

UNIVERSITEIT VAN DIE VRYSTAAT
UNIVERSITY OF THE FREE STATE

STSB 6815

WISKUNDIGE STATISTIEK & AKTUARIËLE WETENSKAP/
MATHEMATICAL STATISTICS & ACTUARIAL SCIENCE

Test 3 — 8 May 2015

MEMORANDUM

TYD/TIME: 270 Minutes

PUNTE/MARKS: 50

INSTRUCTIONS:

- Answer all questions in a single document and submit as a .pdf file.
- Label questions clearly, as it is done on this question paper.
- All calculations accurate to 4 decimal places.
- Show all formulas, code and reasoning.
- Use a significance level of 5% for all hypothesis tests.
- You may make use of any software readily available as well as pre-existing knowledge from the internet.
- You may NOT make any attempt to communicate or share information with fellow students.

Question 1

In a classical experiment carried out from 1918 to 1934, apple trees of different rootstocks were compared (Andrews and Herzberg, 1985, pp. 357–360). The data for eight trees from each of six rootstocks are available (see ‘rootstocks.csv’ on eLearn). The variables are

y_1 = trunk girth at 4 years (mm \times 100),

y_2 = extension growth at 4 years (m),

y_3 = trunk girth at 15 years (mm \times 100),

y_4 = weight of tree above ground at 15 years (lb \times 1000).

- (a) **MODEL A:** Consider only variable y_4 . Assume each observation of this variable follows a *Gamma* distribution with unknown parameters α and λ . Assign each parameter a vague *Gamma*(0.001, 0.001) prior.


- i. Simulate the joint posterior of the two parameters. At least 20,000 simulations in total. [7]

```
alldata <- read.csv('rootstock.csv')
y <- alldata$y4
n <- length(y)
library(R2openBUGS)
rootmodel <- function() {
  for (i in 1:n) {
    y[i] ~ dgamma(alpha, lambda)
  }
  alpha ~ dgamma(0.001, 0.001)
  lambda ~ dgamma(0.001, 0.001)
  mu <- alpha/lambda
}
write.model(rootmodel, 'BUGStemp.txt')
BUGSdata <- list(y=y, n=n)
inits <- function() {return(list(alpha=(mean(y)^2/var(y)), lambda=(mean(y)/var(y))))}
fittedrootsmodel <- bugs(BUGSdata, inits, 'BUGStemp.txt', parameters = c('alpha', 'lambda', 'mu'), n.chains=2, n.iter=22000, n.burnin=2000, n.thin=2, debug=TRUE)
```

Code ✓✓✓✓✓✓✓✓✓.

- ii. Produce a neat scatterplot of the joint distribution. Based on your inspection of the plot, explain the dependence structure (or lack thereof) between the two parameters. [4]


```
plot(fittedrootsmodel$sims.list$alpha, fittedrootsmodel$sims.list$lambda, pch='.', main='Scatter plot of joint posterior 1ai', xlab='alpha', ylab='lambda')
```



Code ✓✓, ✓, observations ✓.

- iii. Draw a histogram of the posterior distribution of the mean of y_4 ($\mu = \frac{\alpha}{\lambda}$) and compare it to the observed global mean of y_4 (\bar{y}_4). [4]

```
ybar <- mean(y)
hist(fittedrootsmodel$sims.list$mu, main='Histogram of mu', xlab='mu')
lines(c(ybar, ybar), c(0, 7000), col='blue', lwd=3)
```



Code ✓✓, ✓, observations ✓.

(b) **MODEL B:** Again consider only variable y_4 . Assume each observation of this variable follows a *Gamma* distribution with unknown parameters α_k and λ_k . These parameters should be allowed to vary between rootstocks ($k = 1, \dots, 6$) but remain constant within each rootstock. Further assume vague *Gamma*(0.001, 0.001) priors for all 12 parameters.

- i. Simulate the joint posterior of the 12 parameters. At least 20,000 simulations in total. Your model must also calculate the ratios of the α 's and λ 's to get the posterior distributions of μ_1, \dots, μ_6 . Finally, calculate the successive differences $\mu_1 - \mu_2, \mu_2 - \mu_3, \dots, \mu_5 - \mu_6$. [12]

```

ymat <- matrix(y,8,6)
rootmodelB <- function() {
  for (j in 1:6) {
    for (i in 1:8) {
      ymat[i,j] ~ dgamma(alpha[j],lambda[j])
    }
    alpha[j] ~ dgamma(0.001,0.001)
    lambda[j] ~ dgamma(0.001,0.001)
    mu[j] <- alpha[j]/lambda[j]
  }
  for (k in 1:5) {
    diffs[k] <- mu[k+1]-mu[k]
  }
}
write.model(rootmodelB,'BUGStemp.txt')
BUGSdata <- list(ymat=ymat)
inits <- function() {return(list(alpha=(colMeans(ymat)^2/apply(ymat,2,var)),lambda=(
  colMeans(ymat)/apply(ymat,2,var))))}
fittedrootmodelB <- bugs(BUGSdata, inits, 'BUGStemp.txt', parameters = c('alpha',
  'lambda','mu','diffs'), n.chains=2,n.iter=22000,n.burnin=2000,n.thin=2, debug=TRUE
)

```

- ii. Calculate 99% intervals of these differences ($\mu_1 - \mu_2, \mu_2 - \mu_3, \dots, \mu_5 - \mu_6$), then state which intervals contain 0 and what this implies. [4]

```
(intervals <- apply(fittedrootmodelB$sims.list$diffs,2,quantile,c(0.005,0.995)))
```

Code ✓✓,

Diffs	[,1]	[,2]	[,3]	[,4]	[,5]
0.5%	0.087	-0.298	-0.748	-0.233	-0.874
99.5%	0.777	0.512	0.007	0.583	-0.117

The first and last intervals do not contain 0 ✓, this implies that there are significant differences between some group means ✓.

