Introductory Econometrics

Lecture 9: Hypothesis testing for linear regression models

Instructor: Ma, Jun

Renmin University of China

October 18, 2022

Null and alternative hypotheses

- Usually, we have two competing hypotheses, and we want to draw a conclusion, based on the data, as to which of the hypotheses is true.
- Null hypothesis , denoted as H_0: A hypothesis that is held to be true, unless the data provides a sufficient evidence against it.
- Alternative hypothesis, denoted as H_1: A hypothesis against which the null is tested. It is held to be true if the null is found false.
- The two hypotheses are not treated symmetrically. H_0 is favored in the sense that it is only rejected if there is strong evidence against it.
- H_0 summarizes our prior belief about the state of the world.
- H_1 corresponds to our belief about how H_0 could be wrong.
- The two hypotheses must be disjoint: it should be the case that either H_0 is true or H_1 but never together simultaneously.

Heads or Tails?

- We believe that a coin is fair, unless we see enough evidence against it.
- In the coin tossing example, let p be the probability of "Heads". Our null hypothesis could be $H_0 : p = \frac{1}{2}$ and our alternative hypothesis could be $H_1 : p \neq \frac{1}{2}$.
- We may believe that the coin is fair, but suspect that it is biased towards heads if it is biased. In this case, we may pick $H_0 : p = \frac{1}{2}$ and $H_1 : p > \frac{1}{2}$.

Hypothesis testing as stochastic proof by contradiction

- Usually, the econometrician has to carry the "burden of proof" and the case that he is interested in is stated as H_1.
- The econometrician has to prove that his assertion (H_1) is true by showing that the data rejects H_0.
- As we perform proof by contradiction, we assume "something" we want to prove that it contradicts to itself.
- In the hypothesis testing, we assume the "null hypothesis" being true and we want to prove that our data contradicts to the hypothesis.

Hypothesis testing procedure

- Suppose we are concerned with a parameter $\theta \in \mathbb{R}$ (e.g. the population mean of some random variable for which we have repeated observations).
- We consider hypotheses like $H_0 : \theta \in I_0$ and $H_1 : \theta \in I_1$ with $I_0 \cap I_1 = \emptyset$. $I_0 \cup I_1$ is the maintained hypothesis.
- The econometrician has to choose between H_0 and H_1.
- We consider a test statistic T, which is a function of the data. T should be informative about the value of θ. T typically has a known distribution under H_0.
- Then we choose a significance level α, with $0 < \alpha < 1$. It is also called the size of the test. By convention, α is chosen to be a small number.
- α can be interpreted as the probability of a "very small-probability" event: How small a chance should be considered small?

Hypothesis testing procedure

- Then one rejects H_0 if the test statistic falls into a critical region Γ_α. A critical region is constructed by taking into account the probability of making a wrong decision, i.e. it typically depends on α. This is usually a set of possible values of the test statistic that contains the test statistic with probability α, under H_0, i.e. $\Pr[T \in \Gamma_\alpha] = \alpha$ under H_0. Notice that this requires knowledge about the distribution of T under H_0.
- If the test statistic assumes a value in the critical region, $T \in \Gamma_\alpha$, this is considered to be strong evidence against H_0. In this case, we reject H_0 in favor of H_1. Otherwise, we fail to reject H_0.
- If α is small and $T \in \Gamma_\alpha$, it means that the sample is unlikely if H_0 is true and so we have evidence against it, there are two possibilities: we observed a very small-probability event or our assumption (H_0) is wrong.

Hypothesis testing as stochastic proof by contradiction

- If we accept the principle that a very small-probability event cannot happen in the current state of the world, we accept the second reasoning: our assumption (H_0) is wrong.
- If $T \notin \Gamma_\alpha$, there is no "significant" contradiction between data and the null hypothesis. On the other hand, if $T \in \Gamma_\alpha$, there is a "significant" contradiction and we reject H_0 like what we conclude in proof by contradiction.
- Notice that "reject H_0" and "fail to reject H_0" are used for conclusions drawn from tests. Such terminology reflects the asymmetry of H_0 and H_1.
- We do not say "accept H_0" instead of "fail to reject H_0" since if we find $T \notin \Gamma_\alpha$, we did not "prove" anything and our test result is non-conclusive, since we did not find a contradiction.
- The decision depends on the significance level α: larger values of α correspond to bigger critical regions Γ_α . It is easier to reject the null for larger values of α.

