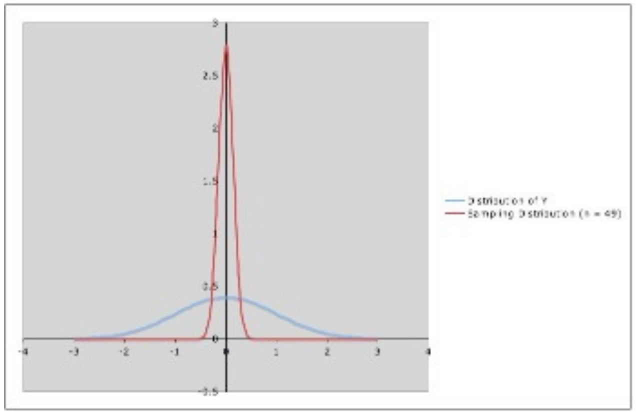

The following chart summarizes which model assumptions are necessary to prove which part of the theorem:

	Conclusions ***about Sampling Distribution*** (Distribution of $\overline{Y}_n$)		
	1: Normal	2: Mean μ	3: Standard deviation $\sigma/\sqrt{n}$
Assumption 1 (Y normal)	√		
Assumption 2 (simple random samples)	√		√

(Note that the conclusion that the sampling distribution $\overline{Y}_n$ has the same mean as Y does not involve either of the model assumptions.)

- The conclusions of the theorem will allow us to do the following:
  - If we specify a probability (we'll use .95 to illustrate), we can find a number *a* so that

    (*) The probability that $\overline{Y}_n$ lies between μ - a and μ + a is approximately 0.95.

    Caution: It is important to get the *reference category* straight here. This amounts to keeping in mind what is a random variable and what is a constant. Here, $\overline{Y}_n$ *is the random variable* (that is, the *sample* is varying), whereas μ *is constant*.

    Note: The z-procedure for confidence intervals is only an approximate procedure; that is why the "approximately" is in (*) and below. Many procedures are "exact"; we don't need the "approximately" for them.

- A little algebraic manipulation allows us to restate (*) as

(**) The probability that μ lies between $\overline{Y}_n$ - a and $\overline{Y}_n$ + a is approximately 0.95

Caution: It is again important to get the reference category correct here. It hasn't changed: it is still the *sample* that is varying, *not* μ. So the probability refers to $\overline{Y}_n$, *not* to μ.

Thinking that the probability refers to μ *is a* **common mistake** *in interpreting confidence intervals*.

It may help to restate (**) as:

(***) The probability that the interval from $\overline{Y}_n$ - a to $\overline{Y}_n$ + a contains μ is approximately 0.95.

- We are now faced with two possibilities (assuming the model assumptions are indeed all true):

1) The sample we have taken is one of the approximately 95% for which the interval from $\overline{Y}_n$- a to $\overline{Y}_n$+ a *does* contain μ. ☺

2) Our sample is one of the approximately 5% for which the interval from $\overline{Y}_n$ - a to $\overline{Y}_n$ + a does *not* contain μ. ☹

- Unfortunately, *we can't know which of these two possibilities is true*. ☹

- Since this is the best we can do, we calculate the values of $\bar{Y}_n$ - a and $\bar{Y}_n$ + a *for the sample we have*, and call the resulting interval an approximate *95% confidence interval for* μ.
  - We ***can*** say that *we have obtained the confidence interval by using a procedure that, for approximately 95% of all simple random samples from Y, of the given size, produces an interval containing the parameter we are estimating*.
  - Unfortunately, we ***can't*** know whether or not the sample we have used is one of the approximately 95% of "good" samples that yield a confidence interval containing the true mean μ, or whether the sample we have is one of the approximately 5% of "bad" samples that yield a confidence interval that does not contain the true mean μ.
  - We can just say that we have used a procedure that "works" about 95% of the time.
  - Various web demos can demonstrate.

