14.30 Introduction to Statistical Methods in Economics
Spring 2009

Problem Set #9

14.30 - Intro. to Statistical Methods in Economics

Instructor: Konrad Menzel

Due: Friday, May 8, 2009

Question One: Confidence Intervals

(Adapted from Bain/Engelhardt p. 384)

Consider a random sample of size n from a normal distribution, $X_i \sim N(\mu, \sigma^2)$.

1. If it is known that $\sigma^2 = 15$, find a 90% confidence interval for μ based on the estimate $\bar{x} = 25.3$ with $n = 16$.

   - Solution to 1: We just plug in the components to the formula for the first of the "Important Cases" from the lecture notes:

     $$[A(X), B(X)] = \left[\hat{\theta} - \sqrt{\sigma^2_{\hat{\theta}}}\Phi^{-1}(1 - \frac{\alpha}{2}),\, \hat{\theta} + \sqrt{\sigma^2_{\hat{\theta}}}\Phi^{-1}(1 - \frac{\alpha}{2})\right]$$

     which gives, for $\hat{\theta} = \bar{x} = 25.3$, $\alpha = .10$, and $\sigma^2_{\hat{\theta}} = \frac{\sigma^2}{n} = \frac{15}{16}$

     $$\begin{aligned} [A(X), B(X)] &= \left[25.3 - \sqrt{\frac{15}{16}}\Phi^{-1}(.95),\, 25.3 + \sqrt{\frac{15}{16}}\Phi^{-1}(.95)\right] \\ \mu &\in [23.71, 26.89] \end{aligned}$$

     with a 90% probability, or the confidence interval covers the truth with a 90% probability (since μ isn't random–the CI is random).

2. Based on the information in (1), find a one-sided lower 90% confidence limit for μ. Also, find a one-sided upper 90% confidence limit for μ.

   - Solution to 2: A one-sided lower 90% confidence limit for μ and one-sided upper 90% confidence limit for μ correspond to a lower/upper bound on μ. We just adjust the probabilities for the upper and lower limit so that the 10% error is on the lower or upper end:

     $$\begin{aligned} [A(X), B(X)] &= \left[25.3 - \sqrt{\frac{15}{16}}\Phi^{-1}(.90),\, 25.3 + \sqrt{\frac{15}{16}}\Phi^{-1}(1)\right] \\ &= [24.06, \infty] \end{aligned}$$

and for the upper end

$$
\begin{aligned}
[A(X), B(X)] &= \left[25.3 - \sqrt{\frac{15}{16}}\Phi^{-1}(1),\ 25.3 + \sqrt{\frac{15}{16}}\Phi^{-1}(.90)\right] \\
&= [-\infty, 26.54] .
\end{aligned}
$$

Notice that the one-sided confidence intervals are actually shorter from $\bar{x}$ than the two-sided version, because we've allocated all of the error to just one side.

3. For a confidence interval of the form given by $(\bar{x} - \Phi^{-1}(1 - \frac{\alpha}{2})\frac{\sigma}{\sqrt{n}},\ \bar{x} + \Phi^{-1}(1 - \frac{\alpha}{2})\frac{\sigma}{\sqrt{n}})$, derive a formula for the sample size required to obtain an interval of specified length, λ. If $\sigma^2 = 9$, then what sample size is needed to achieve a 90% confidence interval of length 2?

   - Solution to 3: The length of a confidence interval is just the difference between the two bounds:

$$
\begin{aligned}
\lambda &= \left(\bar{x} + \Phi^{-1}(1 - \frac{\alpha}{2})\frac{\sigma}{\sqrt{n}}\right) - \left(\bar{x} - \Phi^{-1}(1 - \frac{\alpha}{2})\frac{\sigma}{\sqrt{n}}\right) \\
&= \Phi^{-1}(1 - \frac{\alpha}{2})\frac{\sigma}{\sqrt{n}} + \Phi^{-1}(1 - \frac{\alpha}{2})\frac{\sigma}{\sqrt{n}} \\
\lambda &= 2\frac{\sigma}{\sqrt{n}}\Phi^{-1}(1 - \frac{\alpha}{2}) \\
n &= \left(\frac{2\sigma}{\lambda}\Phi^{-1}(1 - \frac{\alpha}{2})\right)^2
\end{aligned}
$$

   If for $\sigma^2 = 9$, $\alpha = .10$, and $\lambda = 2$, we will need a sample size as follows:

$$
\begin{aligned}
n &= \left\lceil \left(\frac{2 \cdot 3}{2}\Phi^{-1}(.95)\right)^2 \right\rceil \\
n &= 25
\end{aligned}
$$

   So, we'll need a sample of size 25 in order to achieve a confidence interval of length 2 (or slightly less).

