Lab 2: Bayesian Inference for Discrete Random Variables

2026

1 Discrete Probability Distribution

Practice 1: Let Y be a discrete random variable with probability function given in the following table.

y_i	$f(y_i)$
1	0.1
2	0.2
3	0.1
4	0.2
5	0.4

1. Calculate $P(1 < Y < 5)$
2. Calculate $E(Y)$ and $\text{Var}(Y)$
3. Suppose $W = 2Y + 5$. Calculate $E(W)$ and $\text{Var}(W)$
4. Calculate $E(Y + W)$ and $\text{Var}(Y + W)$

Answer

1. Calculate $P(1 < Y < 5)$

$$P(1 < Y < 5) = P(Y = 2) + P(Y = 3) + P(Y = 4) = 0.2 + 0.1 + 0.2 = 0.5$$

2. Calculate $E(Y)$ **and** $\text{Var}(Y)$

$$E(Y) = \sum_i y_i f(y_i) = 1 \times 0.1 + 2 \times 0.2 + 3 \times 0.1 + 4 \times 0.2 + 5 \times 0.4 = 3.6$$
$$E(Y^2) = \sum_i y_i^2 f(y_i) = 1^2 \times 0.1 + 2^2 \times 0.2 + 3^2 \times 0.1 + 4^2 \times 0.2 + 5^2 \times 0.4 = 15$$
$$\text{Var}(Y) = E(Y^2) - E(Y)^2 = 15 - 3.6^2 = 2.04$$

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

3. Suppose $W = 2Y + 5$. Calculate $E(W)$ and $\text{Var}(W)$

$$E(W) = 2E(Y) + 5 = 2 \times 3.6 + 5 = 12.2$$
$$\text{Var}(W) = 2^2\text{Var}(Y) = 4 \times 2.04 = 8.16$$

4. Calculate $E(Y + W)$ and $\text{Var}(Y + W)$

$$E(Y + W) = E(Y) + E(W) = 3.6 + 12.2 = 15.8$$
$$\text{Var}(Y + W) = \text{Var}(Y + 2Y + 5) = \text{Var}(3Y + 5) = 9\,\text{Var}(Y) = 9 \times 2.04 = 18.36$$

R Practice

Calculate $E(Y)$ and $\text{Var}(Y)$

```
# Create values of y
y <- seq(1, 5)

# Probability of each value
py <- c(0.1, 0.2, 0.1, 0.2, 0.4)

# P(1 < Y < 5)
which(y > 1 & y < 5)
sum(py[which(y > 1 & y < 5)])

# E(Y)
my <- sum(y * py)
my

# Var(Y)
vary <- sum(y^2 * py) - my^2
vary
```

Simulate a sample with n observations from this discrete probability distribution.

```
# Create a sample with replacement using
# the above discrete probability distribution

# set.seed(42): Set the seed for random simulation
# Number of random observations
size <- 10^5

# Sample with replacement
ys <- sample(y, size, replace = TRUE, prob = py)
```

Compare the R calculations with the above results

1. Calculate $P(1 < Y < 5)$

```
# Method 1: P(1 < Y < 5) = P(Y <= 4) - P(Y <= 1)
p1 <- which(ys <= 4)
p2 <- which(ys <= 1)
```

```r
P1 <- (length(p1) - length(p2)) / length(ys)
P1

# Method 2: P(1<Y<5) = P(Y=2) + P(Y=3) + P(Y=4)
d2 <- which(ys == 2)
d3 <- which(ys == 3)
d4 <- which(ys == 4)
P2 <- (length(d2) + length(d3) + length(d4)) / length(ys)
P2

# Method 3: P(1<Y<5)
P3 <- length(which(ys < 5 & ys > 1)) / length(ys)
P3
```

2. Calculate $E(Y)$ and $\text{Var}(Y)$

```r
# Calculate mean Y
ms <- mean(ys)
ms

# Calculate variance of Y
vars <- var(ys)
vars
```

