Math 5301 Homework 6

Use R to do the problems for this assignment.

1. Let $Z \sim N(0,1)$, and calculate the following.

   (a) $P(Z \leq 1.96)$. The `pnorm` command can help with this.

   (b) $P(-1 \leq Z \leq 1)$.

   (c) If $X \sim N(500, 100^2)$, find $P(400 \leq X \leq 600)$.

2. If X is a continuous random variable, its *quantile* of order p, x_p, is the unique point where

$$P(X < x_p) = p.$$

   For example, if $X \sim N(0,1)$, then its quantile of order $p = 0.9$ is $x_{0.9} = 1.28$, because

$$P(X < 1.28) = 0.9.$$

   One could also say that the 90th percentile of a standard normal distribution is 1.28. The quantile of order p is the same thing as the $100p$th percentile.

   (a) Verify that the 90th percentile of a standard normal distribution is 1.28. The `qnorm` command can help with this.

   (b) Find the quantile of order 0.975 of a standard normal distribution.

   (c) Find the 90th percentile of a $N(500, 100^2)$ distribution.

3. If $Z \sim N(0,1)$, another common notation is defining z_α to be the point where

$$P(Z > z_\alpha) = \alpha.$$

   (a) Find $z_{0.025}$ and $z_{0.1}$.

   (b) Write a function called `zalpha` that accepts a number α as input and returns z_α as output. Test it using the examples in part (a).

4. Let $T \sim t(30)$, the t-distribution with 30 degrees of freedom.

   (a) Find $P(-2.042 \leq T \leq 2.042)$. (`pt` command)

   (b) Find $t_{0.025}(30)$. (`qt` command)

5. Let $X \sim$ binomial$(n = 2000, p = 0.1)$.

   (a) Find $P(150 \leq X \leq 250)$. (`pbinom` command)

   (b) Find the 70th percentile of a binomial distribution with parameters $n = 2000$ and $p = 0.1$. (`qbinom` command)

   Summary: For the `<name>` distribution, `p<name>` calculates probabilities, and `q<name>` calculates quantiles. We will now see that `r<name>` generates random observations.

6. (a) Generate a sample of size 1000 from a $N(500, 100^2)$ distribution. (`rnorm` command)
   (b) Find the mean $\overline{X}$ and standard deviation S for this sample. (`mean` and `sd` commands)
   (c) Find the 95% confidence interval for μ based on your sample. Is $\mu = 500$ in this interval?
   (d) Write a function called `confint` that accepts a vector `x` containing a sample of data and a significance level `alpha` and returns the corresponding $1-\alpha$ confidence interval. You might have the function return a vector of length two, where the first component is the lower bound L of this interval, and the second component is the upper bound U. Test your function on the sample you generated.

7. (a) Generate $10,000$ samples of size 1000 from a $N(500, 100^2)$ distribution.
   (b) For each sample, find the corresponding 95% confidence interval for μ, and determine if $\mu = 500$ is actually in this interval or not.
   (c) What percentage of the confidence intervals generated actually contained μ?