Activity 2: Simulation and sampling distributions

In this activity, we'll use a mix of theory and simulation to derive important distributions we'll use all semester. By the end of this activity, you should have some familiarity with the normal, t-, chi-squared, and F-distributions. You'll be able to define each distribution and know when and how to calculate probabilities under each distribution.

Scenario: How many hours do you study in a typical week? How does that compare to the time you spent studying as a freshman?

In September of 2013, as part of the MAP-Works survey, first-year students at St. Ambrose were asked:

- In an average week, how many hours do you spend studying? (question S1_NA150)
- In an average day, how many hours do you spend spend relaxing or socializing? (S1_D148)

Some quick visualizations of the responses from 528 students are displayed below:

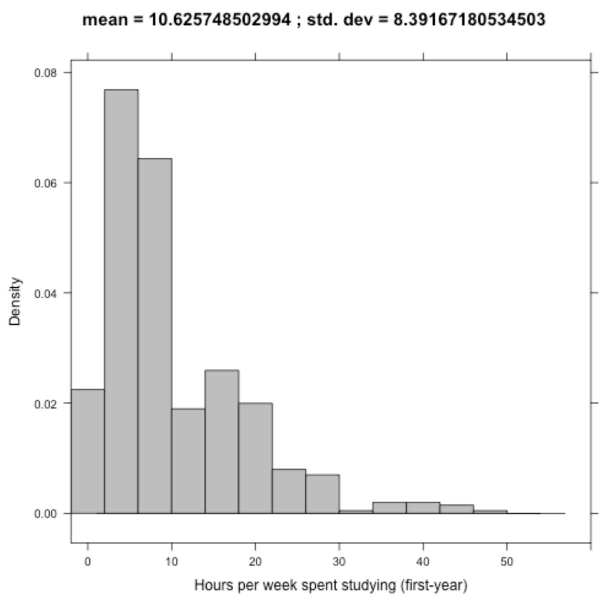

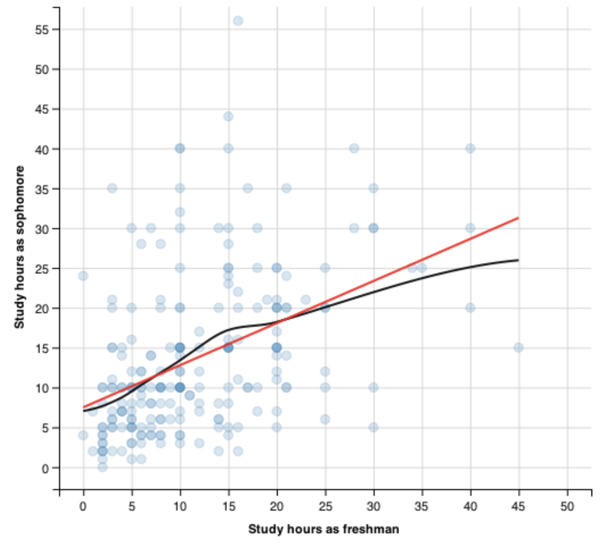

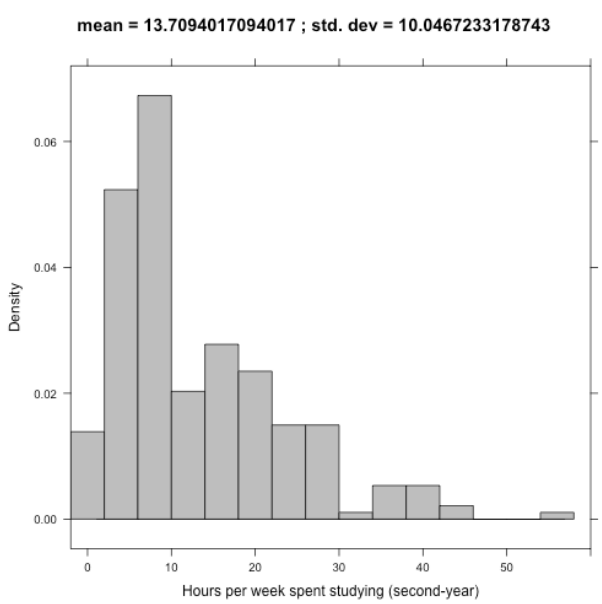

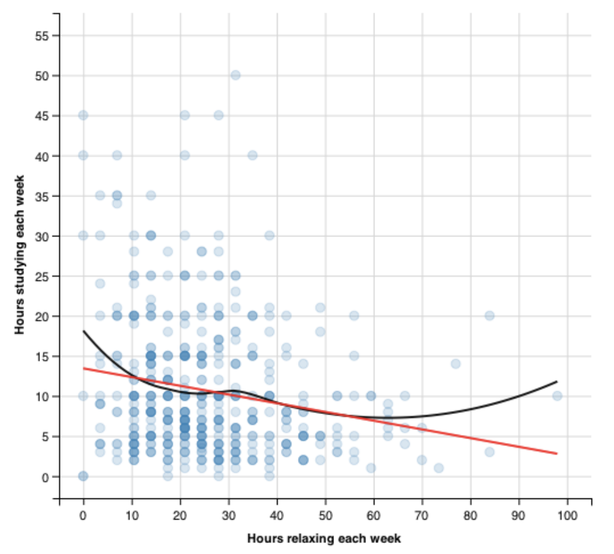

1. We're going to focus on the distribution of hours per week freshmen reported studying in 2013. We'll pretend as though it represents our population of interest.

   Based on the kernel density plot to the right, describe this distribution:

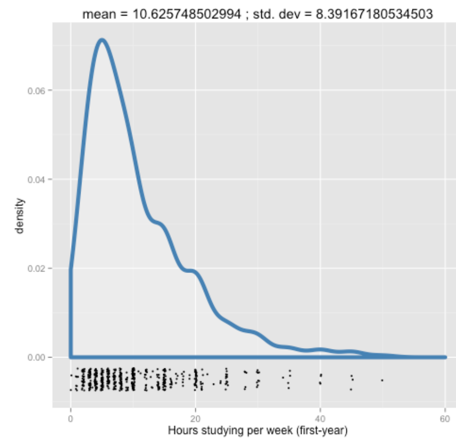

2. Obviously, these study hours do not follow a normal distribution. Describe what the following plots display:

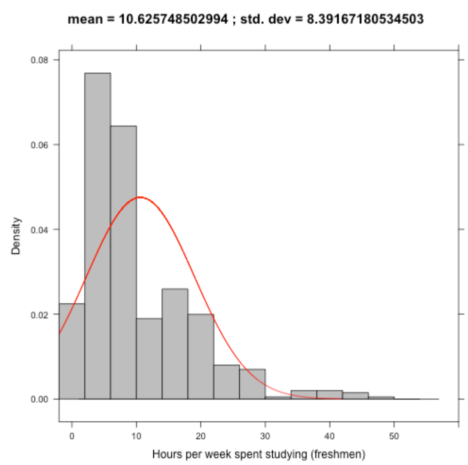

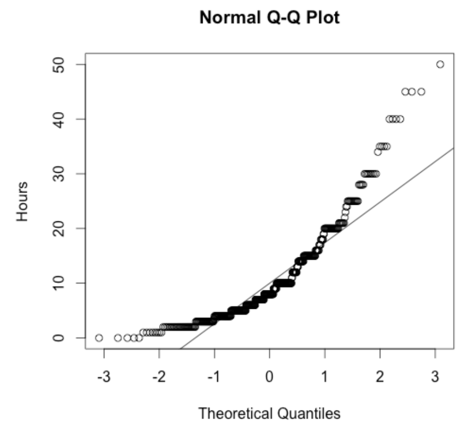

3. Suppose we repeatedly take independent, random samples of size *n=3* from this population and calculate $\overline{X}$ for each sample. Describe the distribution of sample averages we'd get if we took an infinite number of samples. What is the average and standard deviation of this sampling distribution?

