MAS3301 Bayesian Statistics

M. Farrow School of Mathematics and Statistics Newcastle University

Semester 2, 2008-9

17 Large Samples: Examples

1. The numbers of telephone calls arriving at a busy exchange in 200 short intervals are observed. Suppose that these are X_1, \ldots, X_n , where n = 200, that $X_i \mid \lambda \sim \text{Poisson}(\lambda)$ and that X_i and X_j are independent, given λ , when $i \neq j$. The likelihood is therefore

$$L = \prod_{i=1}^{n} \frac{e^{-\lambda} \lambda^{x_i}}{x_i!} \propto e^{-n\lambda} \lambda^s$$

where $s = \sum_{i=1}^{n} x_i$.

Suppose that the prior for λ is a gamma(a, b) distribution. The posterior is then gamma(s + a, n+b). This has mean (s+a)/(n+b) and variance $(s+a)/(n+b)^2$ and we can approximate this with a normal distribution

$$N\left(\frac{s+a}{n+b}, \frac{s+a}{(n+b)^2}\right).$$

Suppose that a = 2, b = 0.2, s = 1723. Then the posterior distribution is approximately $N(8.6164, 0.2075^2)$ and an approximate 95% interval for λ is $8.6164 \pm 1.96 \times 0.2075$. This gives

$$8.210 < \lambda < 9.023.$$

Alternatively we can use the posterior mode to find the mean of the normal distribution. The posterior density is proportional to

$$\lambda^{s+a-1}e^{-(n+b)\lambda}.$$

The log of this is

$$g(\lambda) = (s + a - 1)\log(\lambda) - (n + b)\lambda$$

Differentiating wr
t λ we obtain

$$g'(\lambda) = \frac{s+a-1}{\lambda} - (n+b).$$

This gives the posterior mode at $\lambda_m = (s + a - 1)/(n + b) = 8.6114$. Differentiating again we obtain

$$g''(\lambda) = -\frac{s+a-1}{\lambda^2}.$$

This gives an approximate variance of

$$\frac{\lambda_m^2}{s+a-1} = \frac{s+a-1}{(n+b)^2} = 0.2075^2$$

So our 95% interval becomes $8.6114 \pm 1.96 \times 0.2075.$ This gives

 $8.205 < \lambda < 9.018.$

Calculating the hpd interval without using the normal approximation (using hpdgamma) gives

 $8.211 < \lambda < 9.024$

which is very close to the first of our approximations.

2. We observe the result of a sample survey in which respondents are asked a simple "Yes/No" question. The number answering "Yes" is x. The sample size n is large and we assume that the effect of the prior distribution is negligible. We also assume that x is an observation from a binomial (n, θ) distribution where

$$\theta = \frac{e^{\mu}}{1 + e^{\mu}}.$$

That is

$$\mu = \log\left(\frac{\theta}{1-\theta}\right).$$

We are interested in the parameter μ . Find an approximation to its posterior distribution.

The likelihood is

$$L(\mu) = \begin{pmatrix} n \\ x \end{pmatrix} \theta^x (1-\theta)^{n-x}.$$

The log likelihood is

$$l(\mu) = \text{constant} + x \log(\theta) + (n - x) \log(1 - \theta)$$

We need to differentiate this with respect to $\mu.$ We note that

$$\frac{dl}{d\mu} = \frac{dl}{d\theta} \frac{d\theta}{d\mu}$$

and that

$$\frac{d\theta}{d\mu} = \frac{(1+e^{\mu})e^{\mu} - e^{\mu}e^{\mu}}{(1+e^{\mu})^2} = \frac{e^{\mu}}{(1+e^{\mu})^2} = \theta(1-\theta).$$

 So

$$\begin{aligned} \frac{dl}{d\theta} &= \frac{x}{\theta} - \frac{n-x}{1-\theta} &= \frac{x-\theta x - n\theta + \theta x}{\theta(1-\theta)} \\ &= \frac{x-n\theta}{\theta(1-\theta)} \\ \frac{dl}{d\mu} &= \frac{dl}{d\theta} \frac{d\theta}{d\mu} &= x-n\theta \\ \frac{d}{d\theta} \left(\frac{dl}{d\mu}\right) &= -n \\ \frac{d^2l}{d\mu^2} &= \frac{d}{d\theta} \left(\frac{dl}{d\mu}\right) \frac{d\theta}{d\mu} &= -n\theta(1-\theta) \end{aligned}$$