3. Suppose $W = 2Y + 5$. Calculate $E(W)$ and $\text{Var}(W)$

```r
# Create W = 2Y + 5
w <- 2 * ys + 5

# Calculate mean W
# Method 1
mw1 <- mean(w)
mw1

# Method 2: E(W) = 2*E(Y) + 5
mw2 <- 2 * ms + 5
mw2

# Compare mw1 and mw2
mw1 == mw2

# Calculate variance of W
# Method 1
varw1 <- var(w)
varw1

# Method 2: Var(W) = 2^2 * Var(Y)
varw2 <- 2^2 * vars
varw2

# Compare varw1 and varw2
varw1 == varw2
```

4. Calculate $E(Y+W)$ and $\text{Var}(Y+W)$

```
# E(W + Y)
mwy <- mean(w + ys)
mwy

# Var(W + Y)
varwy <- var(w + ys)
varwy
```

Practice 2: Bayesian Inference with Balls

Practice 2: There is a box containing a total of 4 balls, some of which may be red and the rest of which are green. You do not know how many of the balls are red, so assume that your prior probability is equal across possible values of X.

1. You draw a ball randomly from the box and the selected ball is red. Calculate the posterior distribution probability. Also prove that:
   - Multiplying all the prior probabilities by a constant does not change the result of Bayes' theorem.
   - Multiplying the likelihood by a constant does not change the result of Bayes' theorem.
2. **Without replacement:** You take the second ball from the box (without replacing the first ball) and it is red again. Calculate the posterior distribution probability in two ways: "Analyzing the observations sequentially one at a time" and "Analyzing the observations all together in a single step". Prove that they produce the same result.
3. **With replacement:** You take the second ball from the box (after putting the first ball back) and it is red again. Calculate the posterior distribution probability in two ways: "Analyzing the observations sequentially one at a time" and "Analyzing the observations all together in a single step". Prove that they produce the same result.

Answer

1. First draw: red ball.

x_i	Prior	Likelihood	Prior × Likelihood	Posterior
0	1/5	0/4	0	0
1	1/5	1/4	1/20	1/10
2	1/5	2/4	2/20	2/10
3	1/5	3/4	3/20	3/10
4	1/5	4/4	4/20	4/10
Sum			10/20 = 1/2	

R Corner: Calculate the posterior distribution probability.

```r
# Values of X
x <- seq(0, 4, 1)

# Prior Distribution
priorx <- rep(1/5, 5)
priorx

# Likelihood
likelihood <- x / 4

# Prior * Likelihood
weightx <- priorx * likelihood

# Posterior Distribution
posteriorx <- weightx / sum(weightx)
posteriorx

# Plot posterior distribution
plot(x, posteriorx,
     col  = "blue",
     ylim = c(0, 1),
     main = "Posterior distribution",
     xlab = "x",
     ylab = "y")

# Overlay prior distribution
lines(x, priorx, type = "p", col = "red")

# Table of results
results1Q2 <- cbind(x, priorx, likelihood, weightx, posteriorx)
colnames(results1Q2)[5] <- "Posteriorx"
results1Q2
```

Prove: multiplying priors by a constant does not change the posterior.

```r
# Prior multiplied by 5
priora <- 5 * rep(1/5, 5)
likelihood <- x / 4

weighta    <- priora * likelihood
posteriora <- weighta / sum(weighta)
posteriora

# Proof:
posteriorx == posteriora
```

Prove: multiplying likelihood by a constant does not change the posterior.

```r
# Likelihood multiplied by 4
priorb      <- rep(1/5, 5)
likelihoodb <- (x / 4) * 4
```

```
weightb     <- priorb * likelihoodb
posteriorb <- weightb / sum(weightb)
posteriorb

# Proof:
posteriorx == posteriorb
```

2. Without replacement: second draw is also red.

Sequential approach: Use the posterior from draw 1 as the prior for draw 2. Note only 3 balls remain.

x_i	Prior	Likelihood	Prior × Likelihood	Posterior
0	0	0/3	0	0
1	1/10	0/3	0	0
2	2/10	1/3	2/30	1/10
3	3/10	2/3	6/30	3/10
4	4/10	3/3	12/30	6/10
Sum			$20/30 = 2/3$	

