In-class Exercise

An inspection system has prior defect rate

$$\pi \sim \mathrm{Beta}(2, 38).$$

A new audit finds $y = 4$ defective items among $n = 20$.

1. Write the posterior distribution.
2. Compute its mean.
3. Define an action threshold using false-alarm cost $C_{FP} = 1$ and missed-defect cost $C_{FN} = 9$.
4. What posterior probability must be estimated to decide whether to intervene when "high risk" means $\pi > 0.10$?
5. Name one model-risk concern that the calculation omits.

Answer:

$$\pi \mid y \sim \text{Beta}(2+4, 38+16) = \text{Beta}(6, 54), \qquad \mathbb{E}(\pi \mid y) = 0.10.$$

Let $q = \text{P}(\pi > 0.10 \mid y)$. Intervene when

$$q > \frac{C_{FP}}{C_{FP} + C_{FN}} = \frac{1}{10} = 0.10$$

Because the posterior mean equals the risk threshold, q is near one-half, so intervention is likely under this asymmetric loss.

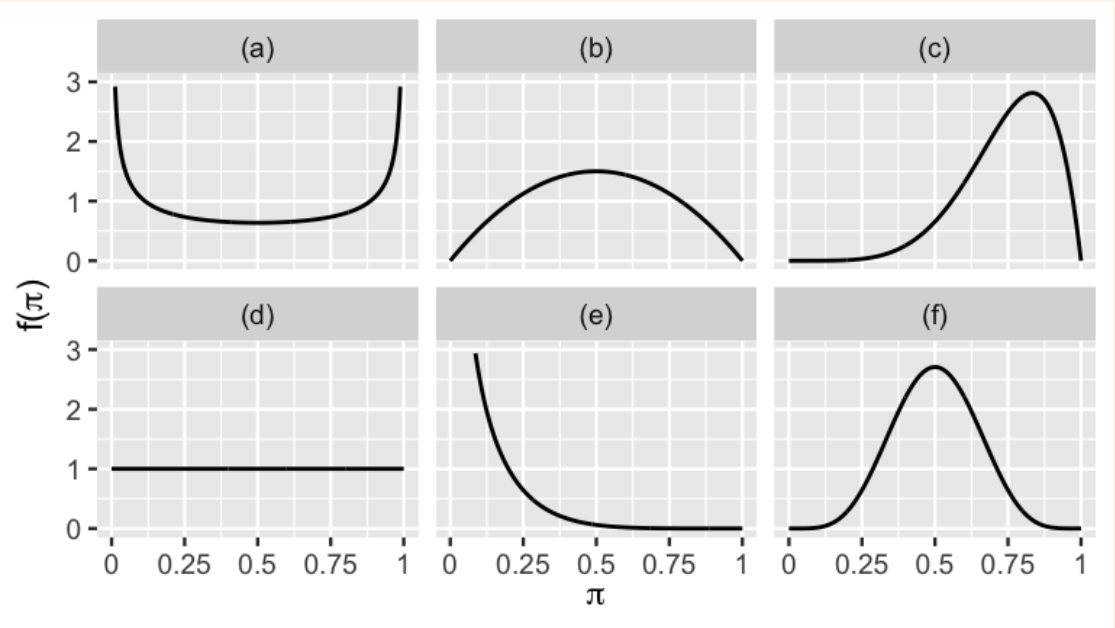
Possible model risks: clustered defects, changing production conditions, imperfect detection, or a nonrepresentative audit.

Practice Questions

1. If prior concentration doubles but its mean stays fixed, what changes?
2. Why is $L(\pi \mid y)$ not $\mathrm{P}(\pi \mid y)$?
3. What does conjugacy buy us computationally?
4. When can two analysts share a posterior but rationally choose different actions?
5. Which assumption would you challenge first in the election-poll example?
6. Identify the most decision-sensitive assumption.
7. Provide one example where false-positive and false-negative costs differ.

Practice Questions

8. Six Beta pdfs are plotted below.
   Match each to one of the following models:
   Beta(0.5,0.5), Beta(1,1), Beta(2,2), Beta(6,6), Beta(6,2), Beta(0.5,6).



9. What do you call a sweet carbonated drink: pop, soda, coke, or something else? Let π be the proportion of U.S. residents that prefer the term "pop." Two different beverage salespeople from different regions of the country have different priors for π. The first salesperson works in North Dakota and specifies a Beta(8,2) prior. The second works in Louisiana and specifies a Beta(1,20) prior.

    a. Calculate the prior mean, mode, and standard deviation of π for both salespeople.
    b. Plot the prior PDFs for both salespeople.
    c. Compare, in words, the salespeople's prior beliefs about the proportion of U.S. residents who say "pop."

10. A university wants to know what proportion of students are regular bike riders, π, so that they can install an appropriate number of bike racks. Since the university is in sunny Texas, staff think that π has a mean of 1 in 4 students, and a mode of $5/22$.

    a. Specify and plot a Beta model that reflects the staff's prior ideas about π.
    b. Among 50 surveyed students, 15 are regular bike riders. What is the posterior model for π?
    c. What are the mean, mode, and standard deviation of the posterior model?
    d. Does the posterior model more closely reflect the prior information or the data? Explain your reasoning.