Bonus Exercises: Bayesian Inference for Conjugate Models

2026

These exercises use the same functions and methods from today's labs. Make sure `binbetaf.R`, `binbetaf_plot.R`, and `binbetaf_plot_CI.R` are saved in your working folder before starting.

Bonus 1: Free Throw Shooting

A basketball player practises free throws. In a training session she attempts **25 free throws** and makes **16**. Let π be her true free throw success probability.

You have no prior information about her ability.

1. What prior distribution does "no prior information" correspond to? Write it down and state its parameters.
2. What is the posterior distribution of π? Identify a_1 and b_1.
3. Use `binbetaf` to find the posterior mean, standard deviation, and a 95% credible interval.
4. Use `binbetaf_plot` to plot the prior and posterior. Describe what you observe.
5. Is the normal approximation to the posterior valid here? Compute and compare both the exact and approximate 95% credible interval.

```
source("binbetaf.R")
source("binbetaf_plot.R")
source("binbetaf_plot_CI.R")

# Fill in the correct values
binbetaf(___, ___, ___, ___, 0.05)
binbetaf_plot(___, ___, ___, ___, "topleft")
```

Bonus 2: PhD Acceptance Rate

You are interested in the proportion π of economics students who are accepted to a PhD programme after applying. Before seeing any data, you believe:

$$E(\pi) = 0.30 \qquad \text{SD}(\pi) = 0.10$$

[0]Email: paulius.kazlauskas@evaf.vu.lt

A survey of **500 applicants** finds that **180** were accepted.

1. Determine the Beta(a, b) prior that matches your belief. Use the formulas:

$$a + b = \frac{\mu_0(1-\mu_0)}{\sigma_0^2} - 1 \qquad a = \mu_0(a+b) \qquad b = (1-\mu_0)(a+b)$$

2. What is the posterior distribution of π?
3. Calculate the prior and data weights in the posterior mean. What do they tell you about the relative influence of your prior belief versus the data?
4. Is the normal approximation to the posterior valid here? Write down the approximate distribution.
5. Use `binbetaf` to compute the posterior summaries, and `binbetaf_plot_CI` to visualise the posterior with the 95% credible interval.

```r
source("binbetaf.R")
source("binbetaf_plot.R")
source("binbetaf_plot_CI.R")

a      <- ___   # from question 1
b      <- ___
n      <- 500
y      <- 180
alpha <- 0.05

binbetaf(a, b, n, y, alpha)
binbetaf_plot_CI(a, b, n, y, "topright", alpha)
```

Bonus 3: Website Hourly Traffic

The number of visitors to a website per hour follows a Poisson(μ) distribution. Before collecting data you believe:

$$E(\mu) = 50 \qquad \text{SD}(\mu) = 5$$

The observed visitor counts over 7 hours are:

$$48,\ 55,\ 52,\ 46,\ 61,\ 50,\ 53$$

1. Determine the Gamma(r, v) prior using:

$$r = \frac{\mu_0^2}{\sigma_0^2} \qquad v = \frac{\mu_0}{\sigma_0^2}$$

2. Find the posterior distribution of μ, recalling that $r_1 = r + \sum y_i$ and $v_1 = v + n$.
3. Compute the posterior mean, standard deviation, prior weight, and data weight. Verify the weighted average decomposition in R.

4. Compute a 95% credible interval using both the exact Gamma quantiles (`qgamma`) and the normal approximation. Are the two intervals similar?

5. Calculate $P(\mu > 55)$ using both the exact Gamma distribution (`pgamma`) and the normal approximation (`pnorm`). Interpret the result.

```r
y <- c(48, 55, 52, 46, 61, 50, 53)
m <- 50
s <- 5

# Prior parameters
r <- ___
v <- ___

# Posterior parameters
sum_y <- sum(y)
n     <- length(y)
r1    <- ___
v1    <- ___

# Posterior mean and SD
pos_mean <- ___
pos_sd   <- ___

# 95% credible interval (exact)
qgamma(___, r1, v1)
qgamma(___, r1, v1)

# 95% credible interval (normal approximation)
pos_mean - 1.96 * pos_sd
pos_mean + 1.96 * pos_sd

# P(mu > 55): exact and normal approximation
1 - pgamma(___, r1, v1)
1 - pnorm(___, pos_mean, pos_sd)
```