In general: We can follow a similar procedure for many other situations to obtain confidence intervals for parameters.

- Each type of confidence interval procedure has its own model assumptions.
  - *If the model assumptions are not true, we can't be sure that the procedure does what is claimed.*
  - However, some procedures are *robust* to some degree to some departures from models assumptions -- i.e., the procedure works pretty closely to what is intended if the model assumption is not too far from true.
  - Robustness depends on the particular procedure; there are no "one size fits all" rules.

- We can decide on the "level of confidence" we want;
  - E.g., we can choose 90%, 99%, etc. rather than 95%.
  - Just which level of confidence is appropriate depends on the circumstances. (More later)
- The *confidence level* is *the percentage of samples for which the procedure results in an interval containing the true parameter*. (Or approximate percentage, if the procedure is not exact.)
- However, *a higher level of confidence will produce a wider confidence interval*. (See demo)
  - i.e., less certainty in our estimate.
  - So *there is a trade-off between level of confidence and degree of certainty*.

- Sometimes the best we can do is a procedure that only gives *approximate* confidence intervals.
  - i.e., the sampling distribution can be described only approximately.
  - i.e., there is one more source of uncertainty.
  - This is the case for the large-sample z-procedure.

- If the sampling distribution is not symmetric, we can't expect the confidence interval to be symmetric around the estimate.
  - In this case, there might be more than one reasonable procedure for calculating the endpoints of the confidence interval.

- There are variations such as "upper confidence limits" or "lower confidence limits" where we are only interested in estimating how large or how small the estimate might be.

V. MORE ON FREQUENTIST HYPOTHESIS TESTS

We'll now continue the discussion of hypothesis tests.

Recall: Most commonly used frequentist hypothesis tests involve the following elements:

1. Model assumptions
2. Null and alternative hypothesis
3. A test statistic (something calculated by a rule from a sample)
   - This needs to have the property that extreme values of the test statistic cast doubt on the null hypothesis.
   - The test statistic will have a certain sampling distribution.
4. A mathematical theorem saying, "If the model assumptions and the null hypothesis are both true, then the sampling distribution of the test statistic has this particular form."

The exact details of these four elements will depend on the particular hypothesis test.

Illustration: One-sided t-test for a Sample Mean

In this situation, the four elements above are:

1. *Model assumptions*:
   - The random variable Y is *normally distributed.*
   - Samples are *simple random samples*.

2. *Null and alternate hypotheses:*
   - *Null hypothesis*: The population mean μ of the random variable *Y is μ_0.*
   - *Alternative hypothesis*: The population mean μ of the random variable Y is greater than μ_0. (i.e., $\mu > \mu_0$)

3. *Test statistic:* For a simple random sample $y_1, y_2, \ldots, y_n$ of size n, we define the *t-statistic as*

$$t = \frac{\bar{y} - \mu_0}{s/\sqrt{n}},$$

where

$$\bar{y} = (y_1 + y_2 + \ldots + y_n)/n \quad \text{(sample mean)},$$

and

$$s = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(\bar{x} - x_i)^2} \quad \text{(sample standard deviation)}$$

The *sampling distribution* for this test is then the distribution of the random variable T_n defined by random process and calculation,

> "Randomly choose a simple random sample of size n and calculate the t-statistic for that sample."

4. The mathematical theorem associated with this inference procedure (one-sided t-test for population mean) says:

> If *the model assumptions are true <u>and</u> the null hypothesis is true*, then the sampling distribution is the *t-distribution with n degrees of freedom*.

(For large values of n, the t-distribution looks very much like the standard normal distribution; but as n gets smaller, the peak gets slightly smaller and the tails go further out.)

The reasoning behind the hypothesis test uses the sampling distribution *and* the value of the test statistic for the sample that has actually been collected (the actual data).

