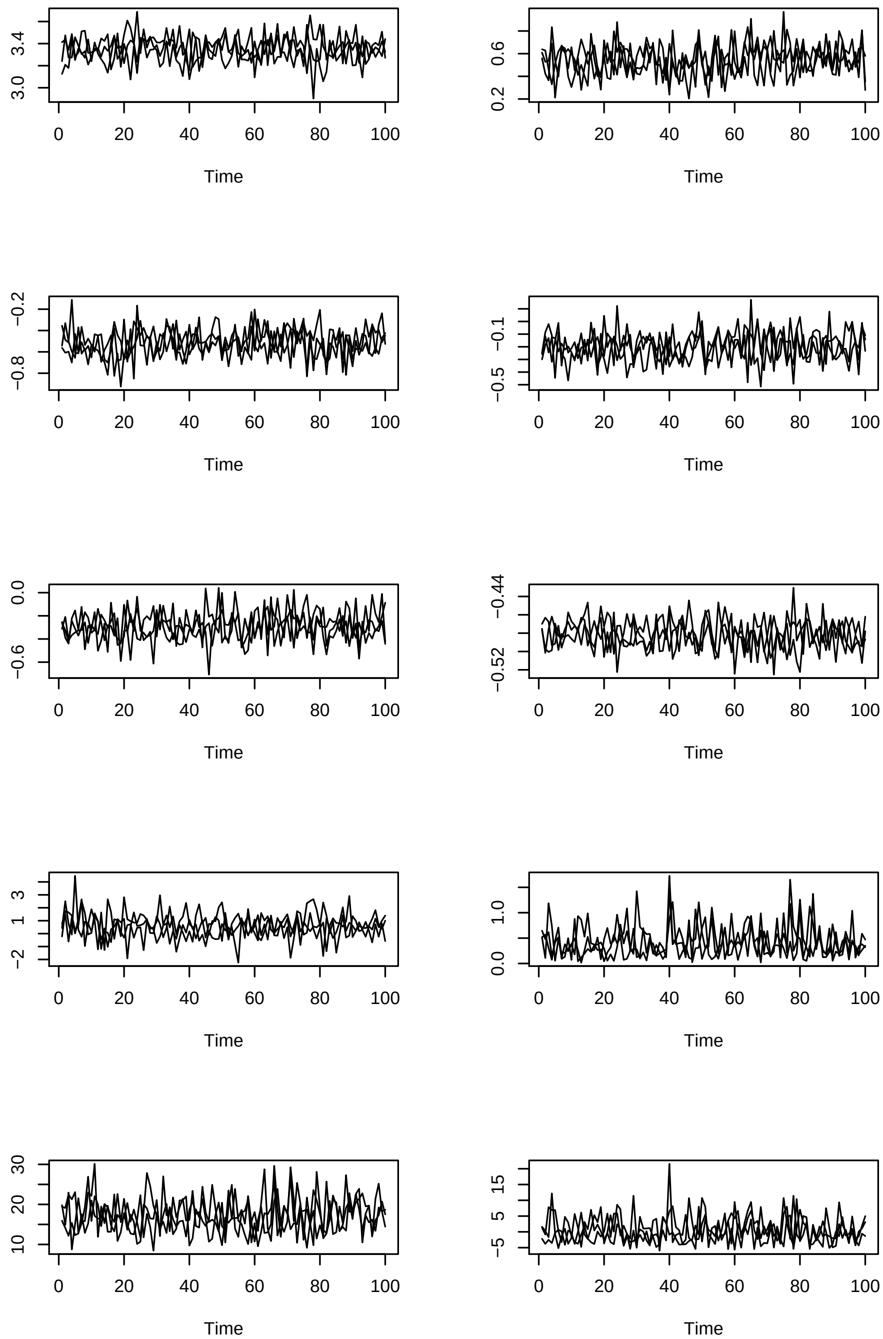

Figure 1: Trace plots.