Solving $dl/d\mu = 0$ we obtain

$$x = n\hat{\theta}$$
 and $\hat{\theta} = \frac{x}{n}$

and therefore

$$\hat{\mu} = \log\left(\frac{x/n}{1-x/n}\right) = \log\left(\frac{x}{n-x}\right).$$

Substituting this into the second derivative we obtain, at $\theta = \hat{\theta}$,

$$\frac{d^2l}{d\mu^2} = -n\hat{\theta}(1-\hat{\theta}) < 0.$$

Therefore $\hat{\mu}$ is a maximum and is approximately the posterior mean and the posterior variance is approximately

$$V = [n\hat{\theta}(1-\hat{\theta})]^{-1}.$$

The posterior distribution of μ is approximately $N(\hat{\mu}, V)$. Suppose that n = 2000 and x = 1456. Then $\hat{\theta} = \overline{x/n} = 0.728$ and

$$\hat{\mu} = \log\left(\frac{1456}{2000 - 1456}\right) = 0.9845.$$

The approximate variance is

$$V = [2000 \times 0.728(1 - 0.728)]^{-1} = [396.032]^{-1} = 0.002525$$

giving a standard deviation of 0.05025. An approximate 95% posterior hpd interval for μ is therefore $0.9845 \pm 1.96 \times 0.05025$. that is

$$0.8860 < \mu < 1.0830.$$

18 Predictive Distributions

18.1 Introduction

In Lecture 3 we introduced parameters as a means of structuring our beliefs and of transferring information from observations to beliefs about future observations. Since then we have spent some time looking at how information from observations affects our beliefs about parameters. Information goes from observations to our beliefs about the values of parameters via the likelihood. In this lecture we will return to thinking about our beliefs for the values of future observations.

Suppose that our beliefs about a (possibly vector) parameter θ are represented by the distribution with density function $f_{\theta}(\theta)$. (We assume here that θ is continuous. It could be discrete. In that case we replace the density function by a probability function and the integration by a summation). The distribution of a future observation Y given θ has density function, if Y is continuous, or probability function, if Y is discrete, $f_{Y|\theta}(y \mid \theta)$. For convenience we will assume for now that Y is continuous.

The joint density of θ and Y is therefore

$$f_{\theta}(\theta) f_{Y\mid\theta}(y\mid\theta).$$

To find the marginal density of Y we simply integrate out θ .

$$f_Y(y) = \int_{-\infty}^{\infty} f_{\theta}(\theta) f_{Y|\theta}(y \mid \theta) \ d\theta.$$

This marginal distribution is called a *predictive distribution* of Y. Its mean is the *predictive mean* and so on.

Very often, of course, we are interested in the case where we observe some data and then make predictions about future observations which we have not yet seen. In this case, before we have seen the data, we can evaluate the *prior predictive distribution*:

$$f_Y^{(0)}(y) = \int_{-\infty}^{\infty} f_{\theta}^{(0)}(\theta) f_{Y|\theta}(y \mid \theta) \ d\theta$$

where $f_{\theta}^{(0)}(\theta)$ is the prior density of θ . After we have seen data D, we can evaluate the *posterior* predictive distribution:

$$f_Y^{(1)}(y \mid D) = \int_{-\infty}^{\infty} f_{\theta \mid D}^{(1)}(\theta \mid D) f_{Y \mid \theta}(y \mid \theta) \ d\theta$$

where $f_{\theta|D}^{(1)}(\theta \mid D)$ is the posterior density of θ given the data D.

Predictive distributions reflect both the *aleatory* uncertainty in the future observations and the *epistemic* uncertainty in the parameters. All (remaining) uncertainty is properly reflected. This is an important difference from non-Bayesian statistics. The prior predictive distribution describes our beliefs before we have seen data and the posterior predictive distribution describes our beliefs afterwards.

Predictive distributions are often used in *model checking* (or *model criticism*) where we examine whether there is evidence that we made invalid assumptions by comparing observations with their predictive distributions.