4. Suppose now that σ^2 is unknown. Find a 90% confidence interval for μ if $\bar{x} = 25.3$ and $s^2 = 15.21$ with $n = 16$.

   - Solution to 4: We now just need to use the formula which using the t distribution's quantiles:

$$
[A(X), B(X)] = \left[\hat{\theta} - \sqrt{\hat{S}^2}t_{n-1}^{-1}(1 - \frac{\alpha}{2}),\ \hat{\theta} + \sqrt{\hat{S}^2}t_{n-1}^{-1}(1 - \frac{\alpha}{2})\right]
$$

   where the $n - 1$ subscript denotes the degrees of freedom parameter of the t distribution. This gives, for $\hat{\theta} = \bar{x} = 25.3$, $\alpha = .10$, and $s^2 = 15.21$

$$
\begin{aligned}
[A(X), B(X)] &= \left[25.3 - \sqrt{\frac{15.21}{16}}t_{n-1}^{-1}(.95),\ 25.3 + \sqrt{\frac{15.21}{16}}t_{n-1}^{-1}(.95)\right] \\
\mu &\in [23.59, 27.01]
\end{aligned}
$$

Note that this 90% confidence interval is slightly wider due to two factors: 15.21>15 and the t distribution's critical values (quantiles) are wider than the Normal's. Most of the difference here is due to the t distribution's wider critical values.

5. Based on the data in (4), find a 99% confidence interval for σ^2. Also show for $n = 14$. (Hint: What is the distribution of $\frac{n \cdot s^2}{\sigma^2}$? It's a χ^2_{n-1}, which happens to approach the Normal distribution, but use the χ^2_{n-1} for this part.)

   - Solution to 5: To construct this confidence interval, we're going to use case 4 from the "Important Cases" in the lecture notes. We want to find constants a and b such that

     $$P_{\sigma^2}(a \leq s^2 \leq b) = F_{s^2}(b) - F_{s^2}(a) = 0.995 - 0.005 = 0.99.$$

     In other words, we're going to focus on a symmetric confidence interval. We just need to find a and b such that

     $$F_{s^2}(b) = 0.995 \, and \, F_{s^2}(a) = 0.005.$$

     So, what is the distribution of s^2? Well, we know what the distribution of $\frac{n \cdot s^2}{\sigma^2}$ is. It is a χ^2_{n-1}. So, what we need to do is go back to where we started and retransform the problem into things we know:

     $$P_{\sigma^2}(a \leq s^2 \leq b) = P_{\sigma^2}\left(\frac{n \cdot a}{\sigma^2} \leq \frac{n \cdot s^2}{\sigma^2} \leq \frac{n \cdot b}{\sigma^2}\right)$$

     which gives us the critical values for $\frac{n \cdot s^2}{\sigma^2}$. But what would then be a reasonable way to obtain the confidence interval for s^2? We can perform the following algebra:

     $$\begin{aligned} F^{-1}_{\chi^2_{n-1}}(.005) &= 4.60 = \frac{n \cdot a}{\sigma^2} \\ F^{-1}_{\chi^2_{n-1}}(.995) &= 32.80 = \frac{n \cdot b}{\sigma^2} \end{aligned}$$

     This suggests the following:

     $$\begin{aligned} P_{\sigma^2}(a \leq s^2 \leq b) &= P_{\sigma^2}\left(4.60 \cdot \frac{\sigma^2}{n} \leq s^2 \leq 32.80 \cdot \frac{\sigma^2}{n}\right) \\ &= P_{\sigma^2}\left(\frac{n \cdot s^2}{32.80} \leq \sigma^2 \leq \frac{n \cdot s^2}{4.60}\right) \\ &= P_{\sigma^2}\left(\frac{n}{32.80} \cdot s^2 \leq \sigma^2 \leq \frac{n}{4.60} \cdot s^2\right) \end{aligned}$$

     which gives us the confidence interval of

     $$[A(X), B(X)] = [7.42, 52.89]$$

     which is very asymmetric, since 15.21 is the point estimate. For $n = 14$ (see next problem), the confidence interval would be $[7.14, 59.73]$.

6. Now use the Normal approximation to get a 99% confidence interval for σ^2. Also show for $n = 14$. Does your estimate make sense for $n = 16$ and $n = 14$? Explain.