   Would the Central Limit Theorem apply in this example? What is the Central Limit Theorem?

4. Using R, I had a computer take 50,000 samples of size *n=3* from that distribution. Here's the sampling distribution of sample averages I obtained. Does it match your expectations? Did the CLT apply?

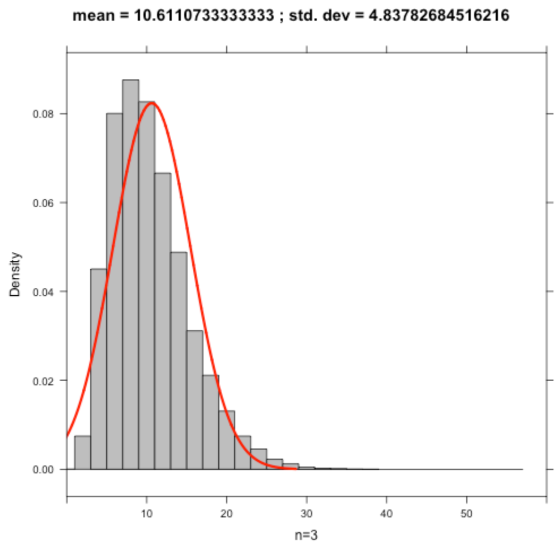

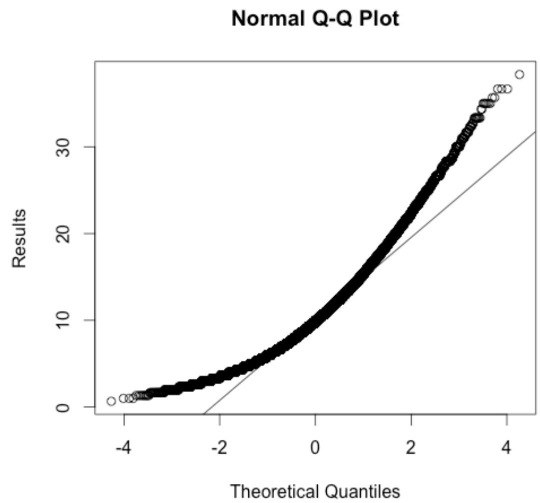

5. I then had the computer take 50,000 samples of size <u>*n=30*</u>. The sampling distribution is displayed below. Verify the mean and standard error of this distribution. Explain *why* this sampling distribution is so much thinner than the one above.

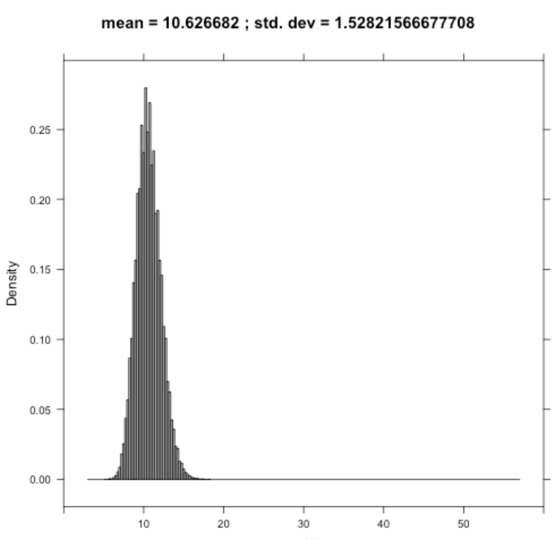

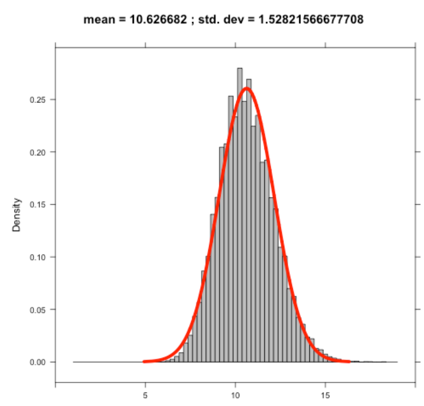

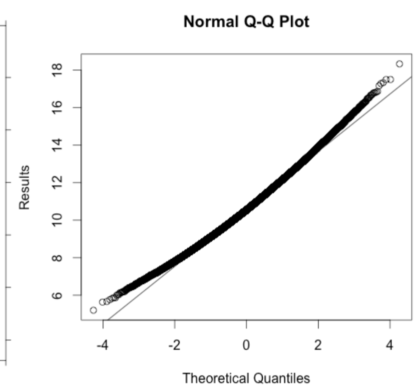

6. Finally, I repeated this one more time with a sample size of n=1000. Verify the results.

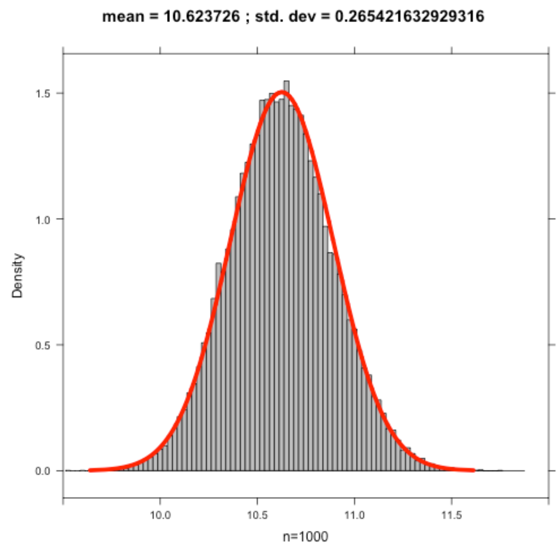

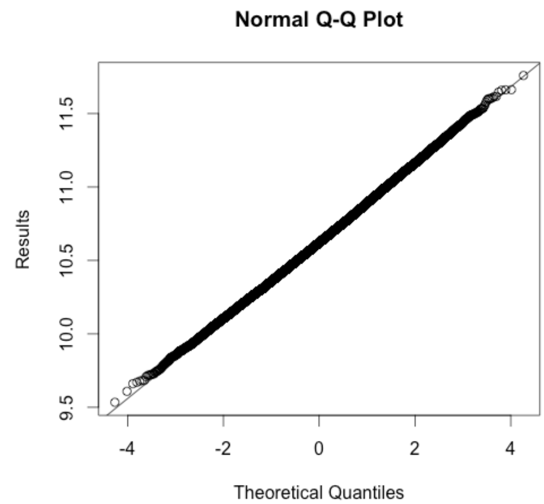

The Central Limit Theorem tells us what to expect from the sampling distribution of the sample mean. Once we know the sampling distribution, we can come up with formulas for confidence intervals and methods for inference.

But what would the sampling distributions of other statistics look like?

You can use the following applet to investigate the sampling distribution of the median, range, or other statistics: http://onlinestatbook.com/stat_sim/sampling_dist/index.html

In this activity (and throughout most of this semester), we're going to focus on distributions of variances.