A simple example

- Consider a normal population with mean μ and variance σ^2 and a random sample from the population. Suppose we know σ^2 for now.
- Consider a one-sided test: $H_0 : \mu = 0$ against $H_1 : \mu > 0$.
- Consider

$$T = \frac{\bar{X}}{\sigma / \sqrt{n}},$$

  which is a $\mathrm{N}(0, 1)$ random variable under H_0. H_1 is more likely to be true if T is large.
- Consider $\Gamma_\alpha = [z_{1-\alpha}, \infty)$, where $z_{1-\alpha}$ is the $(1 - \alpha)$-th quantile of $\mathrm{N}(0, 1)$.
- Under H_0, $\Pr[T \in \Gamma_\alpha] = \alpha$. And notice that Γ_α is a sub-interval of $\Gamma_{\alpha'}$ if $\alpha < \alpha'$.

Summary of the steps

- The following are the steps of the hypothesis testing:
  1. Specify H_0 and H_1.
  2. Choose the significance level α.
  3. Define a decision rule (a test statistic and a critical region).
  4. Perform the test using the data: given the data compute the test statistic and see if it falls into the critical region.

p-Value

- Given a realization t of the test statistic T, we can compute the lowest significance level consistent with rejecting H_0. This is called the p-value:

$$p\text{-value} \quad = \quad \min\left\{0 < \alpha < 1 : t \in \Gamma_{\alpha}\right\}.$$

- The p-value could be viewed as a measure of contradiction. The smaller the p-value is, the larger the contradiction is. In a probabilistic model, we use the way of reasoning in proof by contradiction with a "measure" of contradiction based on data to carry out the similar reasoning.
- Now if the p-value is smaller than our tolerance(significance level), then we reject null hypothesis like what we conclude in proof by contradiction.

The simple example

- In the simple example, the critical region is $\Gamma_\alpha = [z_{1-\alpha}, \infty)$. Suppose we have a realization t of T.
- Suppose we find that $z_{1-p} = t$. We would reject H_0 for all significance levels $\alpha \geq p$. p is the realized p-value.

Errors

- There are two types of errors that the econometrician can make:

		Truth	
		H_0	H_1
Decision	H_0	✓	Type II error
	H_1	Type I error	✓

- Type I error is the error of rejecting H_0 when H_0 is true.

$$\Pr(\text{Type I error}) = \Pr(\text{reject } H_0 \mid H_0 \text{ is true}) = \alpha.$$

- Type II error is the error of not rejecting H_0 when H_1 is true.

Errors

- The real line is split into two regions: acceptance region and rejection region (critical region).
  - When $T \notin \Gamma_\alpha$, we accept H_0 (and risk making a Type II error).
  - When $T \in \Gamma_\alpha$, we reject H_0 (and risk making a Type I error).
  - Unfortunately, the probabilities of Type I and II errors are inversely related.
  - By decreasing the probability of Type I error α, one makes the critical region smaller, which increases the probability of the Type II error. Thus it is impossible to make both errors arbitrary small.
- We require that the test we carry out has to be valid: the probability of type-I error ($\Pr[T \in \Gamma_\alpha]$ under H_0) must be α.
- We want the probability of a type-II error to be as small as possible for a given probability of a type-I error.

Power of tests

- Power of a test:

$$1 - \Pr(\text{Type II error}) = 1 - \Pr(\text{Accept } H_0 \mid H_0 \text{ is false}).$$

- We want the power of a test to be as large as possible, for a given significance level.
- We may not know the distribution of the test statistic under H_1. The distribution typically depends on θ.
- The power function $\pi(\theta)$ of the test is the probability that H_0 is rejected as a function of the true parameter value θ.

The simple example

- For any given value of μ, define

$$T_\mu = T - \frac{\mu}{\sigma/\sqrt{n}} = \frac{\bar{X} - \mu}{\sigma/\sqrt{n}},$$

  which is always $\mathrm{N}\left(0, 1\right)$ if the true mean is μ.