   - Solution to 6: The Normal approximation for the 99% confidence interval for σ^2 will involve just pretending that χ^2_{n-1} is approximately Normal, and constructing the confidence intervals from that approximation. If you remember from one of our previous problem sets, $\chi^2_k \to N(k, 2k)$. Using this fact, we simply construct the confidence intervals as before, but where we use the Normal approximation rather than the χ^2_{n-1} as follows:

$$P_{\sigma^2}(a \le s^2 \le b) = P_{\sigma^2}\left(\frac{n \cdot a}{\sigma^2} \le \frac{n \cdot s^2}{\sigma^2} \le \frac{n \cdot b}{\sigma^2}\right)$$

     which gives us the critical values for $\frac{n \cdot s^2}{\sigma^2}$. But what would then be a reasonable way to obtain the confidence interval for s^2? We can perform the following algebra:

$$\begin{aligned} F^{-1}(.005; n-1, 2(n-1)) &= 0.892 = \frac{n \cdot a}{\sigma^2} \\ F^{-1}(.995; n-1, 2(n-1)) &= 29.11 = \frac{n \cdot b}{\sigma^2} \end{aligned}$$

     This suggests the following:

$$\begin{aligned} P_{\sigma^2}(a \le s^2 \le b) &= P_{\sigma^2}\left(0.892 \cdot \frac{\sigma^2}{n} \le s^2 \le 29.11 \cdot \frac{\sigma^2}{n}\right) \\ &= P_{\sigma^2}\left(\frac{n \cdot s^2}{29.11} \le \sigma^2 \le \frac{n \cdot s^2}{0.892}\right) \\ &= P_{\sigma^2}\left(\frac{n}{29.11} \cdot s^2 \le \sigma^2 \le \frac{n}{0.892} \cdot s^2\right) \end{aligned}$$

     which gives us the confidence interval of

$$[A(X), B(X)] = [8.36, 272.95]$$

     which is very asymmetric, since 15.21 is the point estimate. Note the very wide right tail. In fact, if I had adjusted the numbers slightly, we would have ended up with a confidence interval which would have had a negative right tail, due to the poorness of the approximation. For $n = 14$, we get a confidence interval of $[8.15, -1586.7]$ which makes no sense at all since the upper endpoint isn't greater than the lower endpoint, and it is negative,which is impossible for a sum of squares. However, for larger values of n, the approximation works much better.

Question Two: Empirical Example

The following data are times (in hours) between failures of air conditioning equipment in a particular airplane: 74, 57, 48, 29, 502, 12, 70, 21, 29, 386, 59, 27, 153, 25, 232. Assume that the data are observed values of a random sample from an exponential distribution $X_i \sim EXP(\theta)$.

1. Find a 90% confidence interval for the mean time between failures, θ.

   - Solution to 1: Computing the sample mean of the observations, I get $\bar{x} = 114.93$. The parameter θ has a maximum likelihood estimator of $\hat{\theta} = \frac{1}{N}\sum_{i=1}^{N} x_i$, the sample mean. So, in order to construct a confidence interval for θ, we need to either use a Normal approximation ($n = 15$ isn't large, but it might be large enough–see the picture below from Question 1 from Problem Set #5 where we derived the finite sample distribution for $EXP(\theta = 1)$),