7. Before we get to variances, suppose we're going to do the following:
   - Take a sample of *n* observations from a population
   - Convert all the observations to z-scores
   - Square each of those z-scores
   - Calculate the sum of those squared z-scores

   What do z-scores represent? How do you convert an observation (x) to a z-score?

$$\sum_{i=1}^{n} z_i^2 =$$

8. Suppose our population follows a normal distribution with a mean near 100 and a standard deviation near 16:

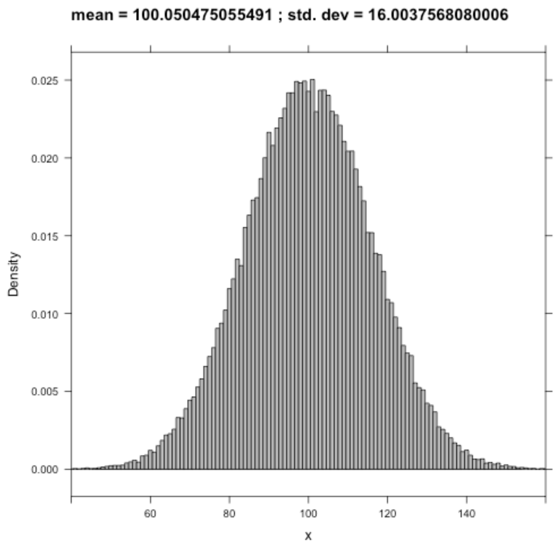

If we repeatedly sampled n=3 observations, these observations would (most likely) come from what part of the distribution?

If we converted those observations to z-scores, squared those z-scores to make them positive, and then summed them, would we expect that sum to be large or small?

Is it possible to get sums that are extremely large?

9. Using that logic, describe the sampling distribution you expect for these sum of squared z-scores (if the number of samples we take approaches infinity). Try to sketch it below.

10. This time, pretend as though we take samples of size n=100. Sketch a prediction for the sampling distribution of the sum of squared z-scores.

11. The next page shows simulated sampling distributions (based on 50,000 samples) of these sum of squared z-scores. Sampling distributions were simulated for samples of size n=3, n=10, and n=100. Fill-in the table below the distributions and see if you can generalize the results to any sample of size *n=k*.

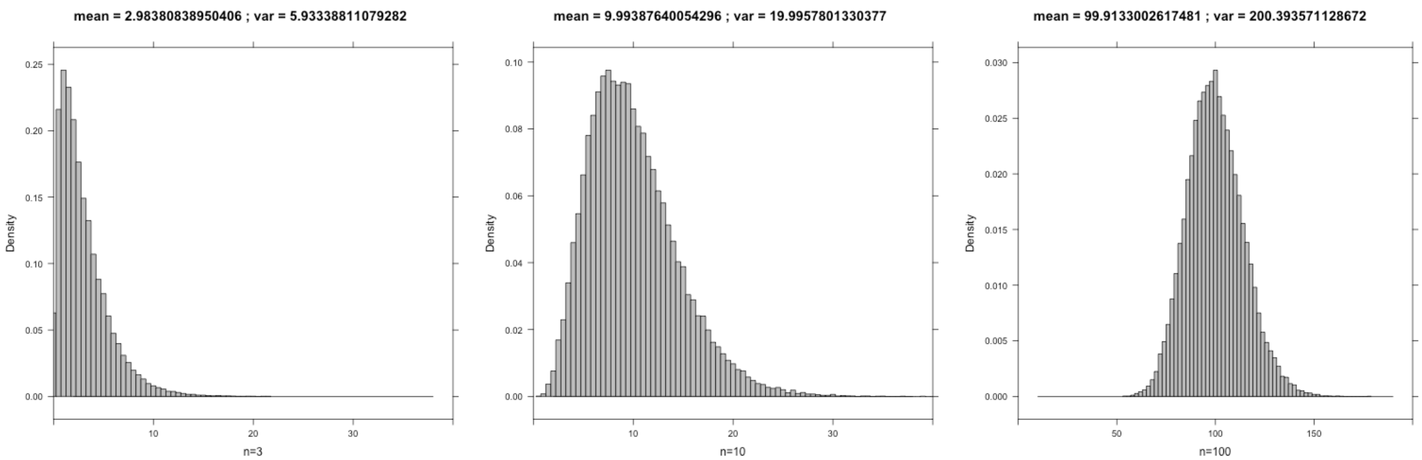

Sample size	Shape of distribution	Mean of distribution	Standard error
n = 3	unimodal, heavy positive skewed	3	6
n = 10	unimodal, moderate positive skewed	10	20
n = 100	unimodal, light positive skewed	100	200
n = k			

These positively skewed distributions of the sum of squared z-scores are called ***Chi-squared (χ^2)*** distributions.

If $Q = z_1^2 + z_2^2 + \ldots + z_k^2$, the distribution of Q would be $\chi^2_{df=k} = \dfrac{x^{\left(\frac{k}{2}-1\right)} e^{\left(-\frac{x}{2}\right)}}{2^{\left(\frac{k}{2}\right)} \Gamma\left(\frac{k}{2}\right)}$ where $E\left[\chi_k^2\right] = k$ and $Var\left[\chi_k^2\right] = 2k$

To help sketch the chi-squared distribution, it might be useful to note the mode is always k – 2.

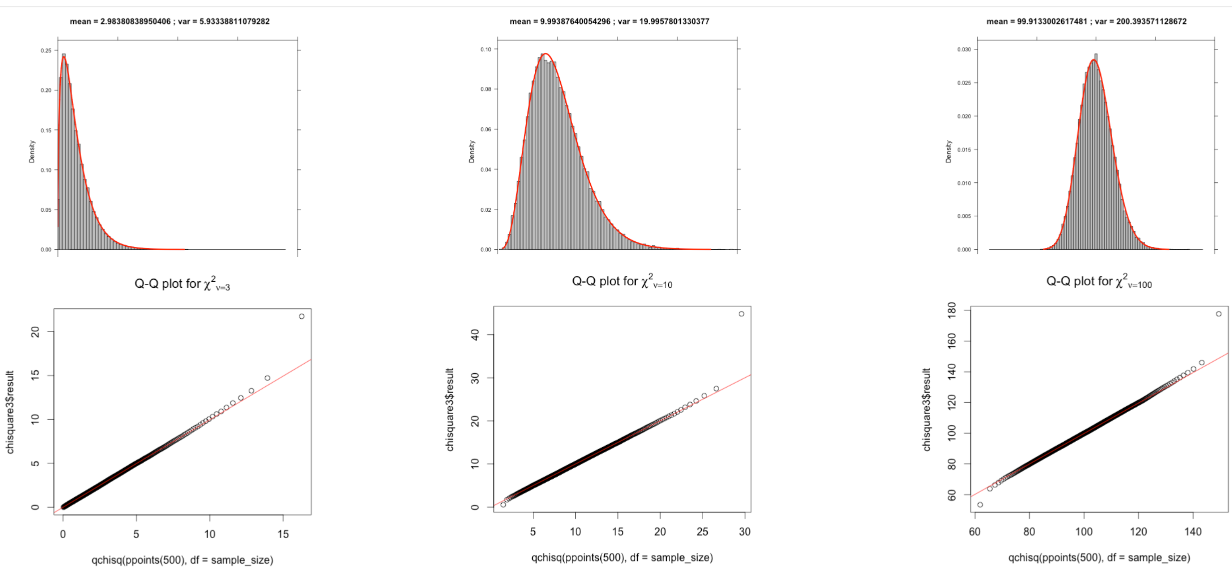

12. Why are we learning about the chi-squared distribution? What possible use could we have for finding the sum of squared z-scores? To answer that, explain what is happening at each step of the following derivation:

$$\chi^2 = \sum_{i=1}^{n} z_i^2 = \sum_{i=1}^{n}\left(\frac{x_i - \bar{X}}{\sigma}\right)^2 = \frac{\sum_{i=1}^{n}\left(x_i - \bar{X}\right)^2}{\sigma^2} = \frac{(n-1)\sum_{i=1}^{n}\frac{\left(x_i - \bar{X}\right)^2}{(n-1)}}{\sigma^2} == \frac{(n-1)s_x^2}{\sigma^2}$$

$$\chi_{n-1}^2 = \frac{(n-1)s_x^2}{\sigma^2}$$

$$\chi_{n-1}^2\sigma^2 = (n-1)s_x^2$$

$$\frac{\chi_{n-1}^2\sigma^2}{(n-1)} = s_x^2$$

We just showed that the chi-squared distribution is directly related to the distribution of sample variances.