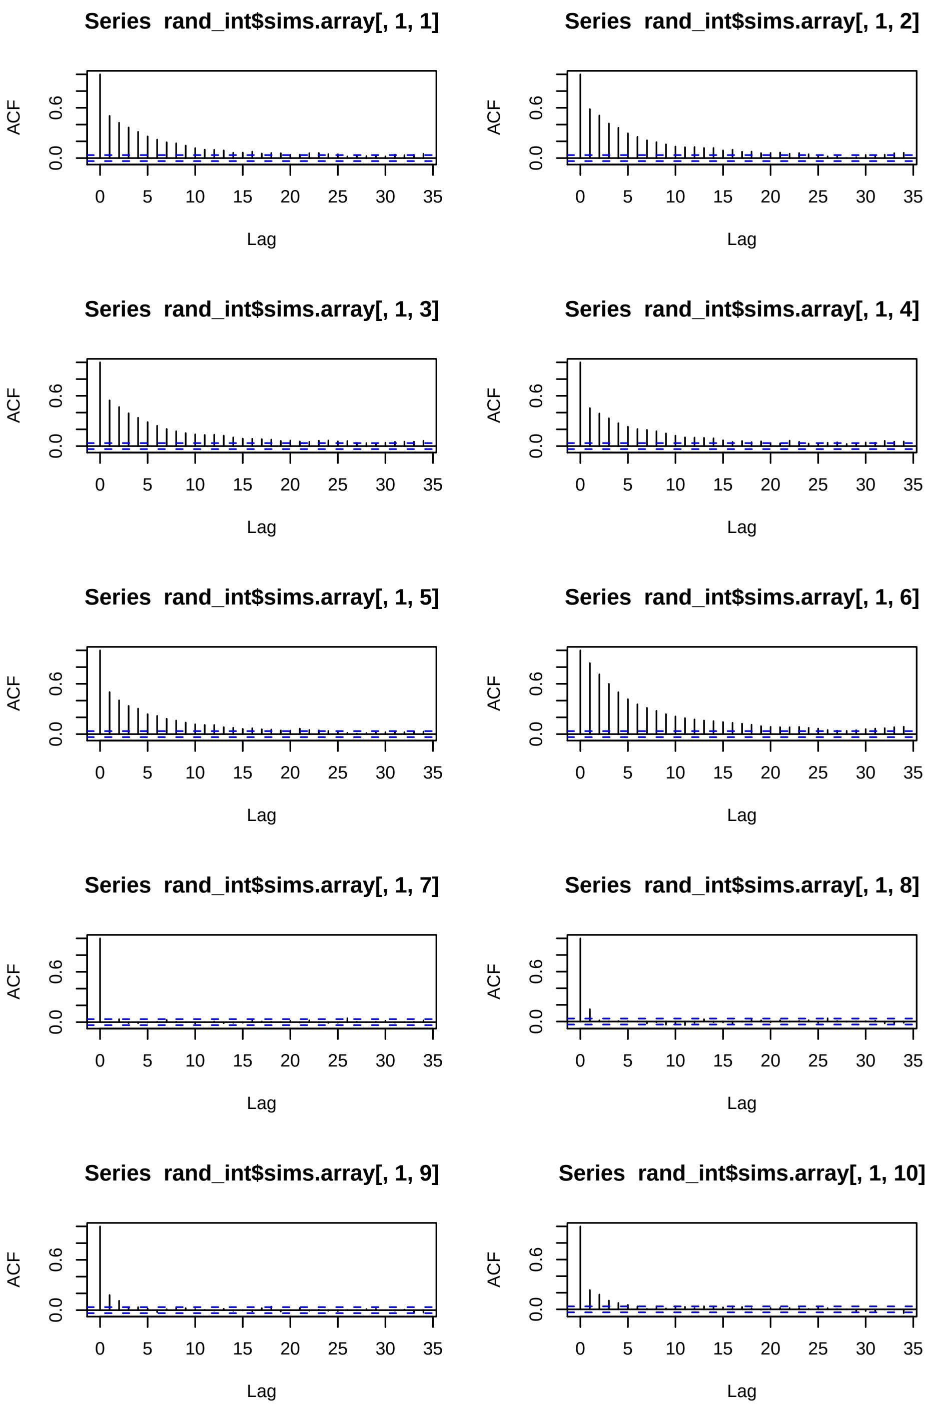

Figure 2: ACF plots.