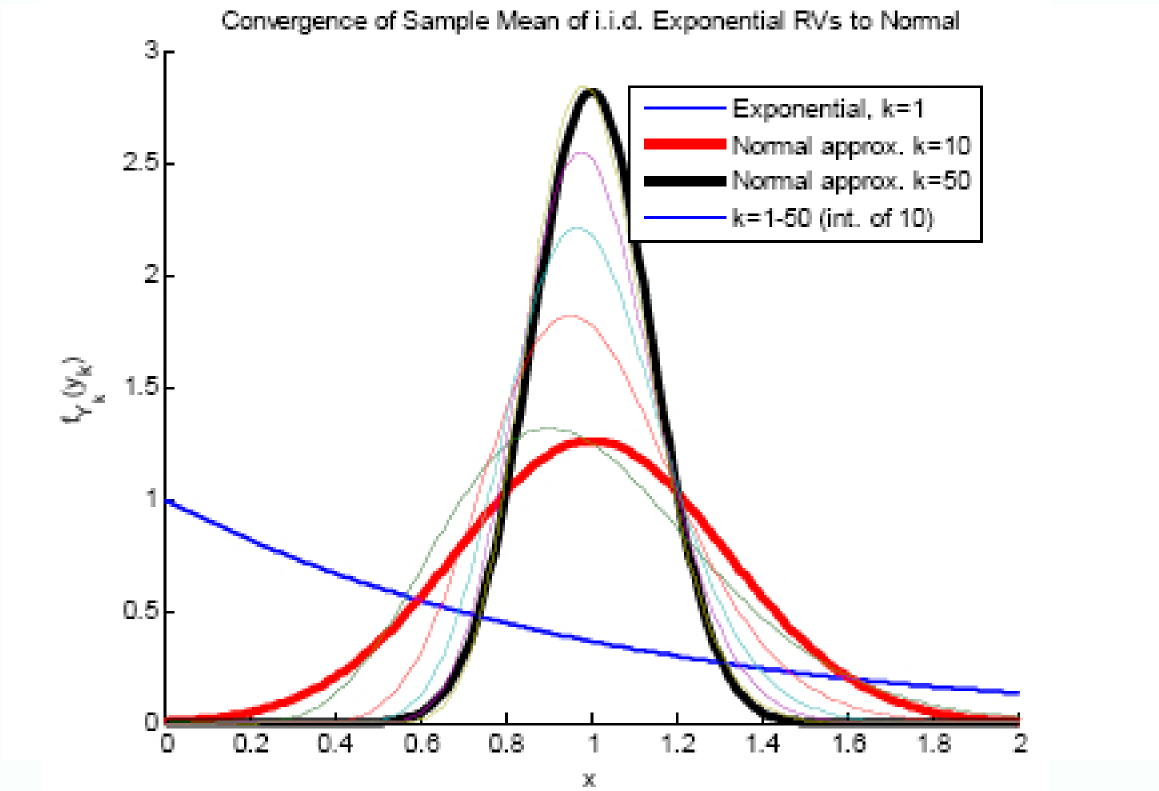

     or the finite sample distribution for the average of n exponentially distribued random variables. I will try the Normal approximation for the mean first, as it is much easier and fully nonparametric (doesn't rely on any distribution, as long as a CLT applies to invoke asymptotic normality):

$$[A(X), B(X)] = \left[\hat{\theta} - \sqrt{\hat{S}^2} t_{n-1}^{-1}(1 - \frac{\alpha}{2}),\ \hat{\theta} + \sqrt{\hat{S}^2} t_{n-1}^{-1}(1 - \frac{\alpha}{2})\right]$$

     where $\alpha = .10$, $\hat{\theta} = \bar{x} = 114.93$, $s^2 = 21,651$, and $n = 15$. The variance of $\hat{\theta}$ is $\frac{\sigma^2}{n}$.

$$\begin{aligned}[A(X), B(X)] &= \left[114.93 - \sqrt{\frac{21,651}{15}} t_{n-1}^{-1}(.95),\ 114.93 + \sqrt{\frac{21,651}{15}} t_{n-1}^{-1}(.95)\right] \\ &= [48.01, 181.85].\end{aligned}$$

     Now, I will attempt to use the finite sample distribution of the sum of exponential random variables. It turns out that the sum of exponential random variables is distributed as a $Gamma(n, \theta)$ (Source: Wikipedia: Exponential Distribution). So, if we want a confidence interval for the average, we'll just have to scale it back dividing by n after we get the confidence interval for the sum (or we could do a transformation of random variables from the start, but I'm going to focus on the sum instead). In particular, we have

$$P_\theta(a \leq \hat{\theta} \leq b) = P_\theta(a \cdot n \leq \hat{\theta} \cdot n \leq b \cdot n)$$

which has a $Gamma(n, \theta)$ distribution. So, we obtain the necessary quantiles of the $Gamma(n, \theta)$ distribution:

$$\begin{aligned} F_\Gamma^{-1}(.05; n, \theta) &= n \cdot a \\ F_\Gamma^{-1}(.95; n, \theta) &= n \cdot b \end{aligned}$$

where we take the inverse with respect to the argument, θ:

$$\begin{aligned} \frac{1}{n} F_\Gamma^{-1}(.05; n, \theta) &= a \\ \frac{1}{n} F_\Gamma^{-1}(.95; n, \theta) &= b \end{aligned}$$

which gives us the probability statement:

$$P_\theta(a \leq \hat{\theta} \leq b) = P_\theta(\frac{1}{n} F_\Gamma^{-1}(.05; n, \theta) \leq \hat{\theta} \leq \frac{1}{n} F_\Gamma^{-1}(.95; n, \theta))$$

which, using some mathematical software, we can invert the bounds by defining a function $G(\theta; \alpha) = \frac{1}{n} F_\Gamma^{-1}(\alpha; n, \theta)$:

$$P_\theta(G^{-1}(\hat{\theta}; 0.95) \leq \theta \leq G^{-1}(\hat{\theta}; 0.05)) = [78.77, 186.45].$$