- We reject the test when $T \geq z_{1-\alpha}$, which holds if and only if

$$T_\mu = T - \frac{\mu}{\sigma/\sqrt{n}} \geq z_{1-\alpha} - \frac{\mu}{\sigma/\sqrt{n}}.$$

  It follows for this simple example, that the power function is

$$\pi\left(\mu\right) = 1 - \Phi\left(z_{1-\alpha} - \frac{\sqrt{n}\mu}{\sigma}\right).$$

- Notice that $\pi\left(0\right) = 1 - \Phi\left(z_{1-\alpha}\right) = \alpha$, which is simply the significance level.

The simple example

- We notice that $\pi(\mu)$ is smaller for all μ if the significance level α is smaller (and hence $z_{1-\alpha}$ is larger). This reflects the trade-off between type-I error and type-II error probabilities: we cannot reduce both simultaneously.
- $\pi(\mu)$ is increasing in μ. For μ's that are farther away from 0, the test can detect such deviation at a higher probability.
- As $\mu \to \infty$, the power converges to 1. The test is very likely to reject H_0 if the true mean is very large.
- $\pi(\mu)$ increases with the sample size n. The test can detect falseness of H_0 at a higher probability if our sample contains more information.

The simple example

- We consider an alternative critical region: $\tilde{\Gamma}_\alpha = (-\infty, -z_{1-\alpha}]$. Notice that under H_0, $\Pr\left[T \in \tilde{\Gamma}_\alpha\right] = \alpha$, so this is a valid test.
- We reject the test when $T \leq -z_{1-\alpha}$, which holds if and only if

$$T_\mu = T - \frac{\mu}{\sigma/\sqrt{n}} \leq -z_{1-\alpha} - \frac{\mu}{\sigma/\sqrt{n}}.$$

- So the power function is

$$\tilde{\pi}(\mu) = \Phi\left(-z_{1-\alpha} - \frac{\sqrt{n}\mu}{\sigma}\right).$$

- $\tilde{\pi}(\mu)$ is decreasing in both μ and n, which is undesirable.
- For all $\mu > 0$ (remember μ is the true mean), the power is smaller than α. At the same significance level, the power of the alternative critical region is much lower.

Linear model assumptions

- β_1 is unknown, and we have to rely on its OLS estimator $\hat{\beta}_1$.
- We need to know the distribution of $\hat{\beta}_1$ or of its certain functions.
- We will assume that the assumptions of the Normal Classical Linear Regression model are satisfied:
    1. $Y_i = \beta_0 + \beta_1 X_i + U_i, \qquad i = 1, \ldots, n.$
    2. $\mathrm{E}\left[U_i \mid X_1, \ldots, X_n\right] = 0$ for all i's.
    3. $\mathrm{E}\left[U_i^2 \mid X_1, \ldots, X_n\right] = \sigma^2$ for all i's.
    4. $\mathrm{E}\left[U_i U_j \mid X_1, \ldots, X_n\right] = 0$ for all $i \neq j$.
    5. U's are jointly normally distributed conditional on X's.
- Recall that, in this case, conditionally on X's:

$$\hat{\beta}_1 \sim \mathrm{N}\left(\beta_1, \mathrm{Var}\left[\hat{\beta}_1\right]\right), \text{ where } \mathrm{Var}\left[\hat{\beta}_1\right] = \frac{\sigma^2}{\sum_{i=1}^{n}\left(X_i - \bar{X}\right)^2}.$$

Two-sided tests

- For $Y_i = \beta_0 + \beta_1 X_i + U_i$, consider testing

  $$H_0 : \beta_1 = \beta_{1,0},$$

  against

  $$H_1 : \beta_1 \neq \beta_{1,0}.$$

- β_1 is the true unknown value of the slope parameter.
- $\beta_{1,0}$ is a known number specified by the econometrician. (For example $\beta_{1,0}$ is zero if you want to test $\beta_1 = 0$).
- Such a test is called two-sided because the alternative hypothesis H_1 does not specify in which direction β_1 can deviate from the asserted value $\beta_{1,0}$.