18.2 Example: Chester Road

In section 4.4 we looked at the rate λ of vehicles arriving along a road (in vehicles per second). Our posterior distribution for λ was a gamma(119, 1220) with pdf

$$f_{\lambda}^{(1)}(\lambda) = \frac{1220^{119}}{118!} \lambda^{118} e^{-1220\lambda}.$$

1. What is the posterior probability now of observing j vehicles in the next t seconds?

Given the value of λ , the distribution of the number of vehicles in the next t seconds would be $Poisson(\lambda t)$.

The joint probability (density) of λ and j is

$$\frac{1220^{119}}{118!}\lambda^{118}e^{-1220\lambda}\frac{\lambda^{j}t^{j}e^{-\lambda t}}{j!} = \frac{1220^{119}}{118!}\frac{t^{j}}{j!}\lambda^{118+j}e^{-(1220+t)\lambda}.$$

We integrate out λ to get the marginal probability of j vehicles. By comparing the function with a gamma p.d.f. we see that

$$\int_0^\infty \lambda^{118+j} e^{-(1220+t)\lambda} d\lambda = \frac{\Gamma(119+j)}{(1220+t)^{119+j}} \\ = \frac{(118+j)!}{(1220+t)^{119+j}}$$

Hence the probability of observing j vehicles in the next t seconds is

$$\frac{(118+j)!}{118!j!} \frac{1220^{119}t^j}{(1220+t)^{119+j}} = \begin{pmatrix} 118+j\\ j \end{pmatrix} \left(\frac{1220}{1220+t}\right)^{119} \left(\frac{t}{1220+t}\right)^j.$$

This represents a negative binomial distribution. For large n and τ this would be approximately a Poisson distribution with mean equal to t times the posterior mean for λ .

2. What is the posterior p.d.f. now for the time T to the next arrival?

Given the value of λ , the waiting time has an exponential(λ) distribution. The joint p.d.f. of λ and T is

$$\frac{1220^{119}}{118!}\lambda^{118}e^{-1220\lambda}\lambda e^{-\lambda t} = \frac{1220^{119}}{118!}\lambda^{119}e^{-(1220+t)\lambda}.$$

We integrate out λ to get the marginal p.d.f. for T. By analogy with the gamma p.d.f.,

$$\int_0^\infty \lambda^{119} e^{-(1220+t)\lambda} d\lambda = \frac{119!}{(1220+t)^{120}}.$$

Hence the predictive p.d.f. for the waiting time is

$$\frac{119}{1220} \left(\frac{1220}{1220+t}\right)^{120}$$

which would be approximately a negative exponential distribution for large n.

18.3 Example: beta-binomial

This is an example where the observation model is a discrete distribution.

- Observation model: $X \mid \theta \sim \operatorname{bin}(n, \theta)$.
- Beliefs about $\theta : \theta \sim \text{beta}(a, b)$.
- Predictive distribution: Pr(X = j) for j = 0, 1, ..., n.

$$\begin{aligned} \Pr(X=j) &= \int_0^1 \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \theta^{a-1} (1-\theta)^{b-1} \begin{pmatrix} n \\ j \end{pmatrix} \theta^j (1-\theta)^{n-j} \, d\theta \\ &= \int_0^1 \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \begin{pmatrix} n \\ j \end{pmatrix} \theta^{a+j-1} (1-\theta)^{b+n-j-1} \, d\theta \\ &= \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \begin{pmatrix} n \\ j \end{pmatrix} \frac{\Gamma(a+j)\Gamma(b+n-j)}{\Gamma(a+b+n)} \\ &\times \int_0^1 \frac{\Gamma(a+b+n)}{\Gamma(a+j)\Gamma(b+n-j)} \theta^{a+j-1} (1-\theta)^{b+n-j-1} \, d\theta \\ &= \frac{\Gamma(a+j)}{\Gamma(a)} \frac{\Gamma(b+n-j)}{\Gamma(b)} \frac{\Gamma(a+b)}{\Gamma(a+b+n)} \begin{pmatrix} n \\ j \end{pmatrix} \end{aligned}$$

E.g. a = 2, b = 2, n = 7, j = 3.

$$\Pr(X=3) = \frac{(5-1)!}{(2-1)!} \frac{(6-1)!}{(2-1)!} \frac{(4-1)!}{(11-1)!} \frac{7!}{3!4!}$$
$$= \frac{4!5!3!7!}{1!1!10!3!4!} = \frac{5!7!}{10!} = 0.16667$$

18.4 Example: normal, known precision

- Observation model: $Y \mid \mu \sim N(\mu, \tau^{-1})$.
- Beliefs about $\mu:$

$$\mu \sim N(m, P^{-1})$$

Let $c = P/\tau$ so $P = c\tau$. Then

$$\mu \sim N(m, \ [c\tau]^{-1})$$

• Predictive distribution has density $f_Y(y)$.

