Exercises: Bayesian Inference for a Poisson Mean

2026

The Gamma–Poisson is a conjugate model for a Poisson rate μ:

- **Prior:** $\mu \sim \text{Gamma}(r, v)$
- **Likelihood:** $f(\mathbf{y} \mid \mu) \propto \mu^{\sum y_i} e^{-n\mu}$
- **Posterior:** $\mu \mid \mathbf{y} \sim \text{Gamma}(r_1, v_1)$, where $\mathbf{r_1} = \mathbf{r} + \sum \mathbf{y_i}$ and $\mathbf{v_1} = \mathbf{v} + \mathbf{n}$

Reference formulas for the Gamma distribution:

$$E(\mu) = \frac{r}{v} \qquad \text{Var}(\mu) = \frac{r}{v^2} \qquad \text{SD}(\mu) = \sqrt{\frac{r}{v^2}}$$

$$w_{\text{prior}} = \frac{v}{v+n} \qquad w_{\text{data}} = \frac{n}{v+n} \qquad \hat{\mu} = E(\mu \mid \mathbf{y}) = \frac{r_1}{v_1}$$

R functions: `pgamma(q, shape, rate)` gives $P(\mu \leq q)$; `qgamma(p, shape, rate)` gives the p-th quantile.

Question 1: BC Žalgiris

The number of points that BC Žalgiris scores in a game follows a Poisson(μ) distribution. The points observed over the last ten games are:

64, 72, 84, 73, 98, 85, 85, 94, 72, 93

Your prior belief about μ is that it has mean 80 and standard deviation 10.

1. Determine the Gamma(r, v) prior

We need r and v such that $E(\mu) = 80$ and $\text{Var}(\mu) = 10^2 = 100$:

$$\frac{r}{v} = 80 \qquad \text{and} \qquad \frac{r}{v^2} = 100$$

Dividing the second equation by the first gives $1/v = 100/80$, so:

$$v = \frac{80}{100} = 0.8 \qquad \text{and} \qquad r = 80 \times 0.8 = 64$$

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

The prior distribution is $\mu \sim \text{Gamma}(64,\ 0.8)$.

R: Fill in the formulas for `r` and `v`.

```r
# Data
y <- c(64, 72, 84, 73, 98, 85, 85, 94, 72, 93)

# Prior belief
m <- 80
s <- 10

# Prior parameters: r = m^2/s^2, v = m/s^2
r <- ???
v <- ???
```

2. Find the posterior distribution of μ

Using $\sum y_i = 820$ and $n = 10$, the posterior parameters are:

$$r_1 = r + \sum y_i = 64 + 820 = 884 \qquad \text{and} \qquad v_1 = v + n = 0.8 + 10 = 10.8$$

The posterior distribution is: $\mu \mid \mathbf{y} \sim \text{Gamma}(884,\ 10.8)$.

R: Compute the posterior parameters.

```r
sum_y <- sum(y)        # 820
n     <- length(y)     # 10

r1 <- ???              # r + sum_y
v1 <- ???              # v + n
```

3. Posterior mean decomposition

$$\text{Posterior mean} = \frac{r_1}{v_1} = \frac{884}{10.8} \approx 81.85$$

	Formula	Value
Prior weight	$v\,/\,(v+n)$	$0.8/10.8 \approx 0.074$
Prior mean	r/v	$64/0.8 = 80$
Data weight	$n\,/\,(v+n)$	$10/10.8 \approx 0.926$
Data mean	$\sum y_i/n$	$820/10 = 82$
$0.074 \times 80 + 0.926 \times 82$		$= 81.85$ ✓

R: Compute each quantity and verify the identity numerically.

```r
# Posterior mean, variance, SD
pos_mean <- ???          # r1 / v1
pos_var  <- ???          # r1 / v1^2
```

```r
pos_sd    <- sqrt(pos_var)

# Weights and component means
w_prior    <- ???         # v / (v + n)
w_data     <- ???         # n / (v + n)
pri_mean   <- r / v
data_mean <- sum_y / n

# Verify: should equal pos_mean
w_prior * pri_mean + w_data * data_mean
```

4. Optimum estimator of μ post-data

The optimum estimator (minimising posterior expected squared error) is the **posterior mean**:

$$\hat{\mu} = \frac{r_1}{v_1} = \frac{884}{10.8} \approx 81.85$$

R:

```r
pos_mean    # the optimum estimator
```

5. 95% Credible interval for μ

Exact (Gamma quantiles):

$$\left[q_{0.025},\ q_{0.975}\right] = \left[\texttt{qgamma}(0.025,\ 884,\ 10.8),\ \texttt{qgamma}(0.975,\ 884,\ 10.8)\right] = [76.54,\ 87.33]$$

Normal approximation: With $m_1 = 81.85$ and $s_1 = \sqrt{884/10.8^2} = 2.75$:

$$m_1 \pm 1.96 \times s_1 = 81.85 \pm 1.96 \times 2.75 = [76.46,\ 87.24]$$

Both intervals are nearly identical, confirming the normal approximation works well.

R: Fill in the correct arguments.

```r
# Exact credible interval
qgamma(???, r1, v1)      # lower bound
qgamma(???, r1, v1)      # upper bound

# Normal approximation CI
pos_mean - 1.96 * pos_sd
pos_mean + 1.96 * pos_sd
```

6. Probability of scoring fewer than 75 points

Exact (Gamma CDF):

$$P(\mu < 75) = \texttt{pgamma}(75,\ 884,\ 10.8) \approx 0.0054$$

Normal approximation: Approximating $\mu \mid \mathbf{y} \approx N(81.85,\ 2.75^2)$:

$$P(\mu < 75) = P\left(Z < \frac{75 - 81.85}{2.75}\right) = P(Z < -2.49) \approx 0.0064$$

Scoring fewer than 75 points would be quite unusual for Žalgiris given the data.

R: Compute both probabilities.

```
# Exact: P(mu < 75)
pgamma(???, r1, v1)

# Normal approximation
pnorm(???, pos_mean, pos_sd)
```