Two-sided tests when σ^2 is known (infeasible test)

- Suppose for a moment that σ^2 is known.
- Consider the following test statistic:

$$T = \frac{\hat{\beta}_1 - \beta_{1,0}}{\sqrt{\mathrm{Var}\left[\hat{\beta}_1\right]}}, \text{ where } \mathrm{Var}\left[\hat{\beta}_1\right] = \frac{\sigma^2}{\sum_{i=1}^{n}\left(X_i - \bar{X}\right)^2}.$$

- Consider the following decision rule (test):

$$\text{Reject } H_0 : \beta_1 = \beta_{1,0} \text{ when } |T| > z_{1-\alpha/2},$$

where $z_{1-\alpha/2}$ is the $(1 - \alpha/2)$ quantile of the standard normal distribution (critical value).

Test validity and power

- We need to establish that:
  1. The test is valid, where the validity of a test means that it has correct size or $\Pr(\text{Type I error}) = \alpha$: if $\beta_1 = \beta_{1,0}$

$$\Pr[|T| > z_{1-\alpha/2}] = \alpha.$$

  2. The test has power: when $\beta_1 \neq \beta_{1,0}$ (H_0 is false), the test rejects H_0 with probability that exceeds α: if $\beta_1 \neq \beta_{1,0}$

$$\Pr[|T| > z_{1-\alpha/2}] > \alpha.$$

- We want $\Pr[|T| > z_{1-\alpha/2}]$ when $\beta_1 \neq \beta_{1,0}$ to be as large as possible.
- Note that $\Pr[|T| > z_{1-\alpha/2}]$ when $\beta_1 \neq \beta_{1,0}$ depends on the true value β_1.

The distribution of T when σ^2 is known (infeasible test)

- Write

$$
\begin{aligned}
T &= \frac{\hat{\beta}_1 - \beta_{1,0}}{\sqrt{\mathrm{Var}\left[\hat{\beta}_1\right]}} = \frac{\hat{\beta}_1 - \beta_1 + \beta_1 - \beta_{1,0}}{\sqrt{\mathrm{Var}\left[\hat{\beta}_1\right]}} \\
&= \frac{\hat{\beta}_1 - \beta_1}{\sqrt{\mathrm{Var}\left[\hat{\beta}_1\right]}} + \frac{\beta_1 - \beta_{1,0}}{\sqrt{\mathrm{Var}\left[\hat{\beta}_1\right]}}.
\end{aligned}
$$

- Under our assumptions and conditionally on X's:

$$
\hat{\beta}_1 \sim \mathrm{N}\left(\beta_1, \mathrm{Var}\left[\hat{\beta}_1\right]\right), \text{ or } \frac{\hat{\beta}_1 - \beta_1}{\sqrt{\mathrm{Var}\left[\hat{\beta}_1\right]}} \sim \mathrm{N}\left(0, 1\right).
$$

- We have that conditionally on X's: $T \sim \mathrm{N}\left(\frac{\beta_1 - \beta_{1,0}}{\sqrt{\mathrm{Var}\left[\hat{\beta}_1\right]}}, 1\right)$.

Validity when σ^2 is known (infeasible test)

- We have that

$$T \sim \mathrm{N}\left(\frac{\beta_1 - \beta_{1,0}}{\sqrt{\mathrm{Var}\left[\hat{\beta}_1\right]}}, 1\right).$$

- When $H_0 : \beta_1 = \beta_{1,0}$ is true, $T \overset{H_0}{\sim} \mathrm{N}(0,1)$.
- We reject H_0 when

$$|T| > z_{1-\alpha/2} \Leftrightarrow T > z_{1-\alpha/2} \text{ or } T < -z_{1-\alpha/2}.$$

- Let $Z \sim \mathrm{N}(0,1)$.

$$\begin{aligned}&\Pr(\text{Reject } H_0 \mid H_0 \text{ is true})\\ &= \Pr\left[Z > z_{1-\alpha/2}\right] + \Pr\left[Z < -z_{1-\alpha/2}\right] = \alpha/2 + \alpha/2 = \alpha.\end{aligned}$$

The distribution of T when σ^2 is known (infeasible test)

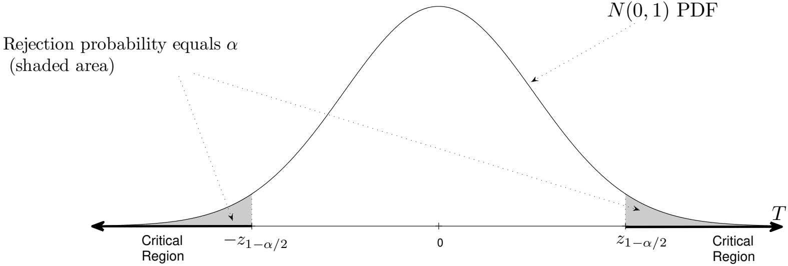

Power of the test when σ^2 is known (infeasible test)

