Note: Include your **R** code and relevant output as part of your submission!

1. Let $X_1, X_2, \ldots, X_n$ be IID $N(\mu, \sigma^2)$ random variables. Let $\hat{\sigma}^2 = \sum(X_i - \bar{X})^2/n$ and $s^2 = \sum(X_i - \bar{X})^2/(n-1)$

   (a) Compute the expected value and variance for $\hat{\sigma}^2$ and s^2. Thus or otherwise, compute the Mean Squared Error (MSE) for each.

   (b) Plot the ratio $MSE(\hat{\sigma}^2)/MSE(s^2)$ against n for `n <- 2:25`.

2. Let X be a Binomial(n, p) random variable. Let $\hat{p} = X/n$, and $\check{p} = (X+2)/(n+4)$. Note: $\check{p}$ would be the posterior mean if you chose a Beta$(2, 2)$ prior. To a Frequentist, it would be called a *shrinkage estimator*.

   (a) Compute the expected value and variance for $\hat{p}$ and $\check{p}$. Thus or otherwise, compute the MSE for each.

   (b) Create a dataset in **R** containing a values for p as follows:

   ```
   P <- seq(0,1,.01).
   ```

   Plot the MSE for the two estimators for that set of values for $n = 5$, $n = 10$, $n = 20$, $n = 50$, and $n = 100$. Note: I'd like a seperate plot for each n. You will need to evaluate the MSE for each value of P, I'd write an **R** function to do this. To make a nice plot in base **R** use something like

   ```
   plot(P,MSE1,type="l",lwd=2,col="red")
   lines(P,MSE2,lwd=2,col="blue")
   ```

   Or use ggplot, in which case I leave it to you to figure out how to do something similar!

   (c) Which estimator seems preferrable under which conditions?

3. Even when we can't analytically evalute the MSE of an estimator, we can always study it by means of a simulation. Generate 100 samples of size 20 from a Gamma$(5, 1/10)$ distribution as follows

   ```
   Data <- matrix(rgamma(2000,5,1/10),ncol=100)
   ```

   The matrix Data will have 20 rows and 100 columns. Think of each column as sample of size 20.

   (a) Write two **R** functions, one to compute the Method of Moments estimators for α and β, the other to compute the MLE's. (Ask for help if you need it!)

   (b) Compute the estimates for each sample, then compute and compare the means, variances, and MSE's for both estimators.