

MAS8303 Modern Bayesian Inference

Part 2

Solutions to Problems 3

M. Farrow
University of Newcastle upon Tyne

Semester 1, 2012-3

1. Prior Elicitation

Some household contents insurance policies require an estimate to be made of what it would cost to replace the existing contents. Suppose that a person has a large collection of books. We might attempt to predict the replacement cost of all of the books by looking at a sample. We might improve this prediction by taking into account an auxiliary variable such as the width of the spine of the book. (We might also distinguish between hardback and paperback books so suppose that we are only considering hardback books). Let C_i be the replacement cost, in £, of book i , and let w_i be its spine width in mm.

Let

$$Y_i = \log_e(C_i) \quad \text{and} \quad x_i = \log_e(w_i).$$

It is believed that Y is related to X by

$$Y_i = \alpha + \beta x_i + \varepsilon_i$$

where Y_i and x_i refer to book i for $i = 1, \dots, n$, $\varepsilon_i \sim N(0, \tau^{-1})$ and $\varepsilon_1, \dots, \varepsilon_n$ are conditionally independent given τ .

We give α and β a bivariate normal prior distribution. Find the parameters of this distribution based on the following prior judgments.

Suppose that we could observe a very large number of books, each of which has a spine $w = 20$ mm wide, and a very large number of books, each of which has a spine $w = 30$ mm wide. Let the median replacement costs at these two spine widths be M_{20} and M_{30} respectively. Our prior median for M_{20} is 25 and our prior median for M_{30} is 35. Our prior upper quartile for M_{20} is 40 and our prior upper quartile for M_{30} is 55.

Let

$$m_{20} = \log_e(M_{20}) \quad \text{and} \quad m_{30} = \log_e(M_{30}).$$

Our prior correlation for m_{20} and m_{30} is 0.75.

We wish to find the prior means, prior variances and prior covariance of α , β .

Solution

The mean for m_{20} is $\mu_{20} = \log(25) = 3.218876$.

The upper quartile for m_{20} is $\mu_{20} + 0.6744898\sigma_{20} = \log(40) = 3.688879$. (Note that the value 0.6744898 can be obtained using the R command `qnorm(0.75, 0, 1)`).

Hence the standard deviation for m_{20} is

$$\sigma_{20} = \frac{\log(40) - \log(25)}{0.6744898} = 0.6968284.$$

Thus the variance for m_{20} is $\sigma_{20}^2 = 0.48556978$.

The mean for m_{30} is $\mu_{30} = \log(35) = 3.555348$.

The upper quartile for m_{30} is $\mu_{30} + 0.6744898\sigma_{30} = \log(55) = 4.007333$.

Hence the standard deviation for m_{30} is

$$\sigma_{30} = \frac{\log(55) - \log(35)}{0.6744898} = 0.6701141.$$

Thus the variance for m_{30} is $\sigma_{30}^2 = 0.4490529$.

The correlation of m_{20} and m_{30} is $\text{Covar}(m_{20}, m_{30})/(\sigma_{20}\sigma_{30})$ so the covariance is

$$\text{Covar}(m_{20}, m_{30}) = \text{Corr}(m_{20}, m_{30})\sigma_{20}\sigma_{30} = 0.3502159.$$

Hence the covariance matrix of $(m_{20}, m_{30})^T$ is

$$V_m = \text{Var} \begin{pmatrix} m_{20} \\ m_{30} \end{pmatrix} = \begin{pmatrix} 0.48556978 & 0.3502159 \\ 0.3502159 & 0.4490529 \end{pmatrix}.$$

Now

$$\begin{aligned} m_{20} &= \alpha + \log(20)\beta \\ \text{and } m_{30} &= \alpha + \log(30)\beta. \end{aligned}$$

so $m_{30} - m_{20} = \log(30/20)\beta$ and

$$\beta = \frac{m_{30}}{\log(1.5)} - \frac{m_{20}}{\log(1.5)}.$$

Substituting, we obtain

$$m_{20} = \alpha + \frac{\log(20)}{\log(1.5)}m_{30} - \frac{\log(20)}{\log(1.5)}m_{20}$$

and

$$\begin{aligned} \alpha &= \left(\frac{\log(20)}{\log(1.5)} + 1 \right) m_{20} - \frac{\log(20)}{\log(1.5)} m_{30} = \left(\frac{\log(20) + \log(1.5)}{\log(1.5)} \right) m_{20} - \frac{\log(20)}{\log(1.5)} m_{30} \\ &= \frac{\log(30)}{\log(1.5)} m_{20} - \frac{\log(20)}{\log(1.5)} m_{30} \end{aligned}$$

