1. The usual arithmetic mean $\bar{y}$ of a sample $(y_1, \dots, y_n)$ is highly sensitive to outliers, so sometimes we analyze $(x_1, \dots, x_n) = (\ln y_1, \dots, \ln y_n)$, or compute the sample mean on this scale.

   (a) Show that $e^{\bar{x}} \leq \bar{y}$.

   (b) Compare magnitudes of three different types of means given by $m(y_1, \dots, y_n) = f^{-1}(\sum f(y_i)/n)$ for $f(y) = 1/y$, $f(y) = \ln y$ and $f(y) = y$.

   (c) For each type of mean in (b), compute its sensitivity to outliers by approximating $m(y_1 + \delta, \dots, y_n)$ with $m(y_1, \dots, y_n) + \delta \times \frac{\partial}{\partial y_1} m(y_1, \dots, y_n)$. Interpret your results.

2. Let $Y_1, \dots, Y_n$ be real-valued random variables with $\mathrm{E}[Y_i] = \mu$ and $\mathrm{Var}[Y_i] = \sigma^2$ for all $i = 1, \dots, n$. Compute the variance of $\bar{Y} = \sum Y_i/n$ when

   (a) $\mathrm{Cor}[Y_i, Y_j] = 0$ for all $i.j$;

   (b) $\mathrm{Cor}[Y_i, Y_j] = \rho$ for all i, j;

   (c) $\mathrm{Cor}[Y_i, Y_j] = \rho$ if $|i - j| = 1$ and is zero if $|i - j| > 1$.

   For each of the three cases, describe what happens to $\Pr(|\bar{Y} - \mu| > \epsilon)$ as $n \to \infty$, that is, evaluate the consistency of $\bar{Y}$ for μ. Discuss your results.

3. Suppose $\mathrm{E}[Y] = \mu$ and $\mathrm{Var}[Y] = \sigma^2$. Consider the estimator $\hat{\mu} = (1 - w)\mu_0 + wY$, where $\mu_0 \neq 0$ and $w \in (0, 1)$ are numbers.

   (a) Find the expectation, variance, bias and MSE of $\hat{\mu}$ as functions of μ.

   (b) For what values of μ does $\hat{\mu}$ have lower MSE than Y? Interpret your results.

4. Let $\hat{\theta}$ be an estimator for some unknown quantity θ. Derive a Chebyshev-like bound on $\Pr(|\hat{\theta} - \theta| > \epsilon)$ in terms of the MSE of $\hat{\theta}$.

5. Let $\bar{Y}_n \sim N(\mu, \sigma^2/n)$ and consider estimating μ with $\hat{\mu}_n = (1-w_n)\mu_0 + w_n\bar{Y}_n$, where μ_0 and w_n are constants, with w_n potentially depending on n.

   (a) Find minimal conditions on w_n that make $\hat{\mu}_n$ a consistent estimator of μ. In other words, find conditions on w_n so that $\Pr(|\hat{\mu}_n - \mu| > \epsilon) \to 0$ as $n \to \infty$ no matter what μ is.

   (b) Now consider a Bayes estimator of μ based on the prior distribution $\mu \sim N(\mu_0, \tau^2)$, and the model that conditional on μ, $\bar{Y}_n \sim N(\mu, \sigma^2/n)$ (so $p(\bar{y}_n|\mu)$ is a $N(\mu, \sigma^2/n)$ density and $p(\mu)$ is a $N(\mu_0, \tau^2)$ density).

      i. Write out the joint density of μ and $\bar{Y}_n$, and by re-arranging terms show that the conditional distribution of μ given $\bar{Y}_n$ is a normal distribution. Specify the expectation and variance of this normal distribution.

      ii. If σ^2, μ, τ^2 are known then the conditional expectation of μ given $\bar{Y}_n$ is some function of $\bar{Y}_n$. Suppose we use this function as an estimator of μ. Using your results from item 2 above, decide if this is a consistent estimator of μ, and explain your answer.