A1/12 B1/15 **Principles of Statistics**

(i) What are the main approaches by which prior distributions are specified in Bayesian inference?

Define the risk function of a decision rule d. Given a prior distribution, define what is meant by a Bayes decision rule and explain how this is obtained from the posterior distribution.

(ii) Dashing late into King's Cross, I discover that Harry must have already boarded the Hogwarts Express. I must therefore make my own way onto platform nine and three-quarters. Unusually, there are two guards on duty, and I will ask one of them for directions. It is safe to assume that one guard is a Wizard, who will certainly be able to direct me, and the other a Muggle, who will certainly not. But which is which? Before choosing one of them to ask for directions to platform nine and three-quarters, I have just enough time to ask one of them "Are you a Wizard?", and on the basis of their answer I must make my choice of which guard to ask for directions. I know that a Wizard will answer this question truthfully, but that a Muggle will, with probability $\frac{1}{3}$, answer it untruthfully.

Failure to catch the Hogwarts Express results in a loss which I measure as 1000 galleons, there being no loss associated with catching up with Harry on the train.

Write down an exhaustive set of non-randomised decision rules for my problem and, by drawing the associated risk set, determine my minimax decision rule.

My prior probability is $\frac{2}{3}$ that the guard I ask "Are you a Wizard?" is indeed a Wizard. What is my Bayes decision rule?

A2/11 B2/16 **Principles of Statistics**

(i) Let $X_1, \ldots, X_n$ be independent, identically-distributed $N(\mu, \mu^2)$ random variables, $\mu > 0$.

Find a minimal sufficient statistic for μ.

Let $T_1 = n^{-1} \sum_{i=1}^n X_i$ and $T_2 = \sqrt{n^{-1} \sum_{i=1}^n X_i^2}$. Write down the distribution of X_i/μ, and hence show that $Z = T_1/T_2$ is ancillary. Explain briefly why the Conditionality Principle would lead to inference about μ being drawn from the conditional distribution of T_2 given Z.

What is the maximum likelihood estimator of μ?

(ii) Describe briefly the Bayesian approach to predictive inference.

Let $Z_1, \ldots, Z_n$ be independent, identically-distributed $N(\mu, \sigma^2)$ random variables, with μ, σ^2 both unknown. Derive the maximum likelihood estimators $\widehat{\mu}, \widehat{\sigma}^2$ of μ, σ^2 based on $Z_1, \ldots, Z_n$, and state, without proof, their joint distribution.

Suppose that it is required to construct a prediction interval $I_{1-\alpha} \equiv I_{1-\alpha}(Z_1, \ldots, Z_n)$ for a future, independent, random variable Z_0 with the same $N(\mu, \sigma^2)$ distribution, such that

$$P(Z_0 \in I_{1-\alpha}) = 1 - \alpha,$$

with the probability over the *joint* distribution of $Z_0, Z_1, \ldots, Z_n$. Let

$$I_{1-\alpha}(Z_1, \ldots, Z_n; \sigma^2) = \left[\bar{Z}_n - z_{\alpha/2}\sigma\sqrt{1+1/n},\ \bar{Z}_n + z_{\alpha/2}\sigma\sqrt{1+1/n}\right],$$

where $\bar{Z}_n = n^{-1} \sum_{i=1}^n Z_i$, and $\Phi(z_\beta) = 1 - \beta$, with Φ the distribution function of $N(0,1)$.

Show that $P(Z_0 \in I_{1-\alpha}(Z_1, \ldots, Z_n; \sigma^2)) = 1 - \alpha$.

By considering the distribution of $(Z_0 - \bar{Z}_n)/\left(\widehat{\sigma}\sqrt{\frac{n+1}{n-1}}\right)$, or otherwise, show that

$$P(Z_0 \in I_{1-\alpha}(Z_1, \ldots, Z_n; \widehat{\sigma}^2)) < 1 - \alpha,$$

and show how to construct an interval $I_{1-\gamma}(Z_1, \ldots, Z_n; \widehat{\sigma}^2)$ with

$$P(Z_0 \in I_{1-\gamma}(Z_1, \ldots, Z_n; \widehat{\sigma}^2)) = 1 - \alpha.$$

[*Hint: if Y has the t-distribution with m degrees of freedom and $t_\beta^{(m)}$ is defined by $P(Y < t_\beta^{(m)}) = 1 - \beta$ then $t_\beta > z_\beta$ for $\beta < \frac{1}{2}$.*]

A3/12 B3/15 **Principles of Statistics**

(i) Explain what is meant by a *uniformly most powerful unbiased test* of a null hypothesis against an alternative.

Let $Y_1, \ldots, Y_n$ be independent, identically distributed $N(\mu, \sigma^2)$ random variables, with σ^2 known. Explain how to construct a uniformly most powerful unbiased size α test of the null hypothesis that $\mu = 0$ against the alternative that $\mu \neq 0$.

(ii) Outline briefly the Bayesian approach to hypothesis testing based on *Bayes factors*.

Let the distribution of $Y_1, \ldots, Y_n$ be as in (i) above, and suppose we wish to test, as in (i), $\mu = 0$ against the alternative $\mu \neq 0$. Suppose we assume a $N(0, \tau^2)$ prior for μ under the alternative. Find the form of the Bayes factor B, and show that, for fixed n, $B \to \infty$ as $\tau \to \infty$.

A4/13 B4/15 **Principles of Statistics**

Write an account, with appropriate examples, of **one** of the following:

(a) Inference in multi-parameter exponential families;

(b) Asymptotic properties of maximum-likelihood estimators and their use in hypothesis testing;

(c) Bootstrap inference.