The joint density of μ , Y is proportional to

$$\exp\left\{-\frac{c\tau}{2}(\mu-m)^{2}\right\}\tau^{1/2}\exp\left\{-\frac{\tau}{2}(y-\mu)^{2}\right\}.$$

We can write $Y = \mu + \varepsilon$ where $\varepsilon \sim N(0, \tau^{-1})$ independently of μ . So

$$Y \sim N(m, [c\tau]^{-1} + \tau^{-1}).$$

That is

$$Y \sim N\left(m, \ [c_p \tau]^{-1}\right)$$

where

$$c_p = \frac{c}{c+1}$$

and the variance is

$$[c_p\tau]^{-1} = \frac{c+1}{c\tau} = \tau^{-1} + [c\tau]^{-1}.$$

E.g. $m = 23.0, \tau = 0.05, P = 0.25$ so c = 5,

$$c_p = \frac{5}{6}, \ c_p \tau = \frac{0.25}{6}, \ \operatorname{var}(Y) = \frac{6}{0.25} = 24.0.$$

Hence

$$\Pr(Y < 20) = \Phi\left(\frac{20 - 23}{\sqrt{24}}\right) = \Phi(-0.612) = 0.270.$$

18.5 Example: normal, conjugate prior

- Observation model: $Y \mid \mu, \tau \sim N(\mu, \tau^{-1})$.
- Beliefs about μ , τ :

$$\tau \sim \operatorname{gamma}(d/2, dv/2)$$

 $\mu \mid \tau \sim N(m, (c\tau)^{-1})$

• Predictive distribution has density $f_Y(y)$.

The joint density of τ , μ , Y is proportional to

$$\tau^{d/2-1} e^{-dv\tau/2} (c\tau)^{1/2} \exp\left\{-\frac{c\tau}{2} (\mu-m)^2\right\} \tau^{1/2} \exp\left\{-\frac{\tau}{2} (y-\mu)^2\right\}.$$

Conditional on τ , we can write $Y = \mu + \varepsilon$ where $\varepsilon \sim N(0, \tau^{-1})$ independently of μ . So

$$Y \mid \tau \sim N(m, [c\tau]^{-1} + \tau^{-1}).$$

That is

$$Y \mid \tau \sim N\left(m, \ [c_p \tau]^{-1}\right)$$

where

$$c_p = \frac{c}{c+1}$$

and the variance is

$$[c_p\tau]^{-1} = \frac{c+1}{c\tau} = \tau^{-1} + [c\tau]^{-1}.$$

So,

$$\frac{Y-m}{\sqrt{v/c_p}} \sim t_d.$$

E.g. d = 4, v = 0.25, m = 15, c = 0.5 so

$$c_p = \frac{0.5}{1.5} = \frac{1}{3}, \quad \frac{v}{c_p} = 0.75,$$

$$\frac{Y-15}{\sqrt{0.75}} \sim t_4.$$

95% predictive (hpd) interval:

 $15 \pm 2.776 \sqrt{0.75}$

That is

12.596 < Y < 17.404.

18.6 Problems 5

1. (Some of this question is also in Problems 4). I recorded the attendance of students at tutorials for a module. Suppose that we can, in some sense, regard the students as a sample from some population of students so that, for example, we can learn about the likely behaviour of next year's students by observing this year's. At the time I recorded the data we had had tutorials in Week 2 and Week 4. Let the probability that a student attends in both weeks be θ_{11} , the probability that a student attends in week 2 but not Week 4 be θ_{10} and so on. The data are as follows.