We'll demonstrate that in a little bit (and explore why this is useful). Fill-in-the-blanks to derive one more thing:

$$E\left[s_x^2\right] = E\left[\frac{\chi_{n-1}^2\sigma^2}{(n-1)}\right] = \frac{\sigma^2}{(n-1)}E\left[\chi_{n-1}^2\right] = \frac{\sigma^2}{(n-1)}(\_\_\_\_\_\_\_\_\_\_) = \_\_\_\_\_\_\_\_\_\_\_\_$$

13. What's the importance (or implication) of what we just derived?

14. Write out the formulas for the population standard deviation and (unbiased) sample standard deviation. What do they represent? How do you convert a standard deviation to a variance?

15. To demonstrate, one last time, that the (n-1) in the denominator of our sample standard deviation makes it unbiased, go to http://onlinestatbook.com/stat_sim/sampling_dist/index.html and click **Begin.**

Let's begin with the default normal distribution. Record the **variance** of this population here: ______________

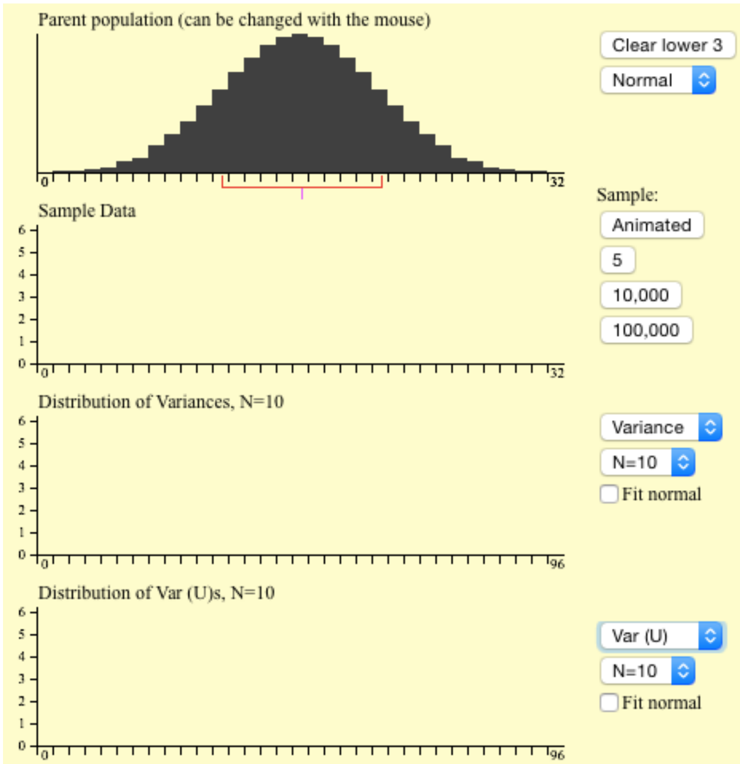

For the first sampling distribution, let's repeatedly sample n=10 observations and calculate the variance using:

$$\sum \frac{\left(x_i - \bar{X}\right)^2}{n}$$

For the second sampling distribution, let's repeatedly sample n=10 observations and calculate the unbiased estimate of the population variance using:

$$\sum \frac{\left(x_i - \bar{X}\right)^2}{n-1}$$

Click **Animated** and watch as 10 observations are randomly selected from our population distribution. The two variances are then calculated from those 10 observations and plotted on the sampling distributions.

16. If you understand what is going on, go ahead and simulate a million samples of size n by clicking **100,000** ten times. Do the results agree with what we derived on the previous page? What *is* an unbiased estimator?

17. On the previous page, we showed that a chi-squared distribution is directly related to the distribution of sample variances. But what, exactly, is this relationship? To find out, let's run some simulations. Let's repeatedly take samples of size n=3, calculate the variance of each sample, and examine the distribution of sample variances.

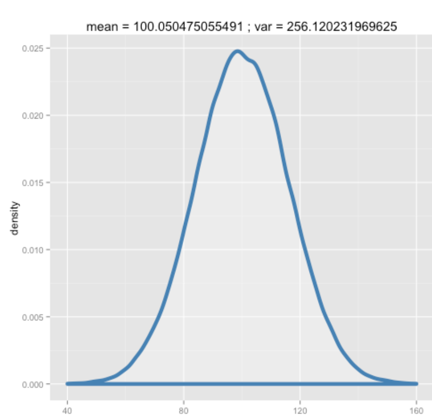

population (mean = 100, var = 256)

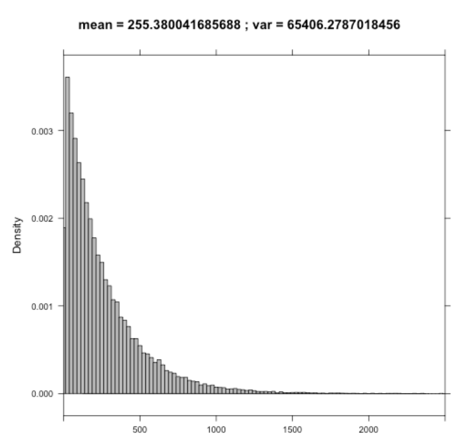

distribution of 50,000 sample variances

How do we know this is **not** a chi-squared distribution?

How many degrees of freedom would a chi-squared distribution have in this case?

18. In question #12, we derived $\chi^2_{n-1} = \frac{(n-1)s_x^2}{\sigma^2}$. That means, in order to get a chi-squared distribution, we must:

- Repeatedly sample n observations
- Calculate the variance of each sample
- Multiply each sample variance by (n-1) and divide each sample variance by the population variance

So that's what I did. I had a computer take 50,000 samples of size n=3 and calculate $\frac{(n-1)s_x^2}{\sigma^2}$ for each sample.

I then compared the distribution of these 50,000 values to the theoretical chi-squared distribution with 2 degrees-of-freedom. Finally, I repeated this process with n=10.