1. First, *calculate the t-statistic for the data*
2. Then *consider where the t-statistic for the data at hand lies on the sampling distribution*. Two possible values are shown in red and green, respectively, in the diagram below.
   - *Remember that this picture depends on the validity of the model assumptions and on the assumption that the null hypothesis is true.*

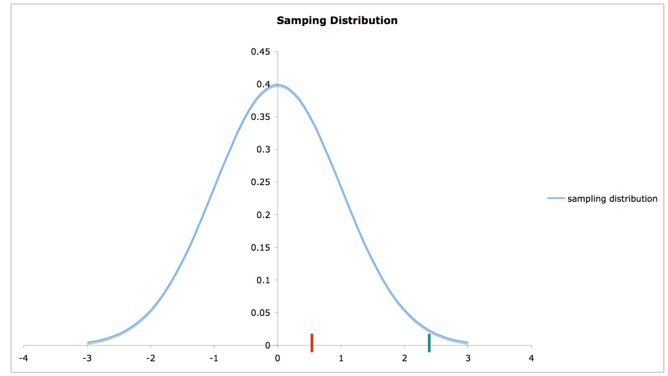

Case 1: If the t-statistic lies at the *red* bar (around 0.5) in the picture, *nothing is unusual*; our data are consistent with the null hypothesis.

Case 2: If the t-statistic lies at the *green* bar (around 2.5), then the data would be fairly *unusual* -- assuming the null hypothesis is true.

So ***a t-statistic at the green bar would cast some reasonable doubt on the null hypothesis***.

A t-statistic even further to the right would cast even more doubt on the null hypothesis.

Note: A little algebra will show that if t = $\dfrac{\bar{y} - \mu_0}{s/\sqrt{n}}$ is unusually large, then so is $\bar{y}$, and vice-versa

p-Values

The *p-value* is a quantitative measure of how unusual a particular test statistic is, with lower p-values indicating more unusual data. The general definition is:

> *p-value* = the probability of obtaining a test statistic *at least as extreme* as the one from the data at hand, assuming the model assumptions and the null hypothesis are all true.

The interpretation of "at least as extreme as" depends on the *alternative* hypothesis.

- For the one-sided alternative hypothesis $\mu > \mu_0$ (as in our example), "at least as extreme as" means "at least as great as".
    - Recalling that the probability of a random variable lying in a certain region is the area under the probability distribution curve over that region, we conclude that for this alternative hypothesis, the p-value is the area under the distribution curve to the *right* of the test statistic calculated from the data.
    - Note that, in the picture, the p-value for the t-statistic at the green bar is much less than that for the t-statistic at the red bar.

- Similarly, for the *other* one-sided alternative, $\mu < \mu_0$, the p-value is the area under the distribution curve to the *left* of the calculated test statistic.
    - Note that for this alternative hypothesis, the p-value for the t-statistic at the green bar would be much greater than the t-statistic at the red bar, but both would be large as p-values go.
- For the *two-sided alternative* $\mu \neq \mu_0$, the p-value would be the area under the curve to the right of the absolute value of the calculated t-statistic, *plus* the area under the curve to the left of the negative of the absolute value of the calculated t-statistic.
    - Since the sampling distribution in the illustration is symmetric about zero, the two-sided p-value of, say the green value, would be twice the area under the curve to the right of the green bar.

Recall that in the sampling distribution, we are only considering samples

- from the *same random variable*,
- that *fit the model assumptions* and
- *of the same size as the one we have.*

So if we spelling everything out, the definition of p-value reads:

p-value = the probability of obtaining a test statistic *at least as extreme* as the one from the data at hand, *assuming*

- *the model assumptions are all true,* and
- *the null hypothesis is true,* and
- *the random variable is the same (including the same population),* and
- *the sample size is the same.*

Comment: The preceding discussion can be summarized as follows:

If we obtain an unusually small p-value, then (at least) one of the following must be true:

- At least one of the model assumptions is not true (in which case the test may be inappropriate).
- The null hypothesis is false.
- The sample we have obtained happens to be one of the small percentage that result in an unusually small p-value.