Attendance	Probability	Observed frequency
Week 2 and Week 4	$ heta_{11}$	$n_{11} = 25$
Week 2 but not Week 4	$ heta_{10}$	$n_{10} = 7$
Week 4 but not Week 2	θ_{01}	$n_{01} = 6$
Neither week	$ heta_{00}$	$n_{00} = 13$

Suppose that the prior distribution for $(\theta_{11}, \theta_{10}, \theta_{01}, \theta_{00})$ is a Dirichlet distribution with density proportional to

$$\theta_{11}^3 \theta_{10} \theta_{01} \theta_{00}^2.$$

- (a) Find the prior means and prior variances of θ_{11} , θ_{10} , θ_{01} , θ_{00} .
- (b) Find the posterior distribution.
- (c) Find the posterior means and posterior variances of θ_{11} , θ_{10} , θ_{01} , θ_{00} .
- (d) Using the R function hpdbeta which may be obtained from the Web page (or otherwise), find a 95% posterior hpd interval, based on the exact posterior distribution, for θ_{00} .
- (e) Find an approximate 95% hpd interval for θ_{00} using a normal approximation based on the posterior mode and the partial second derivatives of the log posterior density. Compare this with the exact hpd interval.

Hint: To find the posterior mode you will need to introduce a Lagrange multiplier.

- (f) The population mean number of attendances out of two is $\mu = 2\theta_{11} + \theta_{10} + \theta_{01}$. Find the posterior mean of μ and an approximation to the posterior standard deviation of μ .
- 2. Samples are taken from twenty wagonloads of an industrial mineral and analysed. The amounts in ppm (parts per million) of an impurity are found to be as follows.

44.3 50.2 51.7 49.4 50.6 55.0 53.5 48.6 48.8 53.3 59.4 51.4 52.0 51.9 51.6 48.3 49.3 54.1 52.4 53.1

We regard these as independent samples from a normal distribution with mean μ and variance $\sigma^2 = \tau^{-1}$.

Find a 95% posterior hpd interval for μ under each of the following two conditions.

- (a) The value of τ is known to be 0.1 and our prior distribution for μ is normal with mean 60.0 and standard deviation 20.0.
- (b) The value of τ is unknown. Our prior distribution for τ is a gamma distribution with mean 0.1 and standard deviation 0.05. Our conditional prior distribution for μ given τ is normal with mean 60.0 and precision 0.025τ (that is, standard deviation $\sqrt{40\tau^{-1/2}}$).
- 3. We observe a sample of 30 observations from a normal distribution with mean μ and precision τ . The data, y_1, \ldots, y_{30} , are such that

$$\sum_{i=1}^{30} y_i = 672 \qquad \text{and} \qquad \sum_{i=1}^{30} y_i^2 = 16193.$$

- (a) Suppose that the value of τ is known to be 0.04 and that our prior distribution for μ is normal with mean 20 and variance 100. Find the posterior distribution of μ and evaluate a posterior 95% hpd interval for μ .
- (b) Suppose that we have a gamma(1, 10) prior distribution for τ and our conditional prior distribution for μ given τ is normal with mean 20 and variance $(0.1\tau)^{-1}$. Find the marginal posterior distribution for τ , the marginal posterior distribution for μ and the marginal posterior 95% hpd interval for μ .
- 4. The following data come from the experiment reported by MacGregor *et al.* (1979). They give the supine systolic blood pressures (mm Hg) for fifteen patients with moderate essential hypertension. The measurements were taken immediately before and two hours after taking a drug.

Patient	1	2	3	4	5	6	7	8
Before	210	169	187	160	167	176	185	206
After	201	165	166	157	147	145	168	180
Patient	9	10	11	12	13	14	15	
Patient Before	9 173	10 146	11 174	12 201	13 198	14 148	$\frac{15}{154}$	

We are interested in the effect of the drug on blood pressure. We assume that, given parameters μ , τ , the changes in blood pressure, from before to after, in the *n* patients are independent and normally distributed with unknown mean μ and unknown precision τ . The fifteen differences are as follows.

-9 -4 -21 -3 -20 -31 -17 -26 -26 -10 -23 -33 -19 -19 -23

Our prior distribution for τ is a gamma(0.35, 1.01) distribution. Our conditional prior distribution for μ given τ is a normal $N(0, [0.003\tau]^{-1})$ distribution.