2. Find a one-sided lower 95% confidence limit for the 10th percentile of the distribution of time between failures.

   - Solution to 2: Now that we have confidence intervals for θ, this is incredibly easy when we recognize that the 10th percentile of the distribution of time between failures is just a function of θ. Also, fortunately, it is a monotonic function, making our exercise even easier. Using the equivariance to monotone transformations property of quantiles, we just apply the following:

     $$P_\theta(a \leq \theta \leq b) = P_{F(\theta)}(F(a) \leq \theta \leq F(b))$$

     for monotone function $F(\cdot)$. We just need to make sure that the 10th percentile function for an exponential random variable is monotonic in θ:

     $$F(x; \theta) = 1 - e^{-\frac{x}{\theta}} = u$$

     where u is the quantile of the CDF. We get

     $$x = -\theta log\,(1 - u)$$

which is clearly monotonic in θ. So, we simply apply what we showed above that for a monotonic function, we just need to evaluate that function at the boundaries of the 90% confidence interval for θ:

$$F^{-1}(.10;\theta) \in [5.06, 19.16]$$

for the normal approximation of the confidence interval for θ or

$$F^{-1}(.10;\theta) \in [8.30, 19.65]$$

for the exact result using the Gamma distribution's quantiles for the 90% confidence interval for θ. For the one-sided lower bounds or one-sided 95% confidence interval, we just evaulate $F^{-1}(.10;\theta)$ at the lower endpoint of the confidence interval from part (1):

$$F^{-1}(.10;\theta) \in [5.06, \infty]$$

for the normal approximation of the one-sided confidence interval and

$$F^{-1}(.10;\theta) \in [8.30, \infty]$$

Question Three: Hypothesis Testing Concepts

1. Define the null, H_0, and alternative, H_a, hypotheses and explain their difference.

   - Solution for 1: The null hypothesis is the hypothesis to be tested, while the alternative hypothesis is the collection of other possible assumptions about the population other than the null. In particular, the null hypothesis is what we assert and try to statistically disprove or reject by gathering evidence or data. Sometimes we're not actually trying to disprove it, but rather we are trying to not accept nearby alternatives. For example, if we were trying to prove that the average return to a year of education is 9%, but our confidence intervals go from -20% and 38%, then while we fail to reject the null that the returns are 9%, we have insufficient evidence to reject other local alternatives such as 15% or -5%.

2. Write down the definition of a Type I error twice: first, mathematically and second, in words.

   - Solution to 2: From the notes, we have

     $$\alpha = P(\text{Type I Error}) = P(\text{reject}|H_0)$$

     or, in words, Type I error is the event that we reject the null hypothesis when it was actually true.

3. Now, write down the definition of a Type II error twice: first, mathematically and second, in words.

- Solution to 3: From the notes, again, we have

$$\beta = P(\text{Type II Error}) = P(\text{fail to reject}|H_a)$$

or, in words, Type II error is the event that we fail to reject the null hypothesis when the alternative is actually true.

4. Fill in where "Type I error" and "Type II error" fall in the following box:

	Reject H_0	*Fail to Reject* H_0
$H_0\, True$	Type I Error	*Good*!
$H_a\, True$	*Good*!	Type II Error

5. In the United States, the criminal justice system claims that a suspect is "innocent until proven guilty." Write down the equivalent null hypothesis and its complementary alternative for this statement. Further, define the Type I and Type II errors. Give your answer in a box as above in (4).

   - Solution to 5: The null hypothesis is H_0 : the defendant is innocent. The alternative is H_a : the defendant is guilty. The Type I error would be to prove that an innocent defendant is guilty and send him or her to jail. The Type II error would be to fail to convict a guilty defendant.

	Prove Guilty	*Fail to Prove Guilty*
Defendant is Innocent	Type I Error	*Good*!
Defendant is Guilty	*Good*!	Type II Error

Question Four: Power Curves

Suppose that a sample of size one is taken from the PDF $f_Y(y) = \frac{1}{\lambda}e^{-\frac{y}{\lambda}}$, with $y > 0$, for the purpose of testing

$$H_0 : \lambda = 1 \, versus \, H_1 : \lambda > 1.$$

The null hypothesis is rejected if $y > 3.20$.

1. Calculate the probability of committing a Type I error.

   - Solution to 1: The probability of committing a Type I error is $P(reject|H_0) = 1 - F(3.20|\lambda = 1) = 0.0408$.

2. Calculate the probability of committing a Type II error when $\lambda = 1,\ \frac{4}{3},\ \frac{3}{2},\ 2,\ \frac{5}{2},\ 3,\ or\, 3.2$.

- Solution to 2: The probability of committing a Type II error is $P(\text{fail to reject}|H_a) = F(3.20|\lambda = \lambda_a)$. For these various values we get the following results:

λ	β
1.00	0.9592
1.33	0.9093
1.50	0.8816
2.00	0.7981
2.50	0.7220
3.00	0.6558
3.20	0.6321

Here's the plot over the entire interval:

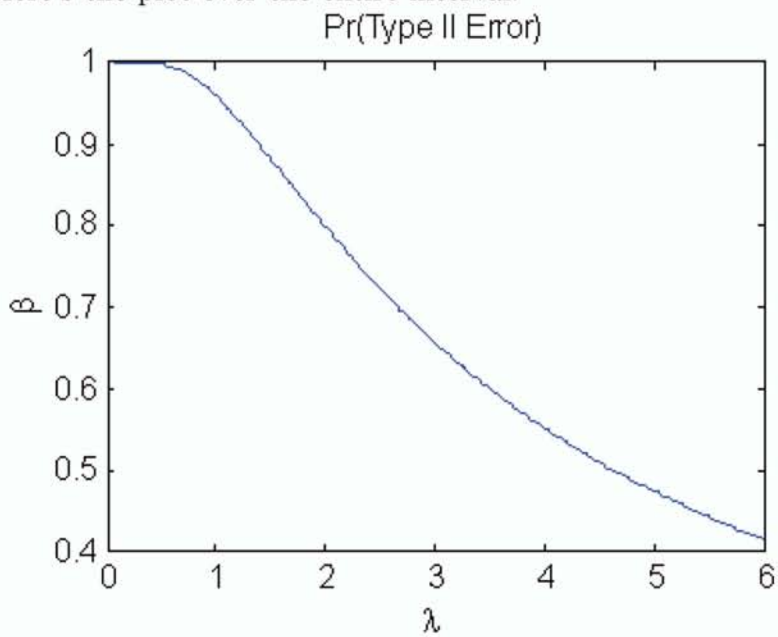

Note that since we have a one-sided hypothesis, the probability of a Type II error approaches 1 for $\lambda < 1$.

3. Let the Type II error be represented by β. Plot the values of $1 - \beta$ from (2) against λ and connect the dots. You just constructed the power curve for the test!

- Solution to 3: For this I'm just going to plot the curve:

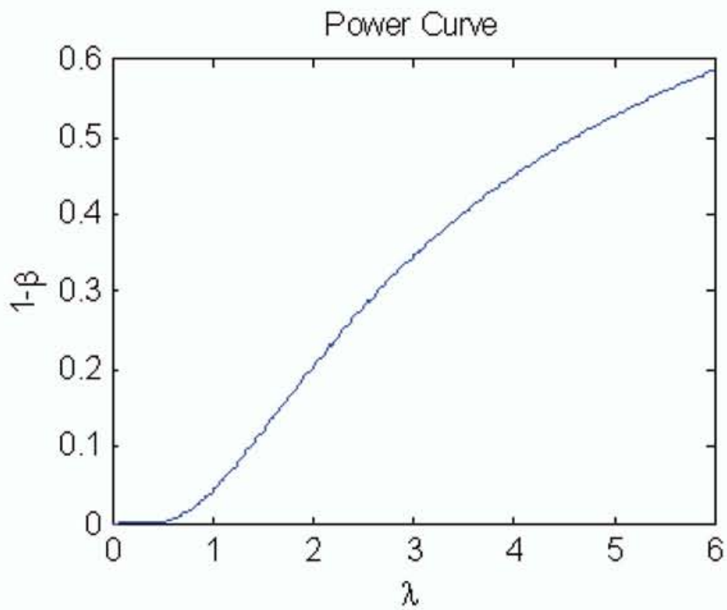

4. Sketch what you think the power curve for the test $H_0 : \lambda = 1 \, versus \, \lambda \neq 1$ will look like if we reject the null hypothesis in the event that $y \notin [0.25, 1.75]$. In particular, will it be symmetric or skewed? Compute the probability of a Type II error if $\lambda = 0.25$ and 1.75 to corroborate your answer.

   - Solution to 4: If we reject for y outside of the interval $[0.25, 1.75]$, then the Type II error probability is $F(1.75) - F(0.25)$. We should expect this to be skewed, since the exponential distribution is skewed.