R Corner: Without replacement – sequential.

```
# Use posterior from draw 1 as prior for draw 2
priorx2 <- posteriorx

# Likelihood of the second draw (without replacement, 3 balls
   left)
likelihood2 <- rep(0, 5)
likelihood2[1] <- 0
for (i in 1:4) {
  likelihood2[i + 1] <- (i - 1) / 3
}
likelihood2

# Prior * Likelihood
weightx2 <- priorx2 * likelihood2

# Posterior distribution
posteriorx2 <- weightx2 / sum(weightx2)
posteriorx2
```

Single-step approach: Use the original flat prior and the joint likelihood of both draws.

x_i	Prior	Likelihood	Prior × Likelihood	Posterior
0	1/5	0	0	0
1	1/5	$\frac{1}{4} \times \frac{0}{3} = 0$	0	0
2	1/5	$\frac{2}{4} \times \frac{1}{3} = \frac{2}{12}$	$\frac{2}{60}$	1/10
3	1/5	$\frac{3}{4} \times \frac{2}{3} = \frac{6}{12}$	$\frac{6}{60}$	3/10
4	1/5	$\frac{4}{4} \times \frac{3}{3} = \frac{12}{12}$	$\frac{12}{60}$	6/10
Sum			$20/60 = 1/3$	

R Corner: Without replacement – single step.

```
# Original flat prior
priorx <- rep(1/5, 5)

# Joint likelihood: P(red1) * P(red2 | red1, no replacement)
likelihood3 <- c(0:4) / 4 * c(0:3) / 3

# Prior * Likelihood
weightx3 <- priorx * likelihood3

# Posterior distribution
posteriorx3 <- weightx3 / sum(weightx3)
posteriorx3

# Proof: both approaches give the same result
posteriorx2 == posteriorx3
```

As can be seen, both approaches yield the same posterior distribution.

3. With replacement: second draw is also red.

Sequential approach: Use the posterior from draw 1 as the prior for draw 2. Since the ball is replaced, all 4 balls are available again.

x_i	Prior	Likelihood	Prior × Likelihood	Posterior
0	0	0/4	0	0
1	1/10	1/4	1/40	1/30
2	2/10	2/4	4/40	4/30
3	3/10	3/4	9/40	9/30
4	4/10	4/4	16/40	16/30
Sum			$30/40 = 3/4$	

R Corner: With replacement – sequential.

```
# Use posterior from draw 1 as prior for draw 2
priorx_wr <- posteriorx

# Likelihood of second draw (with replacement, still 4 balls)
likelihood_wr2 <- x / 4

# Prior * Likelihood
weightx_wr2 <- priorx_wr * likelihood_wr2
```

```
# Posterior distribution
posteriorx_wr2 <- weightx_wr2 / sum(weightx_wr2)
posteriorx_wr2
```

Single-step approach: Use the original flat prior and the joint likelihood of both draws (independent draws with replacement).

x_i	Prior	Likelihood	Prior × Likelihood	Posterior
0	1/5	$(0/4)^2 = 0$	0	0
1	1/5	$(1/4)^2 = 1/16$	1/80	1/30
2	1/5	$(2/4)^2 = 4/16$	4/80	4/30
3	1/5	$(3/4)^2 = 9/16$	9/80	9/30
4	1/5	$(4/4)^2 = 16/16$	16/80	16/30
Sum			$30/80 = 3/8$	

R Corner: With replacement – single step.

```
# Original flat prior
priorx <- rep(1/5, 5)

# Joint likelihood: draws are independent (with replacement)
likelihood_wr3 <- (x / 4)^2

# Prior * Likelihood
weightx_wr3 <- priorx * likelihood_wr3

# Posterior distribution
posteriorx_wr3 <- weightx_wr3 / sum(weightx_wr3)
posteriorx_wr3

# Proof: both approaches give the same result
posteriorx_wr2 == posteriorx_wr3
```

As can be seen, both approaches again yield the same posterior distribution.