- (a) Find the marginal posterior distribution of τ .
- (b) Find the marginal posterior distribution of μ .
- (c) Find the marginal posterior 95% hpd interval for μ .
- (d) Comment on what you can conclude about the effect of the drug.
- 5. The lifetimes of certain components are supposed to follow a Weibull distribution with known shape parameter $\alpha = 2$. The probability density function of the lifetime distribution is

$$f(t) = \alpha \rho^2 t \exp[-(\rho t)^2]$$

for $0 < t < \infty$.

We will observe a sample of n such lifetimes where n is large.

- (a) Assuming that the prior density is nonzero and reasonably flat so that it may be disregarded, find an approximation to the posterior distribution of ρ . Find an approximate 95% hpd interval for ρ when n = 300, $\sum \log(t) = 1305.165$ and $\sum t^2 = 3161776$.
- (b) Assuming that the prior distribution is a gamma(a, b) distribution, find an approximate 95% hpd interval for ρ , taking into account this prior, when $a = 2, b = 100, n = 300, \sum \log(t) = 1305.165$ and $\sum t^2 = 3161776$.
- 6. Given the value of λ , the number X_i of transactions made by customer i at an online store in a year has a Poisson(λ) distribution, with X_i independent of X_j for $i \neq j$. The value of λ is unknown. Our prior distribution for λ is a gamma(5,1) distribution.

We observe the numbers of transactions in a year for 45 customers and

$$\sum_{i=1}^{45} x_i = 182.$$

(a) Using a χ^2 table (i.e. without a computer) find the lower 2.5% point and the upper 2.5% point of the prior distribution of λ .

(These bound a 95% symmetric prior credible interval).

- (b) Find the posterior distribution of λ .
- (c) Using a normal approximation to the posterior distribution, based on the posterior mean and variance, find a 95% symmetric posterior credible interval for λ .
- (d) Find an expression for the posterior predictive probability that a customer makes m transactions in a year.
- (e) As well as these "ordinary customers," we believe that there is a second group of individuals. The number of transactions in a year for a member of this second group has, given θ , a Poisson(θ) distribution and our beliefs about the value of θ are represented by a gamma(1,0.05) distribution.

A new individual is observed who makes 10 transactions in a year. Given that our prior probability that this is an ordinary customer is 0.9, find our posterior probability that this is an ordinary customer.

Hint: You may find it best to calculate the logarithms of the predictive probabilities before exponentiating these. For this you might find the R function lgamma useful. It calculates the log of the gamma function. Alternatively it is possible to do the calculation using the R function dnbinom.

(N.B. In reality a slightly more complicated model is used in this type of application).

7. The following data give the heights in cm of 25 ten-year-old children. We assume that, given the values of μ and τ , these are independent observations from a normal distribution with mean μ and variance τ^{-1} .

66	66	69	61	58	53	78	71	49	57	54	61	49
64	63	60	53	51	65	70	55	55	74	70	42	

- (a) Assuming that the value of τ^{-1} is known to be 64 and our prior distribution for μ is normal with mean 55 and standard deviation 5, find a 95% hpd interval for the height in cm of another ten-year-old child drawn from the same population.
- (b) Assume now that the value of τ is unknown but we have a prior distribution for it which is a gamma(2,128) distribution and our conditional prior distribution for μ given τ is normal with mean 55 and variance $(2.56\tau)^{-1}$. Find a 95% hpd interval for the height in cm of another ten-year-old child drawn from the same population.
- 8. A random sample of n = 1000 people was chosen from a large population. Each person was asked whether they approved of a proposed new law. The number answering "Yes" was x = 372. (For the purpose of this exercise all other responses and non-responses are teated

as simply "Not Yes"). Assume that x is an observation from the binomial(n, p) distribution where p is the unknown proportion of people in the population who would answer "Yes."

Our prior distribution for p is a uniform distribution on (0, 1).

Let $p = \Phi(\theta)$ so $\theta = \Phi^{-1}(p)$ where $\Phi(y)$ is the standard normal distribution function and $\Phi^{-1}(z)$ is its inverse.