Hence

$$\begin{pmatrix} \alpha \\ \beta \end{pmatrix} = H \begin{pmatrix} m_{20} \\ m_{30} \end{pmatrix}$$

where

$$H = \frac{1}{\log(1.5)} \begin{pmatrix} \log(30) & -\log(20) \\ -1 & 1 \end{pmatrix}.$$

So the α and β have a bivariate normal prior distribution with means given by

$$E \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = H \begin{pmatrix} \mu_{20} \\ \mu_{30} \end{pmatrix} = \begin{pmatrix} 0.73289 \\ 0.82984 \end{pmatrix}$$

and variances and covariances given by

$$\text{Var} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = H V_m H^T = \begin{pmatrix} 15.2647 & -4.6013 \\ -4.6013 & 1.4245 \end{pmatrix}.$$

2. Lowering blood pressure during surgery

It is sometimes necessary to lower a patient's blood pressure during surgery, using a hypotensive drug. The length of time over which the drug is administered varies and therefore so does the total dose. This, in turn, might affect the time it takes for the patient's blood pressure to return to normal.

The data provided are as follows, for $n = 53$ patients.

- The natural logarithm of the recovery time, T , in minutes.
- The natural logarithm of the dose, d , in milligrams.
- The average systolic blood pressure, b , in millimetres of mercury, during administration.

Let $Y = \ln(T)$, $x_1 = \ln(d) - 5$ and $x_2 = b - 60$. We will use a regression model with

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \varepsilon$$

where $\beta_0, \beta_1, \beta_2$ are unknown parameters and, conditional on the values of the parameters, $\varepsilon_1, \dots, \varepsilon_{53}$ are independent with $\varepsilon_i \sim N(0, \tau^{-1})$.

Our prior distribution is as follows.

We give τ a gamma prior, $\tau \sim \text{Ga}(1.5, 0.6)$. Conditional on τ we give $\underline{\beta} = (\beta_0, \beta_1, \beta_2)^T$ a multivariate normal prior distribution with mean vector $\underline{b}_0 = (3.0, -0.03, 0.5)^T$ and precision matrix τC_0 where $C_0 = (V_0/v_0)^{-1}$ and

$$V_0 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 10^{-4} & 0 \\ 0 & 0 & 0.04 \end{pmatrix}.$$

Solution

- (a) The prior is conjugate so we can compute the parameters of the posterior distribution using the R function `linmod`. The commands used in the computation are as follows. Data set 30 is used here.

```
> surgery<-read.table("http://www.mas.ncl.ac.uk/~nmf16/teaching/mas8303/surgerydata.txt")
> y<-surgery[,30]
> x1<-surgery[,1]-5
> x2<-surgery[,2]-60
> x0<-rep(1,53)
> X<-cbind(x0,x1,x2)
> V0<-diag(c(1,0.0001,0.04))
> V0
      [,1] [,2] [,3]
[1,]  1 0e+00 0.00
[2,]  0 1e-04 0.00
[3,]  0 0e+00 0.04
> b0<-matrix(c(3,-0.03,0.5),ncol=1)
> d0<-3
> v0<-0.4
> surgeryprior<-list(d=d0,v=v0,b=b0,V=V0)
> surgerypost<-linmod(surgeryprior,X,y)
```

This gives the following output.

```

> surgerypost
$d
[1] 56

$v
[1] 0.4611439

$b
      [,1]
x0 2.899024705
x1 -0.027134909
x2 -0.005778664

$V
      x0      x1      x2
x0 0.0145219367 8.064500e-05 -9.302471e-04
x1 0.0000806450 1.145669e-04 -5.356306e-06
x2 -0.0009302471 -5.356306e-06 1.474940e-04

```

Thus the marginal posterior distribution of τ is $\tau \sim \text{Ga}(28, 12.91203)$.