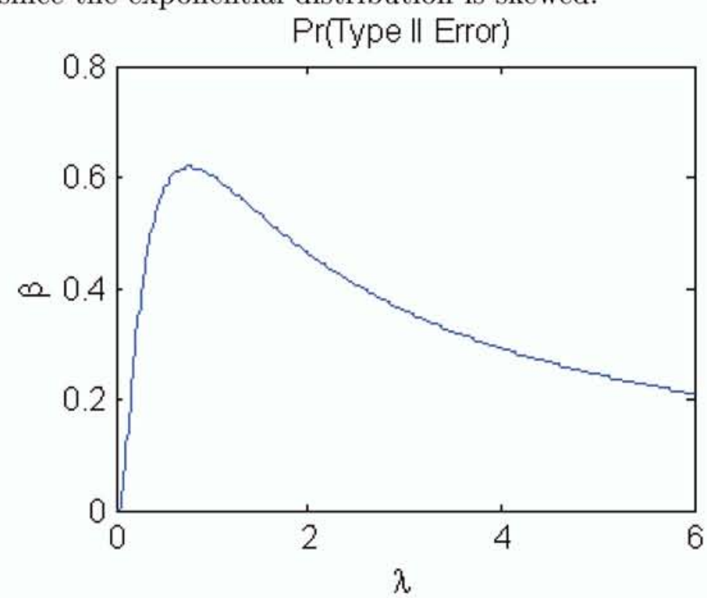

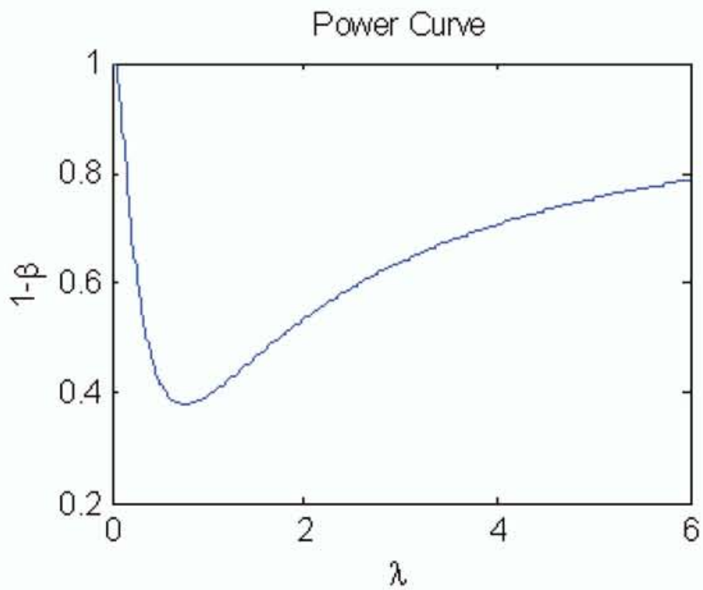

Question Five: Application of Hypothesis Testing

Commercial fisherman working certain parts of the Atlantic Ocean sometimes find their efforts being hindered by the presence of whales. Ideally, they would like to scare away the whales without frightening the fish. One of the strategies being experiment with is to transmit underwater the sounds of a killer whale. On the 52 occasions that this technique has been tried, it worked 24 times (that is, the whales got lost). Experience has shown, though, that 40% of all whales sighted near fishing boats leave of their own accord anyway, probably just to get away from the noise of the boat and the bad smell.

1. Let p be the probability that a whale leaves the area after hearing the sounds of a killer whale. Test $H_0 : p = 0.40$ versus $H_1 : p > 0.40$ at the $\alpha = 0.05$ level of significance. Can it be argued on the basis of these data that transmitting underwater predator sounds it an effective technique for clearing fishing waters of unwanted whales?

   - Solution to 1: We just want to test whether $\hat{p} = \frac{24}{52}$ is significantly different from 0.40. We can do this two different ways, either by computation of a p-value or construction of a 95% confidence interval. We will first compute a p-value by looking at the likelihood of getting something at least as extreme under the null:

     $$p - value = 1 - F(23|H_0) = 1 - \sum_{i=0}^{52} 0.4^i (1-0.4)^{52-i} = 0.2213.$$

     Since the p-value 0.2213>0.05, we don't have sufficient evidence to reject the null. For $\alpha = .25$ we could reject, but this is not a very stringent requirement.
     Using a binomial or normal approximation, we can construct a one-sided 95% confidence interval:

     $$p \ \in \ \left[\frac{24}{52} - 1.645 \cdot \sqrt{\frac{\frac{24}{52}(1-\frac{24}{52})}{52}}, 1\right]$$

$$p \in [0.3478, 1]$$

if we use the Normal distribution's critical values, but

$$\begin{aligned} p &\in \left[\frac{24}{52} - 1.6753 \cdot \sqrt{\frac{\frac{24}{52}(1-\frac{24}{52})}{52}}, 1\right] \\ p &\in [0.3457, 1] \end{aligned}$$

if we use the t_{52-1} critical values, which produces a small, unqualitatively important difference. Thus, regardless of the approximation, our confidence interval covers the null hypothesis, so we fail to reject. The p-values for the null hypothesis using the Normal and t-distribution approximations are 0.1867 and 0.1888, respectively. So, we are unable to say whether the underwater transmission of killer whale sounds has any effect. Maybe we should keep experimenting. We at least know that it isn't a perfect problem solver, as two-sided 95% confidence intervals go basically all the way to 0.60:

$$\begin{aligned} p &\in \left[\frac{24}{52} - 1.96 \cdot \sqrt{\frac{\frac{24}{52}(1-\frac{24}{52})}{52}}, \frac{24}{52} + 1.96 \cdot \sqrt{\frac{\frac{24}{52}(1-\frac{24}{52})}{52}}\right] \\ p &\in [0.3260, 0.5970] \end{aligned}$$

if we use the Normal distribution's critical values, but

$$\begin{aligned} p &\in \left[\frac{24}{52} - 2.01 \cdot \sqrt{\frac{\frac{24}{52}(1-\frac{24}{52})}{52}}, \frac{24}{52} + 2.01 \cdot \sqrt{\frac{\frac{24}{52}(1-\frac{24}{52})}{52}}\right] \\ p &\in [0.3457, 0.6003] \end{aligned}$$

if we use the t_{52-1} critical values.