2 Binomial and Poisson Distribution

Each probability distribution is associated with a core name: `binom` for the binomial distribution and `pois` for the Poisson distribution.

The four basic associated functions are the pdf (probability density function), the cdf (cumulative distribution function), the quantile function, and the simulation procedure: adding the prefixes `d`, `p`, `q`, and `r`, respectively, to the core name. For example:

- Binomial distribution: `dbinom`, `pbinom`, `qbinom`, `rbinom`.
- Poisson distribution: `dpois`, `ppois`, `qpois`, `rpois`.

Practice 3: Let $Y \sim \text{Poisson}(\mu)$. Suppose we believe there are only four possible values for μ: $1, 1.5, 2, 2.5$. Suppose we consider that the two values 1 and 2 are twice as likely as

1.5 and 2.5. Suppose $Y = 1$ was observed. Calculate the posterior distribution probability.

Answer

Using the PMF of the Poisson distribution with $y = 1$:

$$f(y|\mu) = \frac{\mu^y e^{-\mu}}{y!} = \mu e^{-\mu}$$

μ	Prior	Likelihood	Prior $\times$ Likelihood	Posterior
1.0	2/6	$1.0 \times e^{-1.0} = 0.3679$	0.1226	0.4049
1.5	1/6	$1.5 \times e^{-1.5} = 0.3347$	0.0558	0.1843
2.0	2/6	$2.0 \times e^{-2.0} = 0.2707$	0.0902	0.2979
2.5	1/6	$2.5 \times e^{-2.5} = 0.2052$	0.0342	0.1129
Sum			0.3028	

R Corner:

```
# Values of mu
mu <- seq(1, 2.5, 0.5)  # or mu <- c(1, 1.5, 2, 2.5)

# Prior Distribution
prior_mu <- c(2/6, 1/6, 2/6, 1/6)
prior_mu

# Likelihood using dpois(y=1, lambda=mu)
likelihood <- dpois(1, mu)
likelihood

# Prior * Likelihood
weight_mu <- prior_mu * likelihood

# Posterior Distribution
posterior_mu <- weight_mu / sum(weight_mu)
posterior_mu
```

Practice 4: Suppose your evil twin brother has two coins: one comes up heads 30% of the time and the other 50% of the time. He comes to you with one of them and wants to make a bet. You set your prior: 0.6 for the loaded coin and 0.4 for the fair coin. You flip it three times and get two heads and one tail. Which coin do you think it is and how sure are you?

Answer

Using the PMF of the binomial distribution with $y = 2$, $n = 3$:

$$f(y|\pi) = \binom{n}{y}\pi^y(1-\pi)^{n-y} = \binom{3}{2}\pi^2(1-\pi) = 3\pi^2(1-\pi)$$

π	Prior	Likelihood	Prior $\times$ Likelihood	Posterior
0.3 (Loaded)	0.6	$3 \times 0.3^2 \times 0.7 = 0.1890$	0.1134	0.4305
0.5 (Fair)	0.4	$3 \times 0.5^2 \times 0.5 = 0.3750$	0.1500	0.5695
Sum			0.2634	

R Corner:

```r
# Values of pi
pi_vals <- c(0.3, 0.5)

# Prior Distribution
prior_pi <- c(0.6, 0.4)
prior_pi

# Likelihood using dbinom(y=2, size=3, prob=pi)
likelihood <- dbinom(2, 3, pi_vals)
likelihood

# Prior * Likelihood
weight_pi <- prior_pi * likelihood

# Posterior Distribution
posterior_pi <- weight_pi / sum(weight_pi)
posterior_pi
```

3 Bonus Problems

Bonus Practice: Bayesian Updating with a Loaded Die

★ Bonus

Your friend claims to have a fair six-sided die ($\pi = 1/6$ per face). You are suspicious: you think there is a 30% chance the die is loaded, where a loaded die rolls a six with probability $\pi = 1/2$ and all other faces equally with the remaining probability. You set your prior as:

$$P(\text{loaded}) = 0.30, \qquad P(\text{fair}) = 0.70.$$

Your friend rolls the die **three times** and gets: **6, 6, 3**.

