https://classroom.github.com/a/vNIkxPaw

Assume $X \sim \text{Normal}(\mu = 3, \sigma^2 = 2)$. The probability density function for the normal distribution is

$$f(x|\mu, \sigma) = (2\pi\sigma^2)^{-1/2} e^{(\frac{-(x-\mu)^2}{2\sigma^2})}$$

for $\mu \in \mathbb{R}$, $\sigma > 0$, and $x \in \mathbb{R}$.

1. Make a plot of the density function. Since the support S_X is unbounded, the bounds of `x` in your plot should be

$$(\max \{\mathbb{E}(X) - 4 * \mathbb{D}(X), \min(S_X)\}, \min \{\mathbb{E}(X) + 4 * \mathbb{D}(X), \max(S_X)\})$$

2. Generate $N = 101$ random observations and store them into a variable `x`.
3. Write an `R` function that evaluates negative one times the simplified log-likelihood for both parameters (μ, σ) based on `x`.
4. Use the above function to produce estimates, $(\hat{\mu}, \hat{\sigma})$.
5. Write an `R` function that evaluates negative one times the simplified log-likelihood for just the parameter μ. Remember to simplify by removing terms that would be lost by taking a derivative and setting equal to zero.
6. Use `optim()` with your above function to produce an estimate for μ.
7. Do calculus by hand to solve for the maximum likelihood estimator for μ, call it $\hat{\mu}$. Write your solution as a function in `R`.
8. Estimate the expected value of X, $\mathbb{E}(X)$, using `x` and `mean()` together.
9. Do your estimates for μ appear to produce the same number?
10. Do calculus by hand to solve for the maximum likelihood estimator for σ, call it $\hat{\sigma}$. Write your solution as a function in `R`. Hint: if all you had was data `x`, what is a reasonable value to put inplace of μ?
11. Use `optim()` with your above function to produce an estimate for σ.
12. Estimate the standard deviation of X, $\mathbb{D}(X)$, using `x` and `sd()` together.
13. Do your estimates for σ appear to produce the same number? Which one seems different from the rest?