- (a) Find the maximum likelihood estimate of p and hence find the maximum likelihood estimate of θ .
- (b) Disregarding the prior distribution, find a large-sample approximation to the posterior distribution of θ .
- (c) Using your approximate posterior distribution for θ , find an approximate 95% hpd interval for θ .
- (d) Use the exact posterior distribution for p to find the actual posterior probability that θ is inside your approximate hpd interval.
- Notes: The standard normal distribution function $\Phi(x) = \int_{-\infty}^{x} \phi(u) \, du$ where $\phi(u) = (2\pi)^{-1/2} \exp\{-u^2/2\}$.
 - Let l be the log-likelihood. Then

$$\frac{dl}{d\theta} = \frac{dl}{dp}\frac{dp}{d\theta}$$

and

$$\frac{d^2l}{d\theta^2} = \frac{d}{d\theta} \left\{ \frac{dl}{dp} \frac{dp}{d\theta} \right\} = \frac{d}{d\theta} \left\{ \frac{dl}{dp} \right\} \frac{dp}{d\theta} + \frac{dl}{dp} \frac{d^2p}{d\theta^2}$$
$$= \frac{d^2l}{dp^2} \left(\frac{dp}{d\theta} \right)^2 + \frac{dl}{dp} \frac{d^2p}{d\theta^2}$$

• Derivatives of p:

$$\frac{dp}{d\theta} = \phi(\theta)$$

$$\frac{d^2p}{d\theta^2} = -\theta\phi(\theta)$$

• You can evaluate $\Phi^{-1}(u)$ using R with

qnorm(u,0,1) and $\phi(u)$ is given by dnorm(u,0,1)

9. The amounts of rice, by weight, in 20 nominally 500g packets are determined. The weights, in g, are as follows.

496	506	495	491	488	492	482	495	493	496
487	490	493	495	492	498	491	493	495	489

Assume that, given the values of parameters μ , τ , the weights are independent and each has a normal $N(\mu, \tau)$ distribution.

The values of μ and τ are unknown. Our prior distribution is as follows. We have a gamma(2, 9) prior distribution for τ and a $N(500, (0.005\tau)^{-1})$ conditional prior distribution for μ given τ .

- (a) Find the posterior probability that $\mu < 495$.
- (b) Find the posterior predictive probability that a new packet of rice will contain less than 500g of rice

10. A machine which is used in a manufacturing process jams from time to time. It is thought that the frequency of jams might change over time as the machine becomes older. Once every three months the number of jams in a day is counted. The results are as follows.

Observation i	1	2	3	4	5	6	7	8
Age of machine t_i (months)	3	6	9	12	15	18	21	24
Number of jams y_i	10	13	24	17	20	22	20	23

Our model is as follows. Given the values of two parameters α , β , the number of jams y_i on a dat when the machine has age t_i months has a Poisson distribution

$$y_i \sim \text{Poisson}(\lambda_i)$$

where

_

$$\log_e(\lambda_i) = \alpha + \beta t_i$$

Assume that the effect of our prior distribution on the posterior distribution is negligible and that large-sample approximations may be used.

(a) Let the values of α and β which maximise the likelihood be $\hat{\alpha}$ and $\hat{\beta}$. Assuming that the likelihood is differentiable at its maximum, show that these satisfy the following two equations

$$\sum_{i=1}^{8} (\hat{\lambda}_i - y_i) = 0$$
$$\sum_{i=1}^{8} t_i (\hat{\lambda}_i - y_i) = 0$$

where

$$\log_e(\hat{\lambda}_i) = \hat{\alpha} + \hat{\beta}t_i$$

and show that these equations are satisfied (to a good approximation) by

$$\hat{\alpha} = 2.552$$
 and $\hat{\beta} = 0.02638$.

(You may use R to help with the calculations, but show your commands).

You may assume from now on that these values maximise the likelihood.

- (b) Find an approximate symmetric 95% posterior interval for $\alpha + 24\beta$.
- (c) Find an approximate symmetric 95% posterior interval for $\exp(\alpha + 24\beta)$, the mean jam-rate per day at age 24 months.

(You may use R to help with the calculations, but show your commands).

Homework 5

Solutions to Questions 9 and 10 of Problems 5 are to be submitted in the Homework Letterbox no later than 4.00pm on **Tuesday** May 5th.

Reference

MacGregor, G.A., Markandu, N.D., Roulston, J.E. and Jones, J.C., 1979. Essential hypertension: effect of an oral inhibitor of angiotensin-converting enzyme. *British Medical Journal*, 2, 1106-1109.