1. Identify the likelihood function for each hypothesis. Use `dbinom` or write out the expression by hand.
2. Construct the posterior table (prior, likelihood, prior × likelihood, posterior).
3. After seeing the data, what is the posterior probability that the die is loaded?
4. Suppose your friend rolls a **fourth** time and gets another **6**. Use your posterior from (3) as the new prior and update again. How has your belief changed?

R: Implement the full updating procedure (both steps) using the `bayes_update` function from the lab.

Hint. For question (1), the two rolls of 6 and one roll of non-6 are independent. Under the fair die hypothesis: $P(\text{data} \mid \text{fair}) = (1/6)^2 \times (5/6)$. Under the loaded die: the non-six faces each have probability $(1/2)/5 = 1/10$, so $P(\text{data} \mid \text{loaded}) = (1/2)^2 \times (1/2) \cdot \ldots$ — think carefully about what "loaded" implies for non-six faces, and set up the binomial expression accordingly.

Bonus 1: What is R Markdown?

★ Bonus

R Markdown is a file format (`.Rmd`) that lets you combine plain prose, mathematics, and executable R code in a single document. When you *knit* the file, R runs every code chunk and weaves the output — numbers, tables, plots — directly into the final document, which can be rendered as PDF, HTML, or Word.

Task.

1. In RStudio, create a new R Markdown file (*File* → *New File* → *R Markdown*). Keep the default HTML output.
2. Delete the template content and reproduce **Practice 1** from today's lab inside the `.Rmd` file:
   - Write a short sentence in prose explaining what the random variable Y is.
   - Insert a code chunk that defines `y` and `py`, computes `E(Y)` and `Var(Y)`, and prints the results.
   - Below the chunk, write one sentence interpreting the numbers.
3. Knit the document and verify that the output values match the hand-

calculated answers from the lab.

Minimal R Markdown skeleton:

```
---
title: "Lab 2 -- Practice 1"
output: html_document
---

## Discrete Random Variable Y

*Describe Y here in one sentence.*

```{r}
y  <- c(1, 2, 3, 4, 5)
py <- c(0.1, 0.2, 0.1, 0.2, 0.4)

my   <- sum(y * py)
vary <- sum(y^2 * py) - my^2

cat("E(Y) =", my, "\n")
cat("Var(Y) =", vary, "\n")
```

Interpret the results here.
```

Bonus 2: The Pipe Operator |>

★ Bonus

Chapter 4 of *R for Data Science* (2nd ed., Wickham et al., `https://r4ds.hadley.nz/workflow-style.html`) introduces the native pipe operator `|>`, available since R 4.1.

The pipe passes the result of the left-hand side as the **first argument** of the right-hand side function. The two lines below are equivalent:

```
# Without pipe
round(mean(c(1, 2, 3, 4, 5)), digits = 2)

# With pipe
c(1, 2, 3, 4, 5) |> mean() |> round(digits = 2)
```

Tasks.

1. Read Section 4.3 of the textbook linked above. In your own words (2–3 sentences), explain what the pipe does and why it improves code readability.
2. The code below computes the posterior from Practice 2 (first draw) without using the pipe. Rewrite it using `|>` wherever natural.

```
x          <- seq(0, 4, 1)
priorx     <- rep(1/5, 5)
```

```r
likelihood <- x / 4
weightx    <- priorx * likelihood
posteriorx <- weightx / sum(weightx)
round(posteriorx, 3)
```

3. Does using the pipe change the result? Explain why or why not in one sentence.

Challenge: Can you rewrite the `bayes_update` function from the lab so its body uses the pipe?

Note on the pipe. The pipe works best when you have a linear sequence of transformations on a single object. It is less natural when you need to save an intermediate result for later use (e.g. you need both `weightx` and `posteriorx`). Good style is knowing when the pipe helps and when it obscures.