Discussion Problems - Week 3

Instructor: Dr. Paul Parker
TA: Antonio Aguirre

Winter 2025

University of California, Santa Cruz
Winter 2025

Problem 1

Suppose that the proportion θ of defective items in a large shipment is unknown and that the prior distribution of θ is the beta distribution with parameters $\alpha = 2$ and $\beta = 200$. If 100 items are selected at random from the shipment and if three of these items are found to be defective, what is the posterior distribution of θ?

Solution

We are given:

- Prior distribution: $\theta \sim \text{Beta}(2, 200)$.
- Observed data: 100 items sampled, 3 defective.

We aim to determine the posterior distribution of θ given this information.

Step 1: Posterior Distribution Formula

The Beta distribution is a conjugate prior for the binomial likelihood. Given a Beta prior $\theta \sim \text{Beta}(\alpha, \beta)$ and binomial likelihood $X \sim \text{Bin}(n, \theta)$ with x observed successes, the posterior is:

$$\theta \mid X \sim \text{Beta}(\alpha + x, \beta + n - x) \tag{1}$$

Step 2: Compute Posterior Parameters

Using our given values:

$$\begin{aligned}
\alpha' &= \alpha + x = 2 + 3 = 5, \\
\beta' &= \beta + n - x = 200 + 100 - 3 = 297.
\end{aligned}$$

Thus, the posterior distribution is:

$$\theta \mid X \sim \text{Beta}(5, 297). \tag{2}$$

Step 3: Interpretation

The posterior mean is given by:

$$\mathbb{E}[\theta \mid X] = \frac{\alpha'}{\alpha' + \beta'} = \frac{5}{5 + 297} = \frac{5}{302} \approx 0.0166. \tag{3}$$

This means our best estimate of the proportion of defective items after observing the sample is about 1.66%.

Final Answer

Posterior Distribution

$$\theta \mid X \sim \text{Beta}(5, 297). \tag{4}$$

Problem 2

Suppose that a random sample of 100 observations is to be taken from a normal distribution for which the value of the mean θ is unknown and the standard deviation is 2, and the prior distribution of θ is a normal distribution. Show that no matter how large the standard deviation of the prior distribution is, the standard deviation of the posterior distribution will be less than $1/5$.

Solution

We are given:

- A normal likelihood function with unknown mean θ and known standard deviation $\sigma = 2$.
- A normal prior distribution for θ.

Step 1: Define the Prior and Likelihood

Let the prior distribution of θ be:

$$\theta \sim \mathcal{N}(\mu_0, \tau^2), \tag{5}$$

where μ_0 is the prior mean and τ^2 is the prior variance.

Given a random sample of size $n = 100$ from $\mathcal{N}(\theta, \sigma^2)$, the likelihood function follows:

$$X_i \mid \theta \sim \mathcal{N}(\theta, \sigma^2), \quad i = 1, \ldots, 100. \tag{6}$$

Since the sample consists of independent observations, the sample mean $\bar{X}$ follows:

$$\bar{X} \mid \theta \sim \mathcal{N}(\theta, \frac{\sigma^2}{n}). \tag{7}$$

Step 2: Compute the Posterior Variance

From Bayesian updating with a normal prior and normal likelihood, the posterior variance $\sigma^2_{\theta|X}$ is given by:

$$\frac{1}{\sigma^2_{\theta|X}} = \frac{1}{\tau^2} + \frac{n}{\sigma^2}. \tag{8}$$

Substituting $\sigma^2 = 2^2 = 4$ and $n = 100$:

$$\frac{1}{\sigma^2_{\theta|X}} = \frac{1}{\tau^2} + \frac{100}{4} = \frac{1}{\tau^2} + 25. \tag{9}$$

Thus, the posterior standard deviation is:

$$\sigma_{\theta|X} = \left(\frac{1}{\frac{1}{\tau^2} + 25} \right)^{1/2}. \tag{10}$$

