1. Suppose $Y_1, \ldots, Y_n$ are i.i.d. P where $\mathrm{E}[Y_i] = \mu$, $\mathrm{Var}[Y_i] = \sigma^2$ and $\mathrm{E}[Y_i^4] = \gamma < \infty$. Show that the usual sample variance $S_n^2 = \sum_{i=1}^n (Y_i - \bar{Y})^2/(n-1)$ is a CAN estimator of σ^2, and obtain the asymptotic variance (Hint: There are several approaches to this problem. One way it to first consider $S_0^2 = \sum(Y_i - \mu)^2/n$, then $S_1^2 = \sum(Y_i - \bar{Y})^2/n$, then S^2).

2. Let $Y_1, \ldots, Y_n$ be independent random variables with $Y_i \sim N(\theta x_i, \sigma^2)$, where x_i are known positive numbers for each i.

   (a) Let $\hat{\theta} = \bar{W}$ where $W_i = Y_i/x_i$. Find the bias and variance of $\hat{\theta}$.

   (b) Find the MLE $\hat{\theta}_{MLE}$ of θ and compute its bias and variance.

   (c) Compare the two estimators in terms of MSE. Under what conditions is one of them better than the other?

   (d) Now suppose the x_i's will be simulated by the experimenters from a distribution with mean μ_X and variance σ_X^2. What sort of values of μ_X and σ_X^2 would you recommend, and why?

3. Consider a p-parameter exponential family model, $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$, with $\Theta \subset \mathbb{R}^p$ and densities $p_\theta(y) = c(y)\exp(\theta^\top t(y) - A(\theta))$. Here, $t(y)$ is a vector-valued function of the data value y.

   (a) Show that the following models are exponential families, and identify $c(y)$, θ, $t(y)$ and $A(\theta)$ in each:

      i. The binomial(n, p) distributions, with $p \in [0, 1]$.

      ii. The Poisson(μ) distributions, with $\mu \in \mathbb{R}^+$.

      iii. The $N(\mu, \sigma^2)$ distributions, with $\mu \in \mathbb{R}$ and $\sigma^2 \in \mathbb{R}^+$.

      iv. The beta(a, b) distributions, with $a > 0$ and $b > 0$.

   (b) For a sample of size n, write out the log-likelihood function of an exponential family model and simplify as much as possible.

   (c) Find the likelihood equation, that is, the equation that sets the derivative of the log-likelihood function equal to zero, with the data on one side of the equation and items involving the MLE on the other.

(d) Now take the derivative of $p(y|\theta)$ with respect to θ and integrate to obtain a formula for the expectation of $t(y)$. Compare this equation to the one obtained on (b), and comment.

4. Let $Y_1, \ldots, Y_n \sim$ i.i.d. binary(θ).

   (a) Find the MLE of θ.

   (b) Rewrite the model in terms of $Y_1, \ldots, Y_n \sim$ i.i.d. binary$(e^{\psi}/(1 + e^{\psi}))$. Find the MLE of ψ, and relate it to the MLE of θ from (a).

5. Let $\mathcal{P}_1 = \{f_\theta(y) : \theta \in \Theta\}$ denote the densities of a model. Let $\mathcal{P}_2$ be a *reparameterization* of this model, that is, $\mathcal{P}_2 = \{g_\psi(y) : \psi \in \Psi\}$, where for every $\theta \in \Theta$ there is a unique $\psi \in \Psi$ such that $f_\theta(y) = g_\psi(y)$ for all y. Note that this means there is a 1-1 function $h : \Theta \to \Psi$ such that $f_\theta(y) = g_{h(\theta)}(y)$ for all y. Find the relationship between the MLE based on $\mathcal{P}_1$ and the MLE based on $\mathcal{P}_2$.