Thus, if the p-value is small enough *and* all the model assumptions are met, then *rejecting the null hypothesis in favor of the alternate hypothesis* can be considered a rational decision, based on the evidence of the data used.

Comments:

1. How small is "small enough" is a judgment call.

2. "Rejecting the null hypothesis" does *not* mean the null hypothesis is false or that the alternate hypothesis is true. (Why?)

MISINTERPRETATIONS AND MISUSES OF P-VALUES

Recall:

p-value = the probability of obtaining a test statistic at least as extreme as the one from the data at hand, *assuming*:

- the model assumptions for the inference procedure used are all true*, and*
- the null hypothesis is true*, and*
- the random variable is the same (including the same population)*, and*
- the sample size is the same.

Notice that this is a *conditional probability*: The probability that something happens, given that various other conditions hold. *One* **common mistake** *is to neglect some or all of the conditions.*

Example A: Researcher 1 conducts a clinical trial to test a drug for a certain medical condition on 30 patients all having that condition.

- The patients are randomly assigned to either the drug or a look-alike placebo (15 each).
- Neither patients nor medical personnel know which patient takes which drug.
- Treatment is exactly the same for both groups, except for whether the drug or placebo is used.
- The hypothesis test has null hypothesis "proportion improving on the drug is the same as proportion improving on the placebo" and alternate hypothesis "proportion improving on the drug is greater than proportion improving on the placebo."
- The resulting p-value is p = 0.15.

Researcher 2 does another clinical trial on the *same drug*, with the *same placebo*, and *everything else the same* <u>*except*</u> that 200 patients are randomized to the treatments, with 100 in each group. The same hypothesis test is conducted with the new data, and the resulting p-value is p = 0.03.

Are these results contradictory? No -- *since the sample sizes are different, the p-values are not comparable, even though everything else is the same.*

Indeed, *a larger sample size typically results in a smaller p-value*.

> The idea of why this is true is illustrated by the case of the z-test, since large n gives a smaller standard deviation of the sampling distribution, hence a narrower sampling distribution.

Comparing p-values for samples of different size is a **common mistake**.

Example B: Researcher 2 from Example A does everything as described above, but for convenience, his patients are all from the student health center of the prestigious university where he works.

- He *cannot* claim that his result applies to patients other than those of the age and socio-economic background, etc. of the ones he used in the study, because his sample was taken from a smaller *population*.

Example C: Researcher 2 proceeds as in Example A, with a sample carefully selected from the population to which he wishes to apply his results, but he is testing for equality of the means of an outcome variable for the two groups.

- The hypothesis test he uses requires that the variance of the outcome variable for each group compared is the same.
- He doesn't check this, and in fact the variance for the treatment group is twenty times as large as the variance for the placebo group.
- He is *not* justified in rejecting the null hypothesis of equal means, no matter how small his p-value (unless by some miracle the statistical test used is robust to such large departures from the model assumption of equality of variances.)

Another **common misunderstanding** *of p-values is the belief that the p-value is "the probability that the null hypothesis is true".*

- This is essentially a case of confusing a conditional probability with the reverse conditional probability: In the definition of p-value, "the null hypothesis is true" is the condition, not the event.
- The basic assumption of frequentist hypothesis testing is that the null hypothesis is either true (in which case the probability that it is true is 1) or false (in which case the probability that it is true is 0) – so unless $p = 0$ or 1, the p-value couldn't possibly be the probability that the null hypothesis is true.

Note: In the *Bayesian* perspective, it makes sense to consider "the probability that the null hypothesis is true" as having values other than 0 or 1.

- In that perspective, we consider "states of nature;" in different states of nature, the null hypothesis may have different probabilities of being true.
- The goal is then to determine the probability that the null hypothesis is true, given the data.
- This is the *reverse conditional probability* from the one considered in frequentist inference (the probability of the data given that the null hypothesis is true).