Step 3: Show the Standard Deviation Bound

Now, consider the case where $\tau^2 \to \infty$, meaning the prior provides minimal information about θ. In this case:

$$\frac{1}{\tau^2} \to 0, \tag{11}$$

so the posterior variance simplifies to:

$$\sigma^2_{\theta|X} = \frac{1}{25}. \tag{12}$$

Taking the square root:

$$\sigma_{\theta|X} = \frac{1}{5}. \tag{13}$$

Thus, regardless of how large the prior variance is, the posterior standard deviation will always be less than or equal to 1/5.

Final Answer

Conclusion

No matter how large the standard deviation of the prior distribution is, the posterior standard deviation of θ will always be less than or equal to 1/5.

Problem 3

Suppose that the time in minutes required to serve a customer at a certain facility has an exponential distribution for which the value of the parameter θ is unknown and that the prior distribution of θ is a gamma distribution for which the mean is 0.2 and the standard deviation is 1. If the average time required to serve a random sample of 20 customers is observed to be 3.8 minutes, what is the posterior distribution of θ?

Solution

We are given that the time required to serve a customer follows an **exponential distribution** with an unknown rate parameter θ. The prior distribution for θ is a **gamma distribution** with:

- **Mean:** $\mathbb{E}[\theta] = 0.2$
- **Standard Deviation:** $\text{SD}(\theta) = 1$

Additionally, a random sample of 20 customers was observed, and the average service time was found to be 3.8 minutes. We aim to determine the **posterior distribution** of θ.

Step 1: Parameterizing the Prior Distribution

A **gamma distribution** $\theta \sim \text{Gamma}(\alpha, \beta)$ has the following properties:

$$\mathbb{E}[\theta] = \frac{\alpha}{\beta}, \quad \text{Var}(\theta) = \frac{\alpha}{\beta^2}, \quad \text{SD}(\theta) = \frac{\sqrt{\alpha}}{\beta} \tag{14}$$

Using the given mean and standard deviation, we solve for α and β:

$$\frac{\alpha}{\beta} = 0.2, \quad \frac{\sqrt{\alpha}}{\beta} = 1 \tag{15}$$

Rewriting β in terms of α:

$$\beta = \frac{\sqrt{\alpha}}{1} = \sqrt{\alpha} \tag{16}$$

Substituting into the first equation:

$$\frac{\alpha}{\sqrt{\alpha}} = 0.2 \quad \Rightarrow \quad \sqrt{\alpha} = 0.2 \quad \Rightarrow \quad \alpha = 0.04, \quad \beta = 0.2 \tag{17}$$

Thus, the prior distribution is:

$$\theta \sim \text{Gamma}(0.04, 0.2) \tag{18}$$

Step 2: Computing the Posterior Distribution

For an exponential likelihood with a gamma prior, the posterior distribution follows:

$$\theta \mid \mathbf{x} \sim \text{Gamma}\left(\alpha + n, \beta + n\bar{x}\right) \tag{19}$$

where:

- $\alpha = 0.04$ (from prior)
- $\beta = 0.2$ (from prior)
- $n = 20$ (sample size)
- $\bar{x} = 3.8$ (sample mean)

Substituting these values:

$$\theta \mid \mathbf{x} \sim \text{Gamma}\left(0.04 + 20, 5 + 0.2(3.8)\right) \tag{20}$$

$$\theta \mid \mathbf{x} \sim \text{Gamma}\left(20.4, 5.76\right) \tag{21}$$

Final Answer

Posterior Distribution of θ

$$\theta \mid \mathbf{x} \sim \text{Gamma}(20.4, 5.76) \tag{22}$$

This is the required posterior distribution for θ.

