Problem Sheet 4: Bayesian Inference, Tests, and Decision Theory

1. **HPD Regions**

   Let Q be a probability distribution on Θ with density g with respect to a measure ν. We begin by defining the level sets for Q. For all $y \geq 0$, we define:

   $$\mathcal{L}(y) = \{\theta \in \Theta : g(\theta) \geq y\}.$$

   The region $\mathcal{L}(y)$ consists of the set of parameters for which the density g exceeds the level y.

   Let $\alpha \in (0, 1)$. The Highest Posterior Density (HPD) region at level $1 - \alpha$ for a distribution Q of density g is the region $\mathcal{H} \subset \Theta$ given by:

   $$\mathcal{H} = \mathcal{L}(y_\alpha),$$

   where
   $$y_\alpha = \sup\{y \in \mathbb{R}_+ : Q(\mathcal{L}(y)) \geq 1 - \alpha\}.$$

   Note that since $\alpha < 1$, we have $y_\alpha < +\infty$.

   (a) If a posterior distribution on $\mathbb{R}$ has a continuous, symmetric density that is strictly increasing on $\mathbb{R}^-$ and strictly decreasing on $\mathbb{R}^+$, show that the HPD regions of level $1-\alpha$ coincide with the intervals defined by the quantiles $\alpha/2$ and $1 - \alpha/2$ of the posterior density.

   (b) Give an example of a posterior density for which the two types of $1 - \alpha$ credible regions from the previous question do not coincide.

2. **Credible Intervals and Confidence Intervals**

   Let $X = (X_1, \dots, X_n)$ with $X_i \sim \text{Bernoulli}(\theta)$ i.i.d. We place a prior $\text{Beta}(a, b)$ on θ, with $a > 0$ and $b > 0$. We give:

   $$\mathbb{E}[\text{Beta}(a, b)] = \frac{a}{a+b}, \quad \text{Var}[\text{Beta}(a, b)] = \frac{ab}{(a+b)^2(a+b+1)}.$$

   (a) Determine the posterior distribution $\Pi[\cdot|X]$. We will denote its mean by m_X and its variance by v_X.

   (b) Construct a credible interval $I^T(X)$ of level at least $1 - \alpha$ (with $\alpha > 0$), centered at m_X, using Chebyshev's inequality.

   (c) We ask whether $I^T(X)$ can be used as an asymptotic confidence interval, in the frequentist sense under P_{θ_0}. Answer this question by seeking an asymptotic lower bound for the level of $I^T(X)$ as a function of α.

3. **Bayesian Testing**

   (a) **Test I:** Let $X = (X_1, \dots, X_n)|\theta \sim \mathcal{N}(\theta, \sigma^2)^{\otimes n}$ and $\theta \sim \Pi = \mathcal{N}(\mu, \tau^2)$, where σ^2, τ^2 are fixed.

      i. Determine the posterior distribution.

      ii. We want to test $H_0 = \{\theta \geq 1\}$ against $H_1 = \{\theta < 1\}$ from a Bayesian perspective. For a test $\varphi = \varphi(X)$ and $\theta \in \mathbb{R}$, we consider the balanced loss function:

$$\ell(\theta, \varphi) = 1_{\theta \in \Theta_0} 1_{\varphi = 1} + 1_{\theta \in \Theta_1} 1_{\varphi = 0}.$$

      Construct the corresponding Bayesian test for the prior Π defined above.

      iii. What does the test become if we replace H_0 with H_1 and vice-versa?

   (b) **Test II:** Let $X = X_1|\theta \sim \mathcal{N}(\theta, 1)$. Consider the two testing problems:

$$H_0^1 : \theta = 0 \quad \text{vs.} \quad H_1^1 : \theta \neq 0$$
$$H_0^2 : |\theta| \leq \epsilon \quad \text{vs.} \quad H_1^2 : |\theta| > \epsilon$$

      i. Propose a prior distribution with a Gaussian part $\mathcal{N}(0, \sigma^2)$ for each situation.

      ii. Compare the corresponding Bayesian tests when ϵ and σ vary, in the case of a balanced loss function.

4. **Bayes and Constant Risk to Minimax**

   (a) Re-prove that a Bayes estimator with constant risk is minimax.

   (b) Let $\mathcal{P} = \{P_\theta = \text{Bin}(n, \theta), \theta \in (0, 1)\}$ and let $X|\theta \sim P_\theta$.

      i. Show that the family of priors $\{\Pi_{a,b} = \text{Beta}(a, b), a > 0, b > 0\}$ is conjugate for this model.

      ii. Give a Bayes estimator $\hat{\theta}_{a,b}(X)$ for $\Pi_{a,b}$ and the quadratic loss.

      iii. Assume $a = b$. Find a minimax estimator for the quadratic loss.

      iv. Is the estimator $T = X/n$ minimax?

5. **Bayes and Unique to Admissible**

   Let $\mathcal{P} = \{P_\theta, \theta \in \Theta \subset \mathbb{R}\}$ be a statistical model with $dP_\theta = f_\theta d\mu$ and let X be an observation following this model. Let T be an estimator of θ and $R_B(\Pi, T)$ its Bayes risk for a prior Π on Θ and the quadratic loss.

   (a) Give a Bayes estimator for Π. We will denote it T_1.

   (b) Let $m^\pi(x) = \int f_\theta(x) d\Pi(\theta)$. How is this quantity interpreted?

   (c) Show that for $T = T(X)$ an estimator of θ,

$$R_B(\Pi, T) = \int \mathbb{E}[(T(X) - \theta)^2 | X = x] m^\pi(x) d\mu(x).$$

   (d) Let T_2 be a Bayes estimator for Π and the quadratic loss, potentially different from T_1. Show that if the distribution $dQ = m^\pi d\mu$ dominates all distributions P_θ, then T_1 and T_2 are *equivalent*, in the sense where

$$R(\theta, T_1) = R(\theta, T_2) \quad \forall \theta \in \Theta.$$

(e) Show that if the Bayes estimator is unique up to equivalence, it is admissible.

(f) **Application:** Let $\mathcal{P} = \{P_\theta = \mathcal{N}(\theta, 1), \theta \in \mathbb{R}\}$ and $X_1, \ldots, X_n$ i.i.d. with law P_θ given θ. Set $\Pi = \mathcal{N}(a, \sigma^2)$, with $a \in \mathbb{R}$ and $\sigma^2 > 0$ fixed.

  i. Calculate the Bayes estimator for the prior Π and quadratic loss.

  ii. Determine the marginal distribution of $X = (X_1, \ldots, X_n)$, denoted Q_n. *[Hint: You may write X as a sum of two Gaussian vectors].*

  iii. Verify that Q_n dominates all laws $P_\theta^{\otimes n}$.

  iv. Show that the estimators $\alpha\overline{X} + \beta$, with $\alpha \in [0, 1)$ and $\beta \in \mathbb{R}$, are admissible.

  v. Show that the estimators $\overline{X} + \beta$, for $\beta \neq 0$, are not admissible.