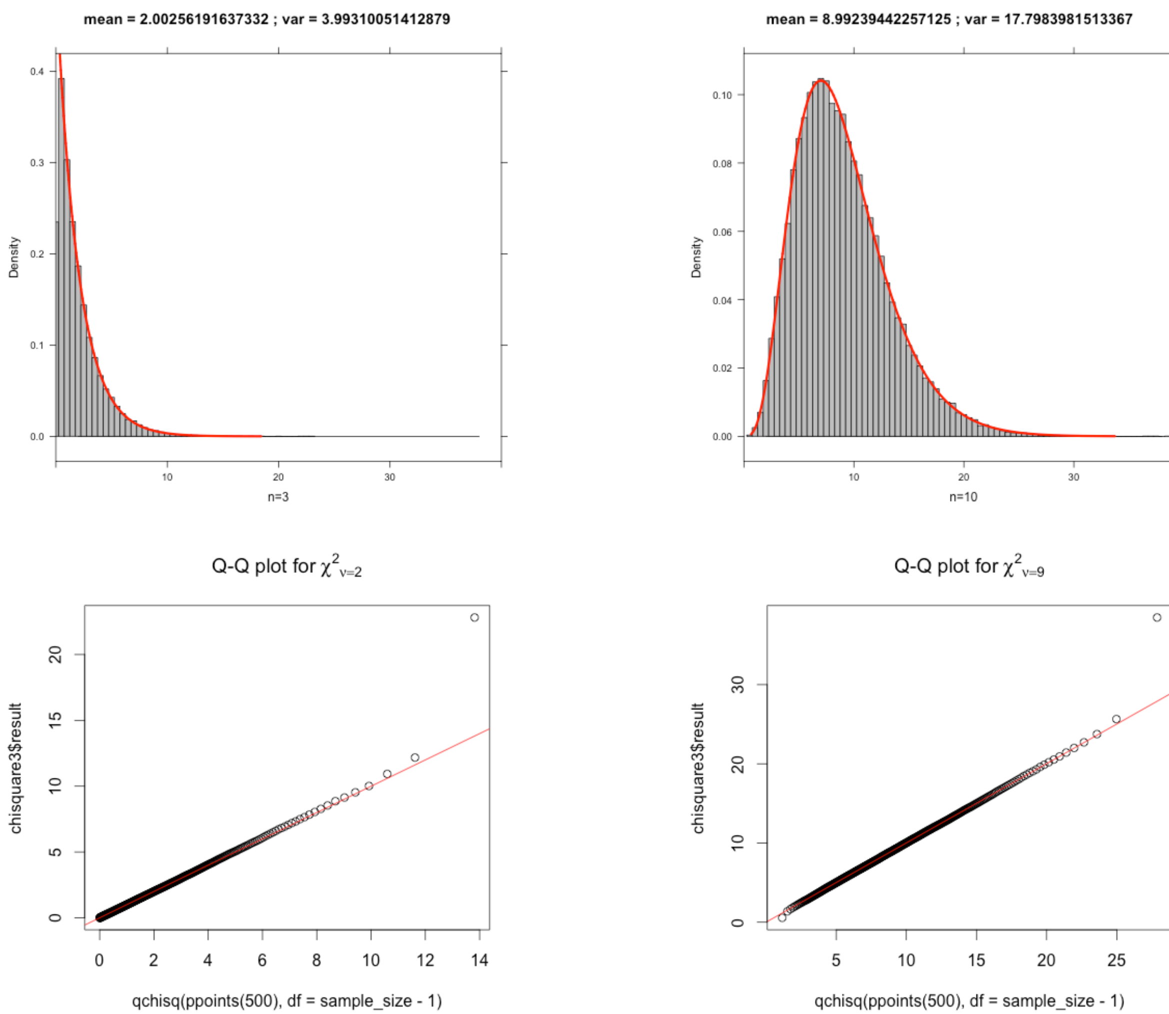

19. It looks like it works – although it's still not obvious how this could be useful. How could we calculate $\frac{(n-1)s_x^2}{\sigma^2}$ if we have an unknown population (and, therefore, an unknown population variance)?

And we only showed this works if we start with a normal population. Does it work if our population is non-normal?

To find out, let's once again take a look at the distribution of hours St. Ambrose freshmen spend studying each week.

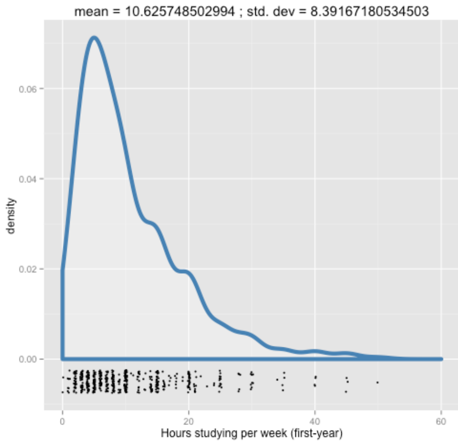

Let's repeatedly take samples of size n=3 from this non-normal distribution and calculate $\frac{(n-1)s_x^2}{\sigma^2}$ for each sample.

If the distribution of these sample values approximates a chi-squared distribution, we'll have evidence it applies to non-normal distributions.

If the sampling distribution does **not** approximate a chi-squared distribution, we'll have evidence it applies **only** to normal populations.

Let's take a look at the simulated sampling distribution (based on 50,000 samples) and compare them to the theoretical chi-squared distributions with 3-1 degrees of freedom:

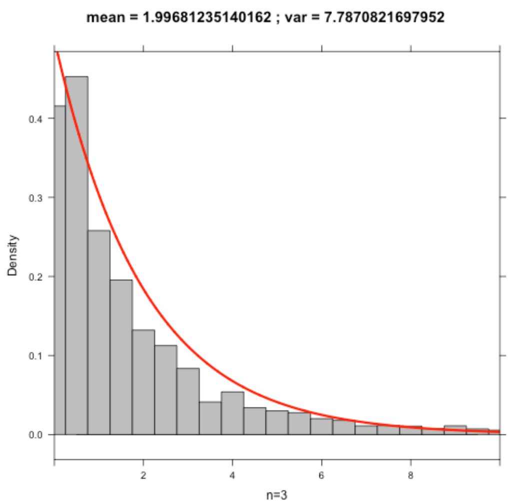

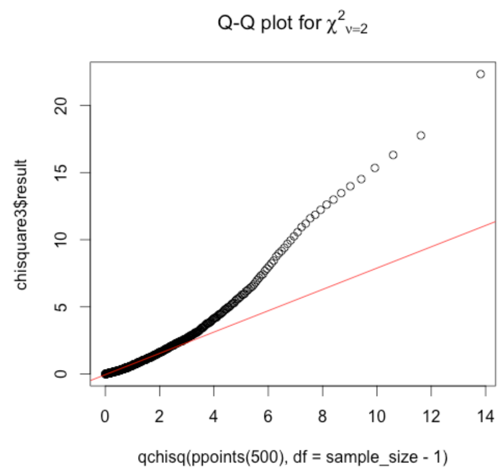

Ok, so that doesn't look like it follows a chi-squared distribution. Maybe it's like the Central Limit Theorem. Maybe it only works if we sample from a normal population or if our sample size is large. Let's try it out with a sample size of n=100.

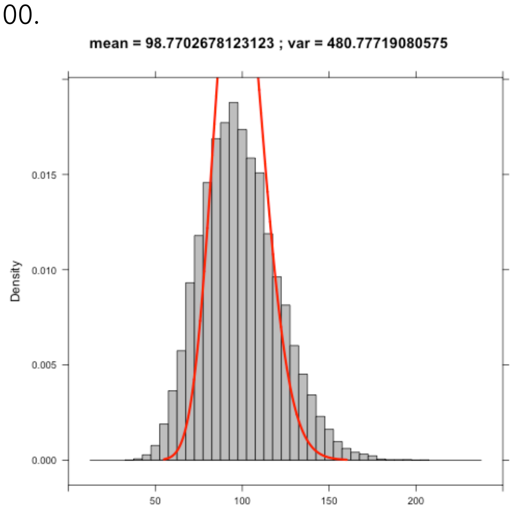

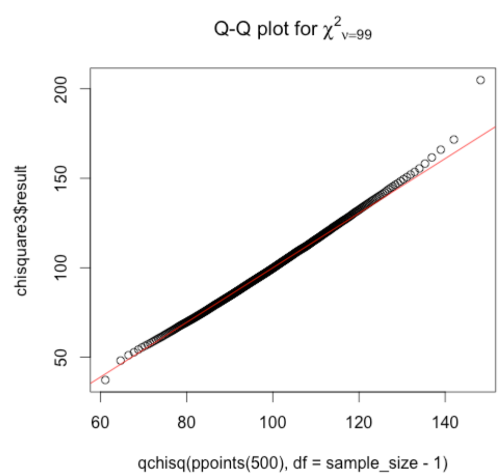

It still doesn't follow a chi-squared distribution (and the mean and variance of our sampling distribution indicates this). So it looks as though this sampling distribution only approximates a chi-squared distribution when we take independent, random samples from a **normal population**.

20. Look how long this activity is and we <u>still</u> haven't gotten to anything useful. I promise that after this activity, this class is almost completely applications-based. Before we see an application of the chi-squared distribution, let's (re-)introduce the concept of degrees-of-freedom and calculate some chi-squared probabilities.