Conditional on τ , the posterior distribution for $\underline{\beta}$ is multivariate normal with mean vector $\underline{b}_1 = (2.89902, -0.027135, -0.0057787)^T$ and precision matrix τC_1 where $C_1 = (V_1/v_1)^{-1}$, $v_1 = 0.461144$ and

$$V_1 = \begin{pmatrix} 0.01452194 & 0.00008065 & -0.00093025 \\ 0.00008065 & 0.00011457 & -0.00000536 \\ -0.00093025 & -0.00000536 & 0.00014749 \end{pmatrix}.$$

We notice that the posterior mean for β_2 is rather smaller than expected in the prior.

- (b) The covariate vector for a patient with log dose 4.0 and blood pressure 70 during administration is $\underline{x}_0 = (1, -1, 10)$. We compute a predictive mean and conditional variance for the log recovery time as $\underline{x}_0 \underline{b}_1 = 2.868373$ and $v_1 + \underline{x}_0 V_1 \underline{x}_0^T = 0.4718707$. these are used in computing and plotting the predictive density using the following R commands.

```

> v1<-surgerypost$v
> V1<-surgerypost$V
> x0<-matrix(c(1,-1,10),nrow=1)
> b1<-surgerypost$b
> mean<-x0%*%b1
> mean
      [,1]
[1,] 2.86197
> var<-v1+x0%*%V1%*%t(x0)
> var
      [,1]
[1,] 0.4702986
> d1<-surgerypost$d
> tpred<-seq(-4,4,0.08)
> tdens<-dt(tpred,d1)
> ypred<-tpred*sqrt(var)+mean
> ydens<-tdens/sqrt(var)
> plot(ypred,ydens,type="l",xlab="Log recovery time",ylab="Density")
> abline(h=0)

```

The resulting plot is shown in Figure 1. Notice that we had to divide the density by the standard deviation because we multiplied out t variable by this.




Figure 1: Predictive density for the log recovery time (minutes) for a patient with log dose 4.0 and blood pressure 70 during administration.

- (c) The recovery time is obtained by exponentiating the log recovery time. As this is a nonlinear transformation we need to obtain the “Jacobian.” So $t = \exp(y)$ giving $dt/dy = \exp(y) = t$. Hence $f_t(t) = f_y(y)/(dt/dy) = f_y(y)/t$. We use the following R commands.

```
> Tpred<-exp(ypred)
> Tdens<-ydens/Tpred
> plot(Tpred,Tdens,type="l",xlab="Recovery time",ylab="Density",xlim=c(0,100))
> abline(h=0)
```

The resulting plot is shown in Figure 2. We see that we are not able to predict the recovery time very precisely.

3. Yields of barley

An experiment was conducted to investigate the effect of manure on the yield of barley. Four different levels of manure were compared: 1: no manure, 2: 0.01 tons per acre, 3: 0.02 tons per acre, 4: 0.04 tons per acre. Three different varieties of barley were used. The experimental plots were arranged in six blocks. (A “block” is an area of land).

Let the design matrix be X . The first four columns of X correspond to the four levels of manure. Columns 5-9 are for the block effects. (There are five degrees of freedom between the six blocks). Columns 10-11 are for the variety effects. (There are two degrees of freedom between the three varieties). We will not fit interaction effects.

Let the parameters corresponding to the eleven columns of X be $\beta_1, \dots, \beta_{11}$. Then the mean yield, $\mu_{m,b,v}$, for manure level m in block b with variety v is defined as follows.




Figure 2: Predictive density for the recovery time in minutes for a patient with log dose 4.0 and blood pressure 70 during administration.

$$\mu_{m,b,v} = \beta_m + \sum_{j=5}^9 \beta_j w_{b,j} + \sum_{j=10}^{11} \beta_j z_{v,j}$$

Here $w_{b,j}$ and $z_{v,j}$ are defined as follows.