(c) **MODEL COMPARISON**

- i. Obtain the value of DIC for both models. Compare these DIC values and judge which model is better, then explain what this implies, with respect to the properties of the weight of the trees of different rootstocks, in plain language.

[6]

```
fittedrootsmodel$DIC  
fittedrootsmodelB$DIC
```

Values are 22.98 ✓ and 0.8215 ✓, since the second model has a much smaller DIC it is the superior model ✓✓, the implication of this is that the data supports the notion that the groups have different means, even accounting for the extra complexity this brings ✓✓.

- ii. For up to 3 bonus marks, compare your judgement above with the results of a classical ANOVA on y_4 .

```
(anova(lm(alldata$y4~alldata$Rootstock)))
```

ANOVA table	Df	SumSq	MeanSq	F	Pr(>F)
alldata\$Rootstock	1	0.203	0.203	2.323	0.134
Residuals	46	4.013	0.087		

Code ✓, the classical ANOVA, with its assumption of Normality, fails to reject the null hypothesis that all means are equal ✓, it may be that the different assumption yields a different conclusion, or perhaps the ANOVA is more conservative. Additional testing would be required to determine the reason for the different results (any 1 of these 3 ✓).

Total for Question 1: 37

Question 2

A set of lecturers assessed the presentations of a set of students (see ‘marks2.csv’ on eLearn). Not all lecturers viewed all presentations, but more importantly, lecturers are not perfect assessors — they have different biases and variances. The goal is to obtain a fair mark for every student, as well as a ranking of individual assessors. Please note that the marks have already been transformed to the logistic scale $(-\infty; \infty)$.

Consider the following model:

- Observed marks are assumed Normal with mean $\mu_{ij} = s_i + l_j$ and precision τ_j .
- s_i , $i = 1, \dots, n = 15$ refers to the ‘true’ mark that each student deserves.
- l_j , $j = 1, \dots, k = 9$ refers to the bias of each lecturer. Note that $l_9 = -\sum_{j=1}^8 l_j$ because we assume that, on average, the lecturers are unbiased.
- A priori, assume $s_i \sim N(1.15, 0.01)$, $l_j \sim N(0, 0.04)$ and $\tau_j \sim \text{Exp}(0.01)$.

(a) Implement the above model by simulating from the joint posterior of all parameters. [9]

```

y <- read.csv('marks2.csv', row.names=1)
n <- nrow(y)
k <- ncol(y)
library(R2OpenBUGS)
marksmodel <- function() {
  for (i in 1:n) {
    for (j in 1:k) {
      y[i,j] ~ dnorm(mu[i,j], taou[j])
      mu[i,j] <- student[i] + lecbias[j]
      er[i,j] <- y[i,j]-student[i]
    }
  }
  student[i] ~ dnorm(1.15, 0.01)
}
for (j in 1:k) {
  taou[j] ~ dexp(0.01)
}
for (j in 1:(k-1)) {
  lecbias[j] ~ dnorm(0, 0.04)
}
lecbias[k] <- 0-sum(lecbias[1:(k-1)])
}
write.model(marksmodel, 'BUGStemp.txt')
BUGSdata <- list(y=as.matrix(y), n=n, k=k)
inits <- function() {return(list(student=rep(1.15, n), lecbias=rep(0, k), taou=rep(2.5, k),
mu=matrix(1.15, n, k)))}
fittedmarksmodel <- bugs(BUGSdata, inits, 'BUGStemp.txt', parameters = c('student', '
lecbias', 'er'), n.chains=2, n.iter=22000, n.burnin=2000, n.thin=2, debug=TRUE)

```

Code 9 marks (note that the `er[i,j]` line is not worth marks here as it is only necessary for Part d, which is bonus marks).

(b) Determine the final mark to be awarded to each student by averaging over the simulations (posterior mean) and then transforming back to the marks scale, using the inverse logistic transform. [2]

```

attach.bugs(fittedmarksmodel)
finalmarks <- round(plogis(colMeans(student))*100, 1)
names(finalmarks) <- rownames(y)
finalmarks

```

Average ✓, inverse logit ✓.

Frikkie&Christelle	79.3	Leon&Carina	86.8	Morapeli	76
Daryl	78.1	Mulanga	68.6	Lucky&Oshomah	63.4
Mariska	75.3	Hilda	63.1	Caldwell	81.5
Seun	63.4	Napo	69.8	Teboho	70.4
Janca&Ehan	84.5	Gao	83.4	Sabelo&Mpendulo	76.9

(c) Sort the lecturers from lowest square average bias to largest square average bias. [2]

```
lecturers <- colnames(y)
lecturers[order(colMeans(lecbias)^2)]
```

Code ✓, “Max” “Michael” “Sean” “Delson” “Morné” “Zani” “Dawie” “Andréhette”
“Martin” ✓

(d) For 2 bonus marks, sort the lecturers from lowest mean square error to largest mean square error, where error refers to the difference between the observed mark and the mark to be awarded to a student.

```
meanerrs <- apply(er^2,3,mean)
lecturers[order(meanerrs)]
```

Code ✓, “Max” “Sean” “Morné” “Dawie” “Andréhette” “Delson” “Michael” “Martin” “Zani” ✓.

Total for Question 2: 13

Question:	1	2	Total
Points:	37	13	50