Problem 4

Suppose that a random sample of size n is taken from the Bernoulli distribution with parameter θ, which is unknown, and that the prior distribution of θ is a beta distribution for which the mean is μ_0. Show that the mean of the posterior distribution of θ will be a weighted average having the form $\lambda_n \bar{X}_n + (1 - \lambda_n)\mu_0$, and show that $\lambda_n \to 1$ as $n \to \infty$.

Solution

We are given:

- A **random sample** of size n from a **Bernoulli distribution** with parameter θ.
- The **prior distribution** of θ is a **Beta distribution** with mean μ_0.

Step 1: Properties of the Beta Distribution

The Beta distribution $\mathbf{Beta}(\alpha, \beta)$ has mean:

$$\mathbb{E}[\theta] = \frac{\alpha}{\alpha + \beta}$$

If the prior for θ follows a $\mathbf{Beta}(\alpha_0, \beta_0)$ distribution, then:

$$\mu_0 = \frac{\alpha_0}{\alpha_0 + \beta_0}$$

Step 2: Posterior Distribution

Given a random sample $X_1, X_2, \dots, X_n \sim \mathbf{Bernoulli}(\theta)$, the **likelihood function** is:

$$L(\theta) = \prod_{i=1}^{n} \theta^{X_i}(1-\theta)^{1-X_i}$$

$$= \theta^{\sum X_i}(1-\theta)^{n-\sum X_i}$$

Using a $\mathbf{Beta}(\alpha_0, \beta_0)$ prior, the **posterior distribution** is:

$$\theta \mid X_1, \dots, X_n \sim \mathbf{Beta}(\alpha_0 + S_n, \beta_0 + n - S_n)$$

where:

$$S_n = \sum_{i=1}^{n} X_i \quad \textbf{(the sum of successes in the sample)}.$$

Step 3: Compute Posterior Mean

The **mean** of a **Beta**(α, β) distribution is:

$$\mathbb{E}[\theta \mid X] = \frac{\alpha_0 + S_n}{\alpha_0 + \beta_0 + n}$$

Rewriting in terms of **sample mean** $\bar{X}_n = \frac{S_n}{n}$:

$$\mathbb{E}[\theta \mid X] = \frac{\alpha_0 + n\bar{X}_n}{\alpha_0 + \beta_0 + n}$$

Factorizing:

$$\mathbb{E}[\theta \mid X] = \lambda_n \bar{X}_n + (1 - \lambda_n)\mu_0$$

where:

$$\lambda_n = \frac{n}{\alpha_0 + \beta_0 + n}.$$

Step 4: Show that $\lambda_n \to 1$ as $n \to \infty$

As $n \to \infty$, we analyze λ_n:

$$\lambda_n = \frac{n}{\alpha_0 + \beta_0 + n} \to 1 \quad \textbf{(since } n \textbf{ dominates the denominator for large } n\textbf{)}.$$

Thus, as the sample size increases, the posterior mean is increasingly dominated by the sample mean $\bar{X}_n$, meaning that:

$$\mathbb{E}[\theta \mid X] \to \bar{X}_n.$$

Final Answer

Posterior Mean of θ

$$\mathbb{E}[\theta \mid X] = \lambda_n \bar{X}_n + (1 - \lambda_n)\mu_0, \quad \textbf{where} \quad \lambda_n = \frac{n}{\alpha_0 + \beta_0 + n}.$$

$$\lambda_n \to 1 \quad \textbf{as} \quad n \to \infty.$$

This shows that for **large sample sizes**, the influence of the prior diminishes, and the posterior mean approaches the sample mean $\bar{X}_n$, demonstrating the **Bayesian learning effect**.

Problem 5

Suppose that the number of defects in a 1200-foot roll of magnetic recording tape has a Poisson distribution for which the value of the mean θ is unknown, and the prior distribution of θ is the gamma distribution with parameters $\alpha = 3$ and $\beta = 1$. When five rolls of this tape are selected at random and inspected, the numbers of defects found on the rolls are 2, 2, 6, 0, and 3. Using squared error loss, what is the Bayes estimate of θ?