$w_{b,j}$	$j = 5$	$j = 6$	$j = 7$	$j = 8$	$j = 9$
$b = 1$	1	2	0	0	0
$b = 2$	1	-1	1	0	0
$b = 3$	1	-1	-1	0	0
$b = 4$	-1	0	0	2	0
$b = 5$	-1	0	0	-1	1
$b = 6$	-1	0	0	-1	-1

$z_{v,j}$	$j = 10$	$j = 11$
$v = 1$	2	0
$v = 2$	-1	1
$v = 3$	-1	-1

The actual yield for for manure level m in block b with variety v is

$$y_{m,b,v} = \mu_{m,b,v} + \varepsilon_{m,b,v}$$

where $\varepsilon_{m,b,v} \sim N(0, \tau^{-1})$ and $\varepsilon_{m,b,v}$ is independent of $\varepsilon_{m',b',v'}$ unless $(m, b, v) = (m', b', v')$.

Our prior distribution is as follows.

We give τ a gamma prior, $\tau \sim \text{Ga}(d_0/2, d_0 v_0/2)$ with $d_0 = 2.1$ and $v_0 = 250$. Conditional on τ we give $\underline{\beta} = (\beta_0, \dots, \beta_{11})^T$ a multivariate normal prior distribution with mean vector

$$\underline{b}_0 = 100(1, 1, 1, 1, 0, 0, 0, 0, 0, 0, 0)^T$$

and precision matrix τC_0 where $C_0 = (V_0/v_0)^{-1}$ and

$$V_0 = \frac{1}{6} \begin{pmatrix} 24 & 12 & 12 & 12 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 12 & 24 & 12 & 12 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 12 & 12 & 24 & 12 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 12 & 12 & 12 & 24 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 6 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 6 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 3 \end{pmatrix}.$$

- (a) We compute the parameters of the posterior distribution using the R function `linmod`. First we construct the design matrix X using the following R commands.

```
> X<-matrix(nrow=72,ncol=11)
> for (col in 1:4)
+   {X[,col]<-ifelse(manure==col,1,0)
+   }
> b<-rep(12,6)
> X[,5]<-rep(c(1, 1, 1,-1,-1,-1),b)
> X[,6]<-rep(c(2,-1,-1, 0, 0, 0),b)
> X[,7]<-rep(c(0, 1,-1, 0, 0, 0),b)
> X[,8]<-rep(c(0, 0, 0, 2,-1,-1),b)
> X[,9]<-rep(c(0, 0, 0, 0, 1,-1),b)
> v<-rep(4,3)
> x<-rep(c(2,-1,-1),v)
> X[,10]<-rep(x,6)
> x<-rep(c(0, 1,-1),v)
> X[,11]<-rep(x,6)
```

The other prior parameters are defined and the posterior parameters are calculated using the following R commands.

```
> d0<-2.1
> v0<-250
> b0<-100*matrix(c(1,1,1,1,0,0,0,0,0,0,0),ncol=1)
> V0<-matrix(0,nrow=11,ncol=11)
> V0[1:4,1:4]<-matrix(2,nrow=4,ncol=4)+diag(2,4)
> V0[5:11,5:11]<-diag(c(2,2,6,2,6,1,3))/6
> barleyprior<-list(d=d0,v=v0,b=b0,V=V0)
> barleypost<-linmod(barleyprior,X,y)
```

The output is as follows.

```
> barleypost
$d
[1] 74.1

$v
[1] 657.3465
```

\$b

```

      [,1]
[1,] 98.04358432
[2,] 101.71491299
[3,] 103.68694097
[4,] 104.76386404
[5,]  0.58759124
[6,]  0.98783455
[7,]  0.35766423
[8,]  0.01703163
[9,] -0.30656934
[10,] -0.19282238
[11,] -0.24270073

