PRINCIPLES OF STATISTICS Part II
Example Sheet 3 (of 4) Michaelmas 2025
By courtesy of Prof Rajen Shah

1. Consider the Bayesian model $X \mid \theta \sim \mathrm{Pois}(\theta)$, $\theta \in \Theta = [0, \infty)$, and suppose the prior for θ is a Gamma distribution with parameters α, λ. Show that the posterior distribution $\theta \mid X$ is also a Gamma distribution and find its parameters.

2. Suppose $X \mid \theta \sim \mathrm{Bin}(n, \theta)$ (where n is known) with $\theta \in \Theta = [0, 1]$.

   (a) Consider a $\mathrm{Beta}(a, b)$ prior for θ. Show that the posterior distribution is $\mathrm{Beta}(a + X, b + n - X)$ and compute the posterior mean $\bar{\theta}_n = \bar{\theta}_n(X)$.

   (b) Show that the maximum likelihood estimator for θ is *not* identical to the posterior mean with 'ignorant' uniform prior $\theta \sim U[0, 1]$.

   (c) Now suppose $X \sim \mathrm{Bin}(n, \theta_0)$ where $\theta_0 \in (0, 1)$ is deterministic. Derive the asymptotic distribution of $\sqrt{n}(\bar{\theta}_n - \theta_0)$.

3. Consider the Bayesian model $X_1, \ldots, X_n \mid \theta \sim N(\theta, 1)$ with prior π such that $\theta \sim N(\mu, v^2)$. Writing $\bar{\theta}_n$ for the posterior mean, for $0 < \alpha < 1$, consider the $(1-\alpha)$-level credible interval

$$\widehat{C}_n = \left\{ \theta \in \mathbb{R} : |\theta - \bar{\theta}_n| \leq R_n \right\}.$$

   Now suppose $X_1, \ldots, X_n \overset{\text{i.i.d.}}{\sim} N(\theta_0, 1)$ for a deterministic $\theta_0 \in \mathbb{R}$. Show that, as $n \to \infty$, $\mathbb{P}_{\theta_0}(\theta_0 \in \widehat{C}_n) \to 1 - \alpha$.

4. Consider estimation of $\theta \in \Theta = [0, 1]$ with data $X \sim \mathrm{Bin}(n, \theta)$ under quadratic risk.

   (a) Find the unique minimax estimator $\widetilde{\theta}_n$ of θ and deduce that the maximum likelihood estimator $\widehat{\theta}_n$ of θ is *not* minimax for any fixed sample size $n \in \mathbb{N}$.

   (b) Show, however, that

$$\lim_{n \to \infty} \frac{\sup_\theta R(\widehat{\theta}_n, \theta)}{\sup_\theta R(\widetilde{\theta}_n, \theta)} = 1$$

   and moreover that the maximum likelihood estimator dominates $\widetilde{\theta}_n$ in the large sample limit in the sense that

$$\lim_{n \to \infty} \frac{R(\widehat{\theta}_n, \theta)}{R(\widetilde{\theta}_n, \theta)} < 1 \quad \text{for all } \theta \in [0, 1],\ \theta \neq \frac{1}{2}.$$

5. Suppose $X_1, \ldots, X_n \overset{\text{i.i.d.}}{\sim} N(\mu, \sigma^2)$.

   (a) Suppose $(\mu, \sigma^2) \in \Theta = \mathbb{R} \times [0, v]$ for some $v > 0$. Show that the sample mean $\bar{X}_n$ is minimax for the risk $R(\bar{X}_n, (\mu, \sigma^2)) = \mathbb{E}[(\bar{X}_n - \mu)^2]$.

   (b) Now suppose it is known that $\sigma^2 = 1$ but $\mu \in \Theta = [0, \infty)$ is unknown. Show that the sample mean $\bar{X}_n$ is inadmissible for quadratic risk, but that it is still minimax. What happens if $\Theta = [a, b]$ for some $0 < a < b < \infty$?

6. Consider a Bayesian version of the normal linear model where $Y \mid \beta \sim N_n(X\beta, I)$, $X \in \mathbb{R}^{n \times p}$ is a fixed matrix of predictors (not necessarily with full column rank) and β has prior $\beta \sim N_p(0, \lambda^{-1} I)$ for a fixed $\lambda > 0$. Find the posterior mean of β.

7. Consider the Bayesian model $X \mid \theta \sim N_p(\theta, I)$ where $p \geq 3$ and $\theta \in \mathbb{R}^p$ has prior $\theta \sim N_p(0, \tau^2 I)$ and τ^2 is deterministic.

   (a) Suppose first that τ^2 is known. Show that the posterior mean $\bar{\theta}$ is given by

   $$\bar{\theta}(X) := (1-\gamma)\, X$$

   where $\gamma := (\tau^2+1)^{-1}$.

