Confidence Interval and z-test Worksheet

Confidence intervals (CI's) and z-tests are two related procedures we use when testing hypotheses. Below is the procedure of testing hypotheses.

Starting matter

Here's what you have at the beginning of a hypothesis test*.

1. A belief about the population mean
2. A known population standard deviation (known parameter)
3. An underlying normal distribution
4. A simple random sample

What are our hypotheses?

After you have the above, you can either be interested in

1. One-sided test
   - Your belief about the population mean is that it is either lower or higher than some value
   - There exists a competing belief that the population mean may point in a different direction than your belief (i.e. instead of greater than as in your belief, it may be actually less than)

$$H_0 : \mu \leq 10 \textbf{ versus } H_1 : \mu > 10$$

2. Two-sided test
   - Your belief about the population mean is that it equals some value exactly
   - There exists a competing belief that the population mean may be greater or less than that exact point

$$H_0 : \mu = 10 \textbf{ versus } H_1 : \mu \neq 10$$

Assumptions

This worksheet focuses on two methods, using a CI and using a z-test. These methods both require the following assumptions to be true.

1. You have a simple random sample
2. That simple random sample was taken from a distribution that is normally distributed
3. You know your population standard deviation

You will be expected to check for these assumptions before continuing on to create a CI or to use a z-test.

*You have these at the beginning of a hypothesis test that can be tested with a normal-based confidence interval or z-test.

Confidence Interval and z-test Worksheet

For both methods, you'll need to calculate the sample mean $\bar{x}$ and the standard error $\textbf{SE} = \dfrac{\sigma}{\sqrt{n}}$.

You may set a significance level. Most of the time this is 0.05.

In order to make a **confidence interval**, follow these steps.

1. Figure out your confidence level (this will probably be given to you). In most situations, the confidence level will be 95%.

Conf	90%	95%	99%
z*	1.645	1.96	2.58

2. Evaluate $\bar{x} \pm z^{*}\textbf{SE}$ by using the sample. This will create an interval of this form:

$$(\bar{x} - z^{*}\textbf{SE}, \bar{x} + z^{*}\textbf{SE})$$

although sometimes it can be the other way around. The lower number will be your lower bound and the greater number will be your upper bound.

In order to conduct a **z-test**, follow these steps.

1. Create your z-statistic by using

$$z = \frac{\bar{x} - \mu_0}{\textbf{SE}}$$ <- read as mu-not

where mu-not is the hypothesized true parameter.

2. Calculate your p-value depending on your hypotheses. We have that z is distributed standard normal, so we can use R to calculate the p-value (the probability of observing our data or more extreme given that the true population mean is the null value.

See the following curves and notice two hypothetical sample means.

