

Applied Bayesian Modeling, ICSPR 2014, Prof. Ryan Bakker

Rose Maier, University of Oregon - Psychology

Assignment 1: Thursday, June 26, 2014

Due: Monday, June 30, 2014

On the first page of your submission, please indicate your *name, home institution, and field of study*.

Please turn in hard copies to the TAs during our course meeting on Monday, June 30.

1. Prove that the gamma distribution,

$$p(\theta|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} e^{-\beta\theta}, \quad \theta, \alpha, \beta > 0 \quad (1)$$

is the conjugate prior distribution for θ in a Poisson probability mass function,

$$p(y_i|\theta) = \frac{e^{-\theta}\theta^{y_i}}{y_i!} \quad (2)$$

that is, calculate a form for the posterior distribution of θ and show that it is also gamma distributed.

First, convert the Poisson probability mass function to a likelihood function by multiplying the probability for each observation in the data together:

$$L(\theta|Y) = \prod_{i=1}^n \frac{e^{-\theta}\theta^{y_i}}{y_i!} \quad (3)$$

Move the terms with θ outside of the product operator.

$$L(\theta|Y) = e^{-n\theta} \theta^{\sum_{i=1}^n y_i} \prod_{i=1}^n \frac{1}{y_i!} \quad (4)$$

Multiply prior and likelihood together.

$$\Pi(\theta|Y) \propto \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} e^{-\beta\theta} e^{-n\theta} \theta^{\sum_{i=1}^n y_i} \prod_{i=1}^n \frac{1}{y_i!} \quad (5)$$

Rearrange to put like terms next to each other.

$$\Pi(\theta|Y) \propto \frac{\beta^\alpha}{\Gamma(\alpha)} \frac{e^{-\beta\theta} e^{-n\theta} \theta^{\alpha-1} \theta^{\sum_{i=1}^n y_i}}{\prod_{i=1}^n y_i!} \quad (6)$$

Combine like terms and drop the constant.

$$\Pi(\theta|Y) \propto e_1^{-(\beta+n)\theta} \theta^{\alpha + \sum_{i=1}^n (y_i)-1} \quad (7)$$

This looks like the kernel of a new gamma distribution, with $\alpha = \alpha + \sum_{i=1}^n (y_i)$ and $\beta = \beta + n$.

2. Now use the Gamma-Poisson conjugate specification to analyze data on the number of presidential appointments from 1960 to 2000. The data are in the file called `hw1.dta` on the Z-drive.
3. The posterior distribution for θ is $\text{gamma}(\delta_1, \delta_2)$ according to some parameters δ_1 and δ_2 that you derived above which of course depend on your choice of the parameters for the gamma prior. You should model θ using two sets of priors. One which specifies a great deal of certainty regarding your best guess as to the value of θ and one that represents ignorance regarding this value.

I choose $\delta_1 = 1, \delta_2 = 1$ for my uncertain prior, and $\delta_1 = 100, \delta_2 = 100$ for my very certain prior.

4. Generate a large number of values from this distribution in R, say 10,000 or so, using the command:

```
posterior.sample <- rgamma(10000,d1,d2)
```

5. Summarize the posteriors with quantities of interest such as means, medians, and variances. Also supply plots of the density of the posterior distributions.

```
setwd("/Users/TARDIS/Documents/ICPSR2014/Bayes/Homeworks/HW 1")
library(foreign)
df <- read.dta("hw1.dta")

n <- nrow(df)

# priors
d1_pri <- c(1,100)
d2_pri <- c(1,100)
mean_pri <- d1_pri / d2_pri
var_pri <- d1_pri / (d2_pri)^2

prior.sample1 <- rgamma(100000,shape=d1_pri[1],rate=d2_pri[1])
mean(prior.sample1)

## [1] 1.001

var(prior.sample1)

## [1] 0.998

prior.sample2 <- rgamma(100000,shape=d1_pri[2],rate=d2_pri[2])
mean(prior.sample2)

## [1] 1

var(prior.sample2)
```

```

## [1] 0.01002

# posteriors
d1_pos <- d1_pri + sum(df$appoints)
d2_pos <- d2_pri + n
mean_pos <- d1_pos / d2_pos
var_pos <- d1_pos / (d2_pos)^2

posterior.sample1 <- rgamma(100000, shape=d1_pos[1], rate=d2_pos[1])
mean(posterior.sample1)

## [1] 1.635

median(posterior.sample1)

## [1] 1.606

var(posterior.sample1)

## [1] 0.1485

posterior.sample2 <- rgamma(100000, shape=d1_pos[2], rate=d2_pos[2])
mean(posterior.sample2)

## [1] 1.064

median(posterior.sample2)

## [1] 1.061

var(posterior.sample2)

## [1] 0.00971

## plotting
plotdf1 <- data.frame(posterior.sample=posterior.sample1,
prior.sample=prior.sample1)
plotdf1$prior <- as.factor(paste("prior d1=", d1_pri[1],
" d2=", d2_pri[1], sep=""))

plotdf2 <- data.frame(posterior.sample=posterior.sample2,
prior.sample=prior.sample2)
plotdf2$prior <- as.factor(paste("prior d1=", d1_pri[2],
" d2=", d2_pri[2], sep=""))

plotdf <- rbind(plotdf1, plotdf2)
plotdf <- melt(plotdf, id.vars="prior")

```

```

library(plyr)
means <- ddply(plotdf, .(prior, variable),
               summarize, means=mean(value))
plotdf <- merge(plotdf,means, by=c("prior", "variable"))

library(ggplot2)
ggplot(plotdf, aes(x=value)) +
  geom_histogram(position="identity",
                 binwidth = .05, alpha=.5, aes(fill=variable)) +
  xlim(0,4) +
  geom_vline(data=plotdf,
             aes(xintercept=means, color=variable), lty=2) +
  facet_wrap(~prior)

```