2. Calculate the $p - values$ for these data. For what values of α would H_0 be rejected?

    - Solution to 2: See part 1.

Question Six: One v. Two Sided Hypotheses; Sample Variance

Suppose that babies' weights at birth are normally distributed with a mean of 7 pounds and a variance of 1 pound. A particular obstetrician is suspected of giving pregnant women poor advice on diet, which would cause babies to be 1 pound lighter on average (but still have the same variance). You observe the weight of $n = 10$ babies that he delivers. The mean weight of the 10 babies is 6.2 pounds.

1. Suppose you want to do a test of the null hypothesis that the obstetrician is not giving poor advice against the alternative hypothesis that he is. Write down the null hypothesis and alternative hypothesis mathematically.

- Solution to 1: The null hypothesis that we want to test is whether the babies that this particular obstetrician delivers are the same weight:

$$H_0 : \mu \geq 7$$
$$H_a : \mu < 7$$

Alternatively, we might write that the hypotheses are that

$$H_0 : \mu = 7$$
$$H_a : \mu = 6$$

which would be a simple alternative hypothesis (also implicitly assuming Normality of birth weights–i.e. the doctor's advice only shifts the mean down to 6).

2. Perform a 5% test of the null hypothesis.

- Solution to 2: The probability of deliver 10 babies with an average weight of 6.2 or less under the null hypothesis is

$$\Phi\left(\frac{6.2-7}{\frac{1}{\sqrt{10}}}\right) = 0.0057.$$

This is our p-value, which is far less than the 5% level that we required to reject the null. Thus, we conclude that this doctor is giving bad advice.
If we used a simple hypothesis, we would just invoke the Neyman-Pearson Lemma and compute the k statistic, which is just a likelihood ratio test. This just involves evaluating

$$-2logT(\bar{x}) = -2log\left(\frac{f_0(\bar{x})}{f_a(\bar{x})}\right)$$

and comparing it to the χ^2_1 distribution's critical values. When we perform the test, using the Normal distribution for the null and alternative, but with different means, we obtain $-2logT(\bar{x}) = 6.00$ which is larger than the χ^2_1's 5% critical value of 3.84. So, we reject the null hypothesis in favor of the simple alternative.

3. Suppose you only knew the mean of the distribution of babies' weight, not the variance, but you did have an estimate of the variance, $s^2 = 1.5$. Perform a 5% test of the null against the alternative.

- Solution to 3: Again, we perform virtually the same test, but using the quantiles of the t distribution:

$$F_t\left(\frac{6.2-7}{\frac{\sqrt{1.5}}{\sqrt{10}}}; 10-1\right) = 0.0344.$$

Thus, at the 5% level, we still reach the same conclusion, although our p-value is much larger now (6 times larger).
We can perform this test using the simple alternative, although it becomes a little bit more complicated, since the ration of t distributions won't necessarily have the nice properties of the likelihood ratio being χ^2_1, except asymptotically. But, we use it anyway and compute the likelihood ratio:

$$-2log\left(\frac{t_9\left(\frac{6.2-7}{\frac{\sqrt{1.5}}{\sqrt{10}}}\right)}{t_9\left(\frac{6.2-6}{\frac{\sqrt{1.5}}{\sqrt{10}}}\right)}\right) = 3.59$$

which is less than the 5% critical value for the χ^2_1 distribution of 3.84, so, again we fail to reject the null in favor of the simple alternative (which shouldn't be very surprising, as the estimate of our variance is larger).

4. Would any of your answers above change if your alternative hypothesis was that the obstetrician was doing something to affect babies' weights, either negatively or positively? If so, adjust your p-values and conclusions accordingly.

- Solution to 4: If we instead had the alternative that he was doing something abnormal which could be represented as $H_a : \mu \neq 7$, we would then have to adjust our test. Instead, our p-values would all double (to account for the possibility that we could have gotten something as extreme in the upper tail that would've rejected that he's doing something to influence birthweights. In particular, our p-values would have become 0.0114 for part (2), which would not have affected our conclusion, but or 0.688 for part (3), which would have affected our conclusion–we would have failed to reject.
This problem does not lend itself well to a simple alternative. Would you define it as $\mu = 8$? The non-simple alternative that $\mu \neq 7$ seems more reasonable for two-sided testing.