   (b) Now suppose τ^2 is unknown. Find the marginal distribution of X (as a function of γ) and show that

   $$\widehat{\gamma} := \frac{p-2}{\|X\|^2}$$

   satisfies $\mathbb{E}_\gamma(\widehat{\gamma}) = \gamma$. [*Hint: If* $Z \sim \chi^2_p$ *then* $\mathbb{E}(Z^{-1}) = (p-2)^{-1}$.] What does this have to do with the James–Stein estimator?

8. Let $X \sim N_p(\theta, I)$ with $p \geq 3$. Show that the risk of the James–Stein estimator $\widehat{\theta}_{\mathrm{JS}}$ satisfies

   $$R(\widehat{\theta}_{\mathrm{JS}}, \theta) \leq p - \frac{(p-2)^2}{p-2+\|\theta\|^2}.$$

   [*Hint: Let* $Z_1, Z_2, \ldots \stackrel{\text{i.i.d.}}{\sim} N(0,1)$. *If* $K \sim \mathrm{Pois}(\mu^2/2)$ *independently of the* Z_j*, then*

   $$(Z_1+\mu)^2 \qquad \text{and} \qquad \sum_{j=1}^{1+2K} Z_j^2$$

   have the same distribution.]

9. Let $X \sim N_p(\theta, I)$ where $\theta \in \Theta = \mathbb{R}^p, p \geq 3$. Consider estimators

   $$\widetilde{\theta}^{(c)} = \left(1 - c\frac{p-2}{\|X\|^2}\right) X, \ \ 0 < c < 2,$$

   for θ, under the risk function $R(\delta, \theta) = \mathbb{E}_\theta \|\delta(X) - \theta\|^2$.

   (a) Show that the James–Stein estimator $\widetilde{\theta}^{(1)}$ dominates all estimators $\widetilde{\theta}^{(c)}, c \neq 1$.

   (b) Let $\widehat{\theta}$ be the maximum likelihood estimator of θ. Show that, for any $0 < c < 2$,

   $$\sup_{\theta\in\Theta} R(\widetilde{\theta}^{(c)}, \theta) = \sup_{\theta\in\Theta} R(\widehat{\theta}, \theta).$$

10. Consider $X_1, \ldots, X_n \stackrel{\text{i.i.d.}}{\sim} N(\theta, 1)$ with $\theta \in \Theta = \mathbb{R}$. The *Hodges' estimator*

    $$\widetilde{\theta}_n := \bar{X}_n \mathbb{1}_{\{|\bar{X}_n| \geq n^{-1/4}\}},$$

    is equal to the maximum likelihood estimator $\bar{X}_n$ of θ whenever $|\bar{X}_n| \geq n^{-1/4}$ and is zero otherwise.

    (a) Find the asymptotic distribution of $\sqrt{n}(\widetilde{\theta}_n - \theta)$ for each $\theta \in \mathbb{R}$ and show moreover that when $\theta = 0$,

    $$\lim_{n\to\infty} n\mathbb{E}_\theta[(\widetilde{\theta}_n - \theta)^2] = 0.$$

    (b) Show however that

    $$\lim_n \sup_{\theta\in\Theta} n\mathbb{E}_\theta[(\widetilde{\theta}_n - \theta)^2] = \infty.$$

11. (i) Let ϕ and Φ denote the standard Gaussian pdf and cdf respectively. If $Z \sim N(\mu, \sigma^2)$, then

$$\mathbb{E}[Z \mid Z \in (a,b)] = \mu + \frac{\phi(\alpha) - \phi(\beta)}{\Phi(\beta) - \Phi(\alpha)}\sigma$$

where

$$\alpha := \frac{a-\mu}{\sigma} \qquad \text{and} \qquad \beta := \frac{b-\mu}{\sigma}.$$

[You need not show this.] Suppose now that $\zeta \sim N(\mu, 1)$ and $a < 0 < b$. Explain why

$$\mu - \frac{\phi(\mu)}{\Phi(-\mu)} \leq \mathbb{E}[\zeta \mid \zeta \in (a,b)] \leq \mu + \frac{\phi(\mu)}{\Phi(\mu)}.$$

[*Hint: Use the fact that $x \mapsto \phi(x)/\Phi(x)$ is decreasing.*]

(ii) In the setting of Question 10 show that the maximum likelihood estimator is "only asymptotically beatable on arbitrarily small sets of θ-values" in the following sense: given $a < b$, any sequence of estimators $\widehat{\theta}_n := \widehat{\theta}_n(X_1, \ldots, X_n)$ has

$$\liminf_{n\to\infty} \sup_{\theta\in(a,b)} n\mathbb{E}_\theta[(\widehat{\theta}_n - \theta)^2] \geq 1.$$

[*Hint: Consider a π-Bayes estimator for an appropriate prior π. You may find the fact that $\int_{-\infty}^{\infty} \phi(x)^3/\Phi(x)^2\,dx < \infty$ useful.*]