- Under H_1, $\beta_1 - \beta_{1,0} \neq 0$ and the distribution of T is not centered zero: $T \sim \mathrm{N}\left(\frac{\beta_1 - \beta_{1,0}}{\sqrt{\mathrm{Var}\left[\hat{\beta}_1\right]}}, 1\right)$.
- When $\beta_1 - \beta_{1,0} > 0$:

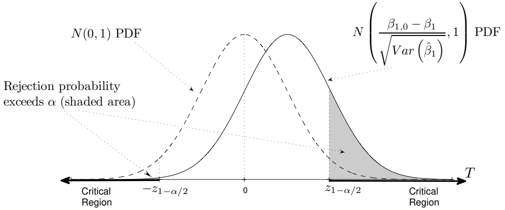

- Rejection probability exceeds α under H_1: power increases with the distance from H_0 ($|\beta_{1,0} - \beta_1|$) and decreases with $\mathrm{Var}\left[\hat{\beta}_1\right]$.

The two-sided t-test

- We are testing $H_0 : \beta_1 = \beta_{1,0}$ against $H_1 : \beta_1 \neq \beta_{1,0}$.
- When σ^2 is unknown, we replace it with $s^2 = \frac{1}{n-2} \sum_{i=1}^{n} \hat{U}_i^2$.
- The t-statistic:

$$T = \frac{\hat{\beta}_1 - \beta_{1,0}}{\sqrt{\widehat{\mathrm{Var}}\left[\hat{\beta}_1\right]}} = \frac{\hat{\beta}_1 - \beta_{1,0}}{\sqrt{\frac{s^2}{\sum_{i=1}^{n}(X_i - \bar{X})^2}}}.$$

- We also replace the standard normal critical values $z_{1-\alpha/2}$ with the t_{n-2} critical values $t_{n-2,1-\alpha/2}$.
  However, for large n, $t_{n-2,1-\alpha/2} \approx z_{1-\alpha/2}$.
- The two-sided$t$-test:

$$\text{Reject } H_0 \text{ when } |T| > t_{n-2,1-\alpha/2}.$$

The two-sided p-value

- The decision to accept or reject H_0 depends on the critical value $t_{n-2,1-\alpha/2}$.
- If $\alpha_1 > \alpha_2$ then $t_{1-\alpha_1/2} < t_{1-\alpha_2/2}$.
- Thus, it is easier to reject H_0 with the significance level α_1 since it corresponds to a smaller acceptance region.
- p-value is the smallest significance level α for which we can reject H_0.

The two-sided p-value

- In order to find p-value:
  1. Compute T.
  2. Find τ such that $|T| = t_{n-2,1-\tau}$.
  3. The p-value$=\tau \times 2$.
- Note that for all $\alpha > p$-value,

$$|T| = t_{n-2,1-(p\text{-value})/2} > t_{n-2,1-\alpha/2}$$

  and we will reject H_0.
- For all $\alpha \leq p$-value,

$$|T| = t_{n-2,1-(p\text{-value})/2} \leq t_{n-2,1-\alpha/2}$$

  and we will accept H_0.

Example of p-value calculation

Suppose a regression with 19 observations produced the following output:

y	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
x	-.6725304	.5804943	-1.16	0.263	-1.897266	.5522055
_cons	10.18197	.2509365	40.58	0.000	9.652542	10.7114

- Here, $\hat{\beta}_1 = -0.6725, \beta_{1,0} = 0$, and in the 4th column $t = -0.6725/0.5804 = -1.16$.
  - Thus, $|T| = 1.16$ and df=17.
  - From the t-table, the closest critical value is $t_{17,1-0.15} = 1.069$.
    (The probability that a random variable with t_{17}-distribution lies on the right of 1.16 is ≈ 0.15.)
  - The p-value is then $\approx 0.15 \times 2 = 0.300$.

