Fourth homework assignment, 10/22/2020 (due 10/29/2020).

The purpose of this assignment is to allow us to explore coverage rates of confidence intervals in four different situations:

- small sample size, underlying normal distribution;
- large sample size, underlying normal distribution;
- small sample size, underlying non-normal distribution;
- large sample size, underlying non-normal distribution.

Results will be aggregated across the class and discussed on November 1. Here is some R code that you will need to complete the assignment:

```
SmallNormal <- NULL
for(i in 1:100) {
   x <- rnorm(3, 50, 10)
   low <- mean(x) - 1.96 * 10/sqrt(3)
   high <- mean(x) + 1.96 * 10/sqrt(3)
   SmallNormal[i] <- 1
   if(50 < low) SmallNormal[i] <-0
   if(50 > high) SmallNormal[i] <- 0
}

LargeNormal <- NULL
for(i in 1:100) {
   x <- rnorm(50, 50, 10)
   low <- mean(x) - 1.96 * 10/sqrt(50)
   high <- mean(x) + 1.96 * 10/sqrt(50)
   LargeNormal[i] <- 1
   if(50 < low) LargeNormal[i] <-0
   if(50 > high) LargeNormal[i] <- 0
}

SmallUniform <- NULL
for(i in 1:100) {
   x <- runif(3)
   low <- mean(x) - 1.96 * sqrt(1/12)/sqrt(3)
   high <- mean(x) + 1.96 * sqrt(1/12)/sqrt(3)
   SmallUniform[i] <- 1
   if(0.5 < low) SmallUniform[i] <-0
   if(0.5 > high) SmallUniform[i] <- 0
}
```

```
LargeUniform <- NULL
for(i in 1:100) {
   x <- runif(50)
   low <- mean(x) - 1.96 * sqrt(1/12)/sqrt(50)
   high <- mean(x) + 1.96 * sqrt(1/12)/sqrt(50)
   LargeUniform[i] <- 1
   if(0.5 < low) LargeUniform[i] <-0
   if(0.5 > high) LargeUniform[i] <- 0
}

mean(SmallNormal)
mean(LargeNormal)
mean(SmallUniform)
mean(LargeUniform)
```

Run the *R* program. The means that you calculated at the end are proportions of confidence intervals that captured the true mean in each condition. Report the proportions you obtained. (If you copy and paste the code, you are likely to get an error because *R* may see the minus signs as special characters. You can correct this by pasting the code into a plain text editor such as notepad and manually replacing the characters with true minus signs.)

Now answer the following questions:

- Assume that all assumptions are met. How many confidence intervals out of your 100 simulated intervals should capture the true mean?
- How do you know?
- Because each of you individually has performed only 100 samples and confidence interval calculations, your results cannot be expected to have exactly the values predicted under the central limit theorem. However, there may also be some systematic deviations from the expected counts that occur because requirements of the central limit theorem are not quite met. Keeping that in mind, discuss your results: What did you expect? What did you observe?