In a previous statistics course, you may have briefly discussed *degrees of freedom*. One working definition is:

The **number of independent scores** used in the estimate **minus** the **number of parameters estimated** in that estimation

We're going to deal with variances quite a bit in this class. Let's look at the formula for the unbiased estimate of the population variance to see if we can figure out the degrees of freedom.

$$\sum \frac{\left(x_i - \bar{X}\right)^2}{n-1}$$

If we have a sample of size n, identify the following:

Number of independent scores used in this formula: ___________

Number of parameters estimated: __________ Degrees of freedom: ______________

21. Quickly sketch and label two chi-square distributions: one with 5 degrees of freedom and another where df = 25.

22. A chi-square distribution with 8 degrees of freedom is displayed to the right. Using an online calculator, find the following:

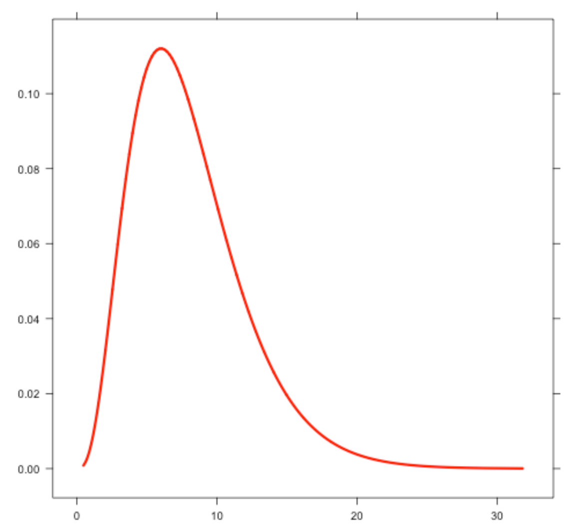

$P\left(\chi_8^2 > 15\right) =$ _________

$P\left(\chi_8^2 > b\right) = 0.025.\ \ b =$ _________

$P\left(\chi_8^2 < a\right) = 0.025.\ \ a =$ _________

$P\left(\_\_\_\_\_\_ < \chi_8^2 < \_\_\_\_\_\_\right) = 0.95$

Calculators:

http://stattrek.com/online-calculator/chi-square.aspx
http://lock5stat.com/statkey/theoretical_distribution/theoretical_distribution.html#chi
http://www.stat.berkeley.edu/~stark/Java/Html/ProbCalc.htm

Table: http://bradthiessen.com/html5/stats/m301/chitable.pdf

23. Let's generalize what we just found:

$$0.95 = P\left(2.180 < \chi_8^2 < 17.534\right)$$

$$0.95 = P\left(\chi^2_{8,\ \alpha=.025} < \chi_8^2 < \chi^2_{8,\ \alpha=.975}\right)$$

Rewriting our values of 2.18 and 17.534 as chi-squares

$$100(1-\alpha) = P\left(\chi^2_{n-1,\ \alpha/2} < \chi^2_{n-1} < \chi^2_{n-1,\ 1-\alpha/2}\right)$$

Rewriting 0.95 as having an alpha error of 0.05

$$100(1-\alpha) = P\left(\chi^2_{n-1,\ \alpha/2} < \frac{(n-1)s_x^2}{\sigma^2} < \chi^2_{n-1,\ 1-\alpha/2}\right)$$

Substituting what we derived earlier for the chi-squared

$$100(1-\alpha) = P\left(\frac{1}{\chi^2_{n-1,\ 1\text{-}\alpha/2}} < \frac{\sigma^2}{(n-1)s_x^2} < \frac{1}{\chi^2_{n-1,\ \alpha/2}}\right)$$

Taking the reciprocal of everything

$$100(1-\alpha) = P\left(\frac{(n-1)s_x^2}{\chi^2_{n-1,\ 1\text{-}\alpha/2}} < \sigma^2 < \frac{(n-1)s_x^2}{\chi^2_{n-1,\ \alpha/2}}\right)$$

Multiplying to isolate the population variance

We just derived the formula for the confidence interval for a population variance.

A 95% confidence interval for the population standard deviation would be:

$$s_x\sqrt{\frac{(n-1)}{\chi^2_{df=n-1,\ \alpha=0.975}}} < \sigma < s_x\sqrt{\frac{(n-1)}{\chi^2_{df=n-1,\ \alpha=0.025}}}$$

Note: This interval relies heavily on the assumption that the population follows a normal distribution

24. The PGA establishes strict rules for golf balls used in its tournaments. In 1942, the PGA established a rule that golf balls must have an initial velocity between 243.75 and 256.25 ft/sec when measured at sea-level in 70° weather. Suppose a golf ball manufacturer has mastered a process in which the initial velocities of its golf balls follow a normal distribution with a mean of 251 and a standard deviation of 4.3 ft/sec.

    You sample 16 golf balls manufactured using an experimental process and find they have a mean of 250 and a standard deviation of 2.4. Does this experimental process have significantly more or less variation than the current process? Construct a 95% confidence interval for the population standard deviation of the experimental process.

25. Diabetic patients monitor their blood sugar levels with a home glucose monitor that analyzes a drop of blood from a finger stick. Although the monitor gives precise results in a laboratory, the results are too variable when it is used by patients. The process has a standard deviation of σ = 10 mg/dl.

A new monitor is developed to improve the precision of the assay results under home use. 25 individuals test the new monitor at home using drops from a sample having a glucose concentration of 118 mg/dl. The readings from the 25 tests are as follows:

Results: 125, 123, 117, 123, 112, 128, 118, 124, 116, 109, 125, 120, 123, 112, 118, 121, 122, 115, 105, 118, 115, 111, 113, 118, 131
mean = 118.48 std. dev = 6.1922

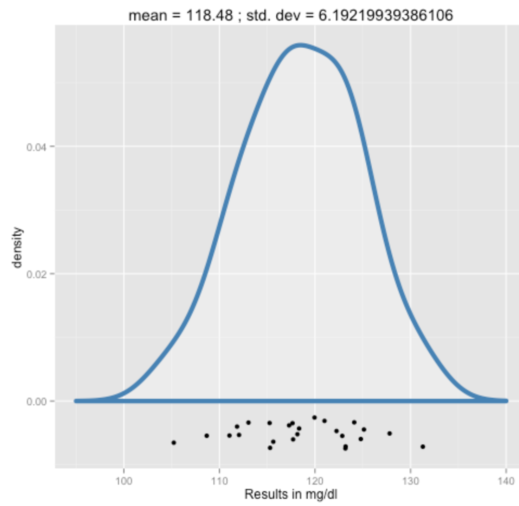

Construct a 95% confidence interval for σ to determine if it is more precise than the previous monitor. Comment on whether you think the normality assumption has been satisfied.

If you're worried about the normality assumption, you could always try a bootstrap approach to find a confidence interval. Recall a bootstrap process involves:

- Treating your sample (the 25 test results) as if they are the entire population.
- Randomly sampling n=25 observations **with replacement** from our population of size n=25.
- Calculating the statistic of interest (variance, in this case) for that sample. We'll call it the bootstrap variance.
- Repeating this process many times (I'll do 100,000 replications).
- Finding the values that cut-off the lowest and highest 2.5% of our bootstrap variances

Here's the code I used to construct this bootstrap CI in R and the output I received. Which method is "better" for constructing confidence intervals for σ: the parametric (chi-squared) method or the bootstrap method?

R-code for 95% CI for population variance using bootstrap methods

```
test <- c(125, 123, 117, …enter all data…, 118, 131)  ## Enter data
trials <- do(100000) * sd(resample(test))            ## Calculate SD of 100,000 bootstrap samples
with(trials, quantile(result, c(0.025, 0.975)))      ## Find the limits of the CI
```

Output:
```
        5%         95%
  4.487761    7.555351
```

26. We won't really use the chi-squared distribution until much later in the semester. The chi-square distribution does relate to other important distributions we'll use all throughout the semester, though.