Stata

- We can compute critical values and p-values using Stata instead of using the tables.
- To compute standard normal critical values use:

$$\text{display invnormal}(\tau),$$

  were τ is a number between 0 and 1.
- For example: display invnormal(1-0.05/2) produces 1.959964.
- For t critical values use

$$\text{display invttail}(df,\tau),$$

  where df is the number of degrees of freedom and τ is a number between 0 and 1.
  Note that here τ is the right-tail probability!
- For example, display invttail(62,0.05/2) produces 1.9989715.

Stata

- To compute two-sided normal p-values use:

$$\text{display } 2 * (\text{1-normal}\,(T))\,.$$

- For example, display 2*(1-normal(1.96)) produces 0.04999579.
- To compute two-sided t-distribution p-values, use

$$\text{display } 2\text{*}\,(\text{ttail}(df, T))\,,$$

  Note that ttail gives the right tail probabilities!
- For example, display 2*(ttail(62, 1.96)) produces 0.05449415.

Example

rent	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
avginc	.01158	.0013084	8.85	0.000	.0089646	.0141954
_cons	148.7764	32.09787	4.64	0.000	84.6137	212.9392

- Stata report the t-statistics and the p-value for $H_0 : \beta = 0$.
- To test H_0 whether the coefficient of AvgInc is zero : $T = 0.01158/0.0013084 = 8.85$.
- The p-value is extremely close to zero, (display 2*(ttail(62, 8.85)) gives 1.345×10^{-12}), so for all reasonable significance levels α, we reject H_0 that the coefficient of AvgInc is zero.
- AvgInc is a statistically significant regressor.

Example

rent	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
avginc	.01158	.0013084	8.85	0.000	.0089646	.0141954
_cons	148.7764	32.09787	4.64	0.000	84.6137	212.9392

- Consider now testing H_0 that the coefficient of AvgInc is 0.009 against the alternative that it is different from 0.009.
- $T = (0.01158 - 0.009) / 0.0013084 \approx 1.97$.
- At 5% significance level, $t_{62,0.975} \approx 1.999 > T$ and we accept H_0.
- At 10% significance level, $t_{62,0.95} \approx 1.67 < T$ and we reject H_0.
- The two sided p-value is $2^*(\text{ttail}(62, 1.97)) \implies \approx 0.053$.
- For $\alpha \leq 0.053$ we will accept H_0 and for $\alpha > 0.053$ we will reject H_0.

Confidence intervals and hypothesis testing

- There is one-to-one correspondence between confidence intervals and hypothesis testing.
- We cannot reject $H_0 : \beta_1 = \beta_{1,0}$ against a two-sided alternative if $|T| \leq t_{n-2,1-\alpha/2}$ or if and only if:

$$-t_{n-2,1-\alpha/2} \leq \frac{\hat{\beta}_1 - \beta_{1,0}}{\sqrt{\widehat{\mathrm{Var}}\left[\hat{\beta}_1\right]}} \leq t_{n-2,1-\alpha/2}$$

$$\Longleftrightarrow$$

$$\hat{\beta}_1 - t_{n-2,1-\alpha/2}\sqrt{\widehat{\mathrm{Var}}\left[\hat{\beta}_1\right]} \leq \beta_{1,0} \leq \hat{\beta}_1 + t_{n-2,1-\alpha/2}\sqrt{\widehat{\mathrm{Var}}\left[\hat{\beta}_1\right]}$$

$$\Longleftrightarrow$$

$$\beta_{1,0} \in CI_{1-\alpha}.$$

- Thus, for any $\beta_{1,0} \in CI_{1-\alpha}$, we cannot reject $H_0 : \beta_1 = \beta_{1,0}$ against $H_1 : \beta_1 \neq \beta_{1,0}$ at significance level α.

Example

```
. regress rent avginc

      Source |       SS       df       MS              Number of obs =      64
-------------+------------------------------           F(  1,    62) =   78.34
       Model |  347069.249     1  347069.249           Prob > F      =  0.0000
    Residual |  274693.188    62  4430.53529           R-squared     =  0.5582
-------------+------------------------------           Adj R-squared =  0.5511
       Total |  621762.438    63  9869.24504           Root MSE      =  66.562

------------------------------------------------------------------------------
        rent |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      avginc |     .01158   .0013084     8.85   0.000     .0089646    .0141954
       _cons |   148.7764   32.09787     4.64   0.000      84.6137    212.9392
------------------------------------------------------------------------------
```

- The 95% confidence interval for the coefficient of AvgInc is [0.0089646,0.0141954].
- A significance level 5% test of $H_0 : \beta_1 = \beta_{1,0}$ against $H_1 : \beta_1 \neq \beta_{1,0}$ will not reject H_0 if $\beta_{1,0} \in [0.0089646, 0.0141954]$.