```

\$V

```

      [,1] [,2] [,3] [,4] [,5] [,6] [,7]
[1,] 7.269403 2.672575 2.672575 2.672575 0.0000000 0.0000000 0.0000000
[2,] 2.672575 7.269403 2.672575 2.672575 0.0000000 0.0000000 0.0000000
[3,] 2.672575 2.672575 7.269403 2.672575 0.0000000 0.0000000 0.0000000
[4,] 2.672575 2.672575 2.672575 7.269403 0.0000000 0.0000000 0.0000000
[5,] 0.0000000 0.0000000 0.0000000 0.0000000 0.7996916 0.0000000 0.0000000
[6,] 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000 0.7996916 0.0000000
[7,] 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000 2.399075
[8,] 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000
[9,] 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000
[10,] 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000
[11,] 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000
      [,8] [,9] [,10] [,11]
[1,] 0.0000000 0.0000000 0.0000000 0.0000000
[2,] 0.0000000 0.0000000 0.0000000 0.0000000
[3,] 0.0000000 0.0000000 0.0000000 0.0000000
[4,] 0.0000000 0.0000000 0.0000000 0.0000000
[5,] 0.0000000 0.0000000 0.0000000 0.0000000
[6,] 0.0000000 0.0000000 0.0000000 0.0000000
[7,] 0.0000000 0.0000000 0.0000000 0.0000000
[8,] 0.7996916 0.0000000 0.0000000 0.0000000
[9,] 0.0000000 2.399075 0.0000000 0.0000000
[10,] 0.0000000 0.0000000 0.3998458 0.0000000
[11,] 0.0000000 0.0000000 0.0000000 1.199537

```

Thus the posterior distribution is as follows. The marginal posterior distribution of τ is a gamma distribution, $\tau \sim \text{Ga}(d_1/2, d_1 v_1/2)$ with $d_1 = 74.1$ and $v_0 = 657.3465$. Conditional on τ , $\underline{\beta} = (\beta_0, \dots, \beta_{11})^T$ has a multivariate normal posterior distribution with mean vector \underline{b}_1 with elements as follows

i	1	2	3	4	
$\underline{b}_{1,i}$	98.04358	101.71491	103.68694	104.763864	
i	5	6	7	8	9
$\underline{b}_{1,i}$	0.58759	0.98783	0.35766423	0.01703163	-0.30656934
i	10	11			
$\underline{b}_{1,i}$	-0.19282	-0.24270			

and precision matrix τC_1 where $C_1 = (V_1/v_1)^{-1}$ and V_1 has elements $v_{1,i,j}$ as follows. For $i = 1, \dots, 4$, $j = 1, \dots, 4$, we have:

	$j = 1$	$j = 2$	$j = 3$	$j = 4$
$i = 1$	7.26940	2.67258	2.67258	2.67258
	2.67258	7.26940	2.67258	2.67258
	2.67258	2.67258	7.26940	2.67258
	2.67258	2.67258	2.67258	7.26940

For $i = 5, \dots, 11$, we have $v_{1,i,i}$ as follows:

i	5	6	7	8	9	10	11
$v_{1,i,i}$	0.79969	0.79969	2.39908	0.79969	2.39908	0.39985	1.19954

Otherwise $v_{1,i,j} = 0$.

- (b) To find the mean yield for Manure level 1 in Block 1 with Variety 1 we introduce the covariate vector $\underline{x}_0 = (1, 0, 0, 0, 1, 2, 0, 0, 0, 2, 0)$. We compute the appropriate mean $\underline{x}_0 \underline{b}_1$ and conditional variance $\underline{x}_0 V_1 \underline{x}_0^T$ using the following R commands.

```
> x0<-matrix(c(1,0,0,0,1,2,0,0,0,2,0),nrow=1)
> v1<-barleypost$v
> d1<-barleypost$d
> V1<-barleypost$V
> b1<-barleypost$b
> mean<-x0%*%b1
> var<-x0%*%V1%*%t(x0)
```

The symmetric posterior 95% interval for the mean yield is then given by $\underline{x}_0 \underline{b}_1 \pm t_{d_1, 0.975} \sqrt{\underline{x}_0 V_1 \underline{x}_0^T}$ where $t_{d_1, 0.975}$ is the upper 97.5% point of the t -distribution on d_1 degrees of freedom. The R commands are as follow.

```
> tval<-qt(0.975,d1)
> mean-tval*sqrt(var)
      [,1]
[1,] 93.07392
> mean+tval*sqrt(var)
      [,1]
[1,] 107.3685
```

The 95% interval is therefore (93.07392, 107.3685).