1. Exercise 9.4 in the text. Do part (a) and assume that the given sample standard deviation of 6.9 is the true standard deviation of the entire population. Replace part (b) with: Is your confidence interval you constructed in (a) exact or approximate? What information did you use about the distribution of the sample mean of student heights to construct you confidence interval?

2. Exercise 9.8 in the text.

3. Exercise 9.10 in the text. Assume that the given sample standard deviation of 7.8 is the true standard deviation of the entire population.

4. Exercise 9.12 in the text. Assume that the given sample standard deviation of 15 is the true standard deviation of the entire population. Additionally, use the confidence interval you construct to test the hypothesis $H_0 : \mu = 250$. What is the Type I error of your test?

5. Exercise 10.2 in the text.

6. Exercise 10.20 in the text. Assume that the weights are normally distributed and that the given sample standard deviation of 0.24 is the true population standard deviation.

7. Guided `R` exercise. This exercise will help you to understand the Central Limit Theorem (CLT) using computer simulations. Below I have written some `R` code to illustrate the CLT in action. The CLT says that the sample mean behaves approximately like a normal random variable when the sample size is large, regardless of the distribution of the data.

Let's test this out using distributions that are very different from the normal distribution. The first distribution we will use is the Gamma distribution. Run the following code to see a picture of this particular Gamma distribution:

```
x = seq(from = 0, to = 10, by =0.01)
y = dgamma(x,2,4)
plot(x,y,type = 'l')
```

Not very normal, is it? The following code simulates 1000 datasets of size n from this Gamma distribution, computes the sample means of each data set, and creates a histogram of these sample means. This histogram approximates the distribution of the sample means. for large enough n, we should see a histogram that looks like a normal distribution. Change the value of n in the code below until the historgram looks normal.

```
# Simulation of sample means of Gamma data


n=1
xbar = matrix(0,1000,1)
for(i in 1:1000){
sample = rgamma(n,2,4)
xbar[i,1] = mean(sample)
}
hist(xbar)
```

The uniform distribution is also very non-normal. It is flat and only nonzero on a finite interval. We will try the same experiment as above using the uniform distribution. Look at the plot of the uniform distribution and then change the value of n in the code below until the historgram looks normal.

```
x = seq(from = 0, to = 1, by =0.01)
y = dunif(x,0,1)
plot(x,y,type = 'l')

# Simulation of sample means of Uniform data
n=1
xbar = matrix(0,1000,1)
for(i in 1:1000){
sample = runif(n,0,1)
xbar[i,1] = mean(sample)
}
hist(xbar)
```

We have seen the normal approximation to the binomial before. It all comes from the CLT. Even though the binomail is discrete and the normal is continuous, the CLT still applies! Try the same experiment with the binomial

```
x = seq(from = 0, to = 20, by =1)
y = dbinom(x,20,0.1)
plot(x,y,type = 'p')

# Simulation of sample means of binomial data
```

```
n=1

xbar = matrix(0,1000,1)

for(i in 1:1000){

sample = rbinom(n,20,.1)

xbar[i,1] = mean(sample)

}

hist(xbar)
```

Tell me how large a value of n you needed for each distribution. Your book says CLT works fine for $n \geq 30$. Based on your experiments, do you agree?