Remember when you learned about the t-distribution in a previous statistics class? You probably learned it is the distribution of sample averages when we have an unknown population mean. In other words:

$$\frac{\bar{X}-\mu}{s_x/\sqrt{n}} \sim t_{n-1}$$

At least when I teach the t-distribution, I kind of wave my hands and explain that sample standard deviations can vary from sample-to-sample so we need a wider distribution (than the standard normal, z, distribution) to reflect that increased variability. Here's a real derivation of the t-distribution:

We'll start with the z-distribution (sampling distribution of sample averages under the CLT) and the definition of chi-squared distribution.

$$z = \frac{\bar{X}-\mu}{\sigma/\sqrt{n}} \quad \text{and} \quad \chi^2_{n-1} = \frac{(n-1)s_x^2}{\sigma^2}$$

$$\sqrt{\frac{\chi^2_{n-1}}{\text{degrees of freedom}}} = \sqrt{\frac{\chi^2_{n-1}}{(n-1)}} = \sqrt{\frac{(n-1)s_x^2}{(n-1)\sigma^2}} = \sqrt{\frac{s_x^2}{\sigma^2}} = \frac{s_x}{\sigma}$$

$$\frac{z}{\sqrt{\frac{\chi^2_{n-1}}{(n-1)}}} = \frac{\left(\frac{\bar{X}-\mu}{\sigma/\sqrt{n}}\right)}{\left(\frac{s_x}{\sigma}\right)} = \frac{\bar{X}-\mu}{s_x/\sqrt{n}} = t_{n-1}$$

27. The chi-squared distribution also contributes to an extremely important distribution in this class: the F-distribution.

As we've seen, the chi-squared distribution is useful when we want to make inferences about a single variance or standard deviation. We can use the chi-squared distribution to construct a confidence interval for σ.

Suppose we're interested in comparing the variances from two groups. Here are two fictitious examples:

a) A hospital must select between two types of blood pressure monitors. After calibrating both types, they test each monitor 21 times. The first type, an inexpensive armband/pump device, had a standard deviation of 6.1 (<u>variance of 37.21 units</u>). The second type, a more expensive automated device, had a standard deviation of 3 (<u>variance of 9 units</u>). The hospital must decide which brand to buy based on the precision of the instruments.

b) One supplier provides upholstery fabric with an average durability of 74,283 DR and a <u>variance of 21,864,976</u>. Another supplier provides a lower average durability of 74,200 DR and a lower <u>variance of 20,250,000</u>. All measurements are based on a sample of n=61. You must decide which supplier to purchase from based on the variance of their fabric.

Which scenario (a or b) represents the bigger discrepancy between group variances? How do we compare variances?

28. Let's start with two normal distributions that have different means but (approximately) the same variances:

Dataset A

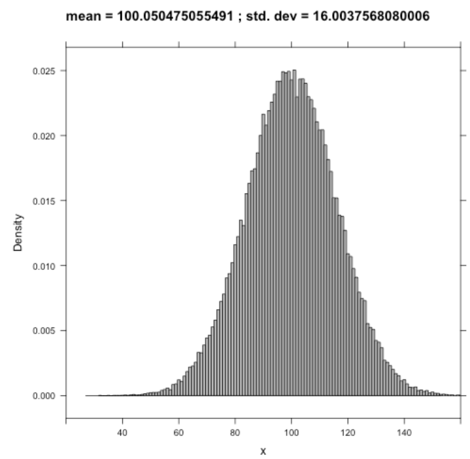

Dataset B

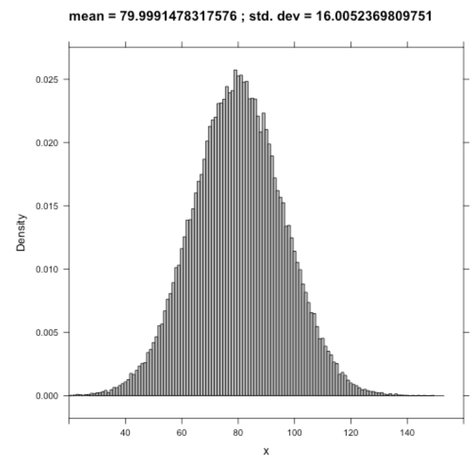

We're going to take a sample of size n=100 from each distribution. For each sample, we'll calculate the variance. We'll then calculate the ratio of variances (bigger variance divided by the smaller variance). We'll repeat this process 25,000 times to simulate the sampling distribution of these variance ratios.

Before we look at the sampling distributions, let's see if we can predict how they'll look. If we take a random sample of n=100 observations from dataset A and then B (both of which have the same population variances), how will those variances (most likely) compare to each other? Will one be much bigger than the other? If we always divide the larger variance by the smaller variance, what's the smallest ratio we could get? What's the largest ratio?

Here's the simulated sampling distribution of the ratios of variances:

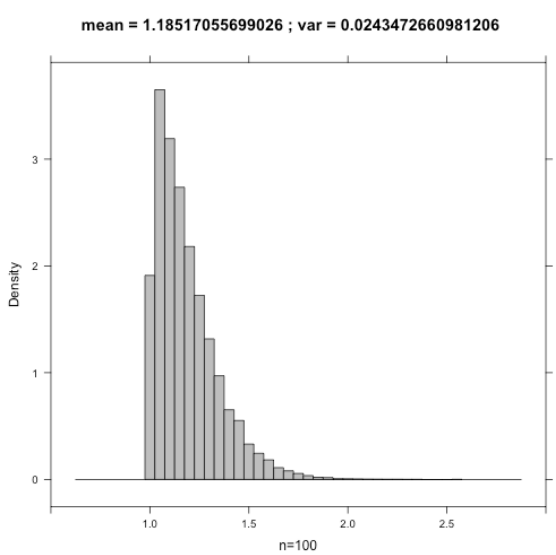

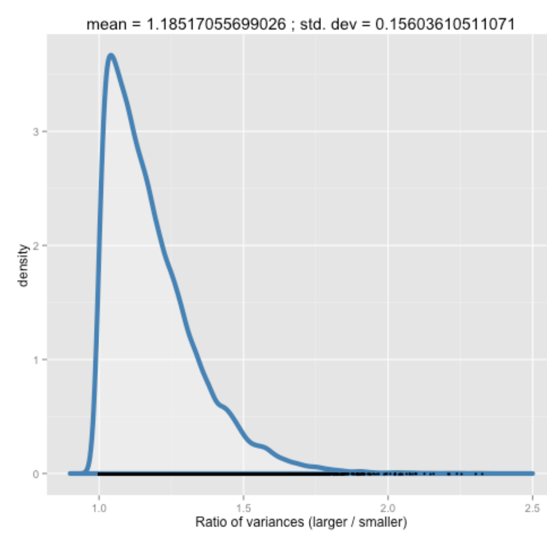

When samples are drawn from populations with <u>equal</u> variances, we expect the ratio of their sample variances to be near: _____

29. Suppose we have populations with different variances. What will the distribution of the ratios of their variances look like?