Solution

We are given:

- The number of defects follows a **Poisson**(θ) distribution.
- The prior distribution of θ follows a **Gamma**(α, β) with parameters $\alpha = 3$ and $\beta = 1$.
- A sample of five rolls is observed with defect counts: $2, 2, 6, 0, 3$.
- The loss function used is the **squared error loss**, implying that the Bayes estimator is the posterior mean.

Step 1: Posterior Distribution

The likelihood function for a Poisson-distributed variable $X_i \sim \text{Poisson}(\theta)$ is:

$$L(\theta) = \prod_{i=1}^{n} \frac{\theta^{x_i} e^{-\theta}}{x_i!}$$

For a sample of size n, the joint likelihood is:

$$L(\theta) = \frac{\theta^{\sum x_i} e^{-n\theta}}{\prod x_i!}.$$

The prior is given by:

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

Using Bayes' theorem, the posterior distribution is proportional to:

$$p(\theta \mid \mathbf{x}) \propto L(\theta) p(\theta).$$

Substituting the Poisson likelihood and gamma prior:

$$p(\theta \mid \mathbf{x}) \propto \theta^{\sum x_i} e^{-n\theta} \times \theta^{\alpha-1} e^{-\beta\theta}.$$

Rewriting:

$$p(\theta \mid \mathbf{x}) \propto \theta^{\alpha+\sum x_i - 1} e^{-(\beta+n)\theta}.$$

This is the kernel of a **Gamma**(α', β') distribution with:

$$\alpha' = \alpha + \sum x_i, \quad \beta' = \beta + n.$$

Step 2: Compute Posterior Parameters

Summing the observed defects:

$$\sum x_i = 2 + 2 + 6 + 0 + 3 = 13.$$

Thus, the updated parameters are:

$$\alpha' = 3 + 13 = 16, \quad \beta' = 1 + 5 = 6.$$

Step 3: Compute Bayes Estimate

Under squared error loss, the Bayes estimate of θ is the **posterior mean**:

$$\hat{\theta}_{\text{Bayes}} = \frac{\alpha'}{\beta'} = \frac{16}{6} = \frac{8}{3} \approx 2.67.$$

Final Answer

Bayes Estimate of θ

$$\hat{\theta}_{\text{Bayes}} = \frac{8}{3} \approx 2.67.$$

This is the Bayes estimate of θ under squared error loss.

Conjugate Priors and Posteriors

The table below highlights common conjugate prior-likelihood relationships, showing how the posterior maintains the same distributional form as the prior, making Bayesian updating computationally convenient.

Prior	Likelihood	Posterior
Beta(α, β)	**Binomial**(n, p)	**Beta**$(\alpha + k, \beta + n - k)$
Normal(μ_0, σ_0^2)	**Normal**(μ, σ^2)	**Normal**$\left(\frac{\frac{\mu_0}{\sigma_0^2} + \frac{k}{\sigma^2}}{\frac{1}{\sigma_0^2} + \frac{n}{\sigma^2}}, \frac{1}{\frac{1}{\sigma_0^2} + \frac{n}{\sigma^2}}\right)$
Gamma(α, β)	**Exponential**(λ)	**Gamma**$(\alpha + 1, \beta + x)$
Gamma(α, β)	**Poisson**(λ)	**Gamma**$(\alpha + \sum x_i, \beta + n)$

Table 1: Conjugate Priors and Their Corresponding Posteriors

Observations:

- The **Beta-Binomial** model updates the shape parameters by adding successes and failures.
- The **Normal-Normal** model updates the mean using a precision-weighted average.
- The **Gamma-Exponential** model increments the shape parameter and scales the rate by adding observed data.
- The **Gamma-Poisson** model increments the shape parameter by the sum of observed counts while updating the rate.

Practice for midterm: Show that given each prior and likelihood, those posteriors follow!