One-sided tests

- Consider testing $H_0 : \beta_1 \leq \beta_{1,0}$ against $H_1 : \beta_1 > \beta_{1,0}$.
- It is reasonable to reject H_0 when $\hat{\beta}_1 - \beta_{1,0}$ is large and positive or when

$$T = \frac{\hat{\beta}_1 - \beta_{1,0}}{\sqrt{\widehat{\mathrm{Var}}\left[\hat{\beta}_1\right]}} > c_{1-\alpha}$$

  where $c_{1-\alpha}$ is a positive constant.
- The null hypothesis H_0 is composite. The probability of rejection under H_0 depends on β_1.
- We pick the critical value $c_{1-\alpha}$ so that

$$\Pr\left[\frac{\hat{\beta}_1 - \beta_{1,0}}{\sqrt{\widehat{\mathrm{Var}}\left[\hat{\beta}_1\right]}} > c_{1-\alpha} \mid \beta_1 \leq \beta_{1,0}\right] \leq \alpha$$

  for all $\beta_1 \leq \beta_{1,0}$.

One-sided tests

- For all $\beta_1 \leq \beta_{1,0}$,
$$\frac{\beta_1 - \beta_{1,0}}{\sqrt{\widehat{\text{Var}}\left[\hat{\beta}_1\right]}} \leq 0,$$
and
$$\begin{aligned} & \Pr\left[\frac{\hat{\beta}_1 - \beta_{1,0}}{\sqrt{\widehat{\text{Var}}\left[\hat{\beta}_1\right]}} > c_{1-\alpha} \mid \beta_1 \leq \beta_{1,0}\right] \\ = \ & \Pr\left[\frac{\hat{\beta}_1 - \beta_1}{\sqrt{\widehat{\text{Var}}\left[\hat{\beta}_1\right]}} + \frac{\beta_1 - \beta_{1,0}}{\sqrt{\widehat{\text{Var}}\left[\hat{\beta}_1\right]}} > c_{1-\alpha} \mid \beta_1 \leq \beta_{1,0}\right] \\ \leq \ & \Pr\left[\frac{\hat{\beta}_1 - \beta_1}{\sqrt{\widehat{\text{Var}}\left[\hat{\beta}_1\right]}} > c_{1-\alpha} \mid \beta_1 \leq \beta_{1,0}\right] \\ = \ & \alpha \text{ if } c_{1-\alpha} = t_{n-2,1-\alpha}. \end{aligned}$$

One-sided tests

- For size α test, we reject $H_0 : \beta_1 \leq \beta_{1,0}$ against $H_1 : \beta_1 > \beta_{1,0}$ when

$$T = \frac{\hat{\beta}_1 - \beta_{1,0}}{\sqrt{\widehat{\text{Var}}\left[\hat{\beta}_1\right]}} > t_{n-2,1-\alpha}.$$

  where $t_{n-2,1-\alpha}$ is the critical value corresponding to t-distribution with $n-2$ degrees of freedom.
  - Note that we use $1-\alpha$ and not $1-\alpha/2$ for choosing critical values in the case of one-sided testing.
  - For size α test, we reject $H_0 : \beta_1 \geq \beta_{1,0}$ against $H_1 : \beta_1 < \beta_{1,0}$ when

$$T = \frac{\hat{\beta}_1 - \beta_{1,0}}{\sqrt{\widehat{\text{Var}}\left[\hat{\beta}_1\right]}} < -t_{n-2,1-\alpha}.$$

One-sided tests

- One-sided p-values for $H_0 : \beta_1 \leq \beta_{1,0}$ against $H_1 : \beta_1 > \beta_{1,0}$:
  1. Compute T.
  2. Find τ such that $T = t_{n-2,1-\tau}$.
  3. The p-value$=\tau$.