Dataset A

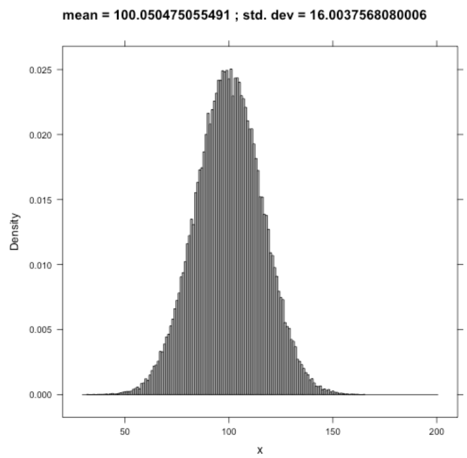

Dataset B

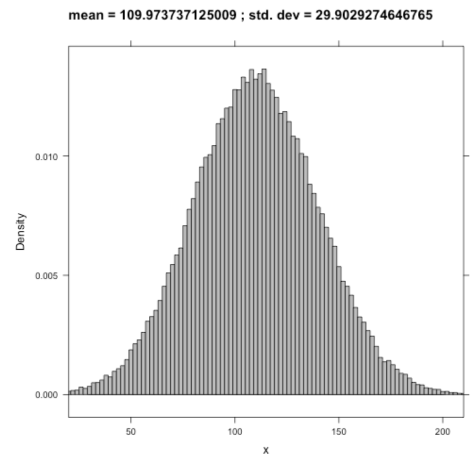

Here's the sampling distribution of the ratios of variances (based on 25,000 simulations of samples of size n=100):

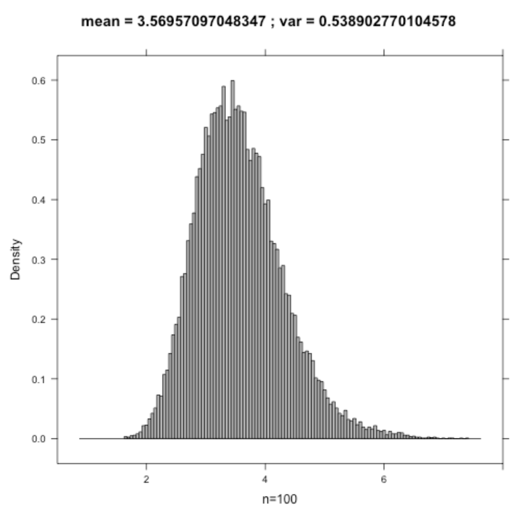

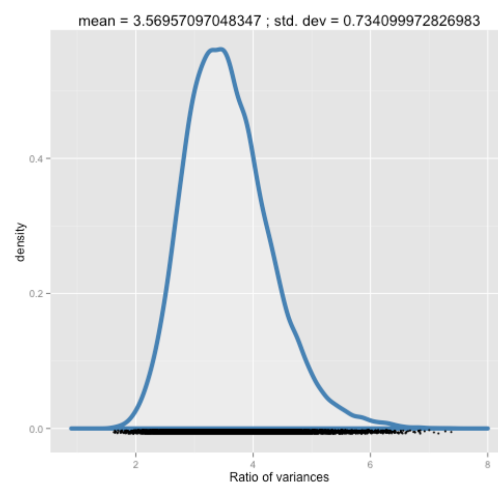

When samples are drawn from populations with <u>unequal</u> variances, we expect the ratio of their sample variances to be: _______

30. Try to follow this derivation and see if you can explain each step. What was derived?

$$\chi^2_{a,\,\mathrm{df}=n\text{-}1} = \frac{(n_a - 1)s_a^2}{\sigma_a^2} \quad \text{and} \quad \chi^2_{b,\,\mathrm{df}=n\text{-}1} = \frac{(n_b - 1)s_a^2}{\sigma_a^2}$$

$$\frac{\chi^2_{a,\,\mathrm{df}=n\text{-}1}}{n_a - 1} = \frac{s_a^2}{\sigma_a^2} \qquad \text{and} \qquad \frac{\chi^2_{b,\,\mathrm{df}=n\text{-}1}}{n_b - 1} = \frac{s_b^2}{\sigma_b^2}$$

$$\frac{s_a^2}{\sigma_a^2} / \frac{s_b^2}{\sigma_b^2} = \frac{s_a^2}{s_b^2} \ \ (\text{if } \sigma_a^2 = \sigma_b^2) = \frac{\chi^2_{a,\,\mathrm{df}=n\text{-}1}}{n_a - 1} / \frac{\chi^2_{b,\,\mathrm{df}=n\text{-}1}}{n_b - 1} \sim F^{\,n_a - 1}_{\,n_b - 1} = F^{\,\nu_1}_{\,\nu_2}$$

We just derived the **F-Distribution**, named after Sir Ronald A. Fisher, a British mathematician and biologist who is credited with discovering p-values and ANOVA). The ratio of two sample variances (or the ratio of two chi-squares each divided by their degrees of freedom) is distributed as an F-distribution with degrees of freedom v1 and v2.:

$$\frac{s_a^2}{\sigma_a^2} / \frac{s_b^2}{\sigma_b^2} = \frac{\chi^2_{a,\,\mathrm{df}=n\text{-}1}}{n_a - 1} / \frac{\chi^2_{b,\,\mathrm{df}=n\text{-}1}}{n_b - 1} \sim F^{\,n_a - 1}_{\,n_b - 1} = F^{\,v_1}_{\,v_2}$$

To graph the F-distribution, you'd use: $F\left(\frac{\chi_1^2 / v_1}{\chi_2^2 / v_2}\right) = \frac{\Gamma \frac{v_1 + v_2}{2}}{\Gamma \frac{v_1}{2} \Gamma \frac{v_2}{2}} \left(\frac{v_1}{v_2}\right)^{\frac{v_1}{2}} \left(\frac{\chi_1^2 / v_1}{\chi_2^2 / v_2}\right)^{\frac{v_1}{2-1}} \left(1 + \frac{v_1}{v_2}\left(\frac{\chi_1^2 / v_1}{\chi_2^2 / v_2}\right)\right)^{\frac{-(v_1 - v_2)}{2}}$

Like we did with the z, t, and χ^2 distributions, we'll use a computer to get probabilities under the F-distribution.

Calculators: http://stattrek.com/online-calculator/f-distribution.aspx
http://lock5stat.com/statkey/theoretical_distribution/theoretical_distribution.html#F

Table: http://bradthiessen.com/html5/stats/m301/ftable.pdf

Key point: If we want to compare two variances, we take their ratio and compare it to an F-distribution.

31. Sketch an F-distribution with 20 degrees of freedom in the numerator and 10 degrees of freedom in the denominator. Calculate the following:

P(F < 1) = ________________ P(F > 3) = ________________ P(1 < F < 3) = ________________

32. Consider situation A from the first page of this hand-out. Conduct an F-test (at a 0.05 significance level) to compare the variances of the two blood pressure monitors.

Recall: The cheaper type of blood pressure monitor was tested 21 times and had a variance of 37.21 units.
The more expensive type was tested 21 times and yielded a variance of 9 units.

a) First, write out our null and alternate hypotheses. We'll assume the null hypothesis is true. Can it be?

H_0:

H_1:

b) Calculate the observed F-statistic (observed ratio of sample variances). To simplify things, we usually divide the larger variance by the smaller variance.

$$F_{v_2}^{v_1} = \frac{s_{\text{larger}}^2}{s_{\text{smaller}}^2} =$$

c) How many degrees of freedom will we have for our observed F-statistic?

d) This is our observed F-statistic for this particular pair of samples of size n=21. If we had selected a different random sample of 21 observations, we would have observed a different F-statistic. Suppose we could repeatedly go back in time, collect random samples of size n=21, and calculate the F-statistic for each of those samples. What would the distribution of all those potential F-statistics look like? Sketch it below and label its mean.

e) Remember that we do not know the population variances for our two types of blood pressure monitors. T Then how could we possibly calculate an F-statistic and use the F-distribution? Remember:

$$F_{n_b-1}^{n_a-1} = \frac{s_a^2}{\sigma_a^2} / \frac{s_b^2}{\sigma_b^2}$$

f) Find the p-value for our observed F-statistic, compare it to our alpha-level, and draw a conclusion.