Exam question for **Fundamental Theory of Statistical Inference**.

(i) What are the key characteristics of a *curved exponential family* model? What are the main statistical principles relevant to inference in such a family?

(ii) 'An ancillary statistic gives no information itself on a parameter of interest, but determines the precision with which it can be estimated'.
Discuss this statement, with detailed reference to the problem of inference for θ based on an independent sample $X_1, \ldots, X_n$ from a uniform distribution over $(\theta - 1/2, \theta + 1/2)$.

(iii) Suppose that, given parameter θ, X is normally distributed $N(\theta, 1)$, and consider estimation of θ with $\Theta = \mathcal{A} = \mathbb{R}$ and loss function

$$L(\theta, a) = \exp(-\lambda a \theta),$$

where λ is a positive constant.
Determine the form of the Bayes rule, under a $N(0, \sigma^2)$ prior for θ. Deduce that the rule $d(x) = x/\lambda$ is minimax. What is the minimum variance unbiased estimator of θ?

(iv) Suppose that $X_1, \ldots, X_n$ are independent, with common $N(0, v)$ distribution, where $v > 0$ is assumed to have a prior density π. In a decision-theoretic approach to estimation of v, we take quadratic loss: $L(v, a) = (v - a)^2$. Write $X = (X_1, \ldots, X_n)$ and $|X| = (X_1^2 + \ldots + X_n^2)^{1/2}$.
By considering estimators of the form $\widehat{v}(X) = \alpha |X|^2$, show that if $\alpha \neq 1/(n+2)$ then the estimator $\widehat{v}(X) = \alpha |X|^2$ is not Bayes, for any choice of prior π.
By considering estimators of the form $\widehat{v}(X) = \alpha |X|^2 + \beta$, show that if $\alpha \neq 1/n$, then the estimator $\widehat{v}(X) = \alpha |X|^2$ is not Bayes, for any choice of prior π.

(v) Show that the normal distribution, $N(\mu, \sigma^2)$, of mean μ and variance σ^2 constitutes *both* a full exponential family *and* a transformation model.
Suppose it is required, on the basis of an independent sample $X_1, \ldots, X_n$ from this distribution, to test $H_0 : \mu \leq \mu_0$ against $H_1 : \mu > \mu_0$. Explain what the full exponential family form of the distribution implies about the form of the optimal test. Discuss carefully the relationship between the optimal test and the elementary $t-$test routinely used in this inference problem.