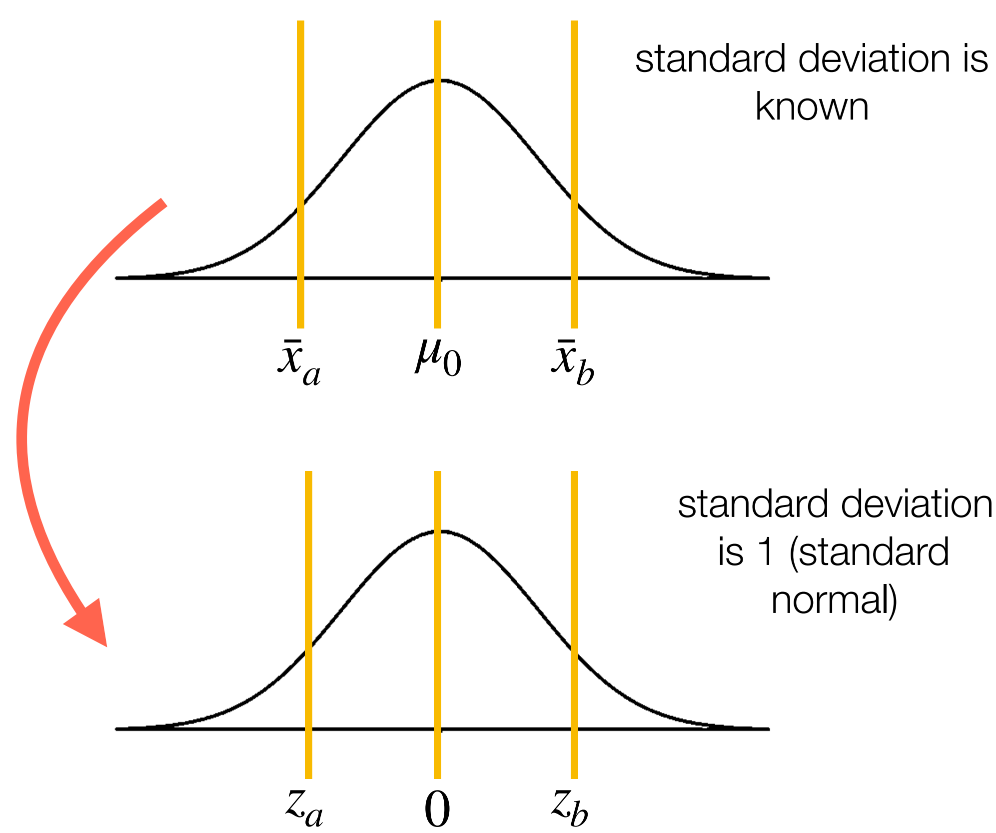

Figure A

This figure shows how we can calculate a p-value given the sample mean we calculated for our data. We use what is called "standardization" to get a corresponding value for our sample mean on the standard normal curve (i.e. the prob to left of observation is the same as each other).

Using R software, it is possible to simply use pnorm(x_bar_a, mu_not, sigma), but for pedagogical purposes, we learn about z-scores.

For these hypothetical sample mean observations, I would write this R code to calculate the p-value. You multiply by 2 if you're testing a two-sided hypothesis.

```
pnorm(za, mean=0, sd=1)
1-pnorm(zb, mean=0, sd=1)
```

Interpreting your confidence interval
State a sentence like this.

My confidence interval for the true mean is [state your interval]. This interval was created using a method that produces confidence intervals that contain the true parameter [confidence level %] of the time.

Concluding based on your confidence interval

1. If the hypothesized mean mu-not is not in our CI, then we reject our null hypothesis

2. If the hypothesized mean mu-not is indeed in our CI, then we fail to reject the null hypothesis.

Interpreting your p-value
State a sentence like this.

The probability of observing my data or more extreme is [p-value in percent] given that the true mean is [mu-not].

Concluding based on your p-value

1. If our p-value is very small (most of the time, we consider this to be less than 5%), then we reject the null hypothesis.

2. If the p-value is greater than 5%, then we fail to reject the null hypothesis.

Consult your course notes to understand and validate cut-offs.

If you have set a significance level, then if your p-value is less than your significance level, reject your null hypothesis. Otherwise, fail to reject.

For both methods, be able to interpret in context what rejecting or failing to reject mean in context.

Practice

Bill Nye, the "science guy" and Electric Daisy Carnival (EDC) 2019's opening ceremony host *(YES, SERIOUSLY!)*, is interested in whether his viewership has changed since his shows dropped on Netflix. He believes that about 15 minutes of his shows are shown in classrooms around the globe to demonstrate scientific theory. The true standard deviation of watching time is 15 minutes. To test his hypothesis of whether or not his shows are being watched on average for about 15 minutes in classrooms worldwide, his team took a random sample from teachers asking how many minutes of his show they presented in their classrooms. He retrieved this data:

```
90 30 120 30 0 90 120 0
```

Assume that the true underlying distribution of minutes watched of Bill Nye is normal. Use (a) a 95% confidence interval and (b) a z-test to test the relevant hypotheses.

Hypotheses
To figure out which type of hypotheses I'm going to write, I'll refer to these excerpts from the prompt.

"interested in whether his viewership has changed since his shows dropped on Netflix"
"whether or not his shows are being watched on average for about 15 minutes"

We can deduce this is a two-sided test because Bill Nye doesn't explicitly state a belief of "greater" or "less" minutes given the Netflix drop of his shows.

Assumptions
- These data are stated to be random.
- We are assuming normality as stated in the prompt.
- We know the true standard deviation is 15 minutes, also stated in the prompt.

Calculating the sample mean
We can calculate the sample mean to be the sum of the data

```
90 30 120 30 0 90 120 0
```

divided by 8. Thus, the sample mean is 60.

Calculating the standard error
The standard error is the true standard deviation divided by the square root of our sample size. We calculate this to be about 5.3.

Confidence Interval
We're going to make a 95% CI, so our critical value (z*) is 1.96. Thus, our CI will be (60-1.96(5.3), 60+1.96(5.3)) = (49.612, 70.388). Our hypothesized value of mu, 15, is not in this confidence interval. We reject the null hypothesis that the average minutes that Bill Nye was watched in classrooms was 15.

z-test
We can use the formula on page 2 to calculate z=(60-15)/5.3=8.49. On a standard normal curve, we know that only 0.3% of the data lie beyond 3 standard deviations from the mean. Therefore, we know that there is a very small p-value associated with this number. We can check by using

```
2*(1-pnorm(8.49))
```

The value that outputs from using the above code is the p-value. Given that the true minutes watched of Bill Nye around the globe in classrooms is 15 minutes, we would with almost 0% probability expect to collect the data that we did. We reject the null hypothesis because this probability is so small.