Module 11: Statistical inference for simple linear regression

(The attached PDF file has better formatting.)

Intuition: *P-VALUES VS CRITICAL VALUES*

Jacob: We often speak of rejecting a null hypothesis at a 95% or a 90% confidence interval. We can also phrase hypothesis testing with *p*-values. Which is better?

Rachel: Statisticians prefer *p*-values. If we reject a null hypothesis at a 5% significance level, we don't know if the *p*-value is 5.1%, and the null hypothesis probably ought to be rejected (or viewed with suspicion) or the *p*-value is 50%, and the null hypothesis should not be rejected. A 5% significance level is an arbitrary choice; it has no greater justification than a 6% level or a 4% level or any other level. Yet social scientists sometimes speak of regression results as absolutes; they say a certain result is significant or is not significant. This is misleading.

Jacob: If the *p*-value is better, why do we use arbitrary confidence intervals?

Rachel: A lay person may have trouble interpreting a *p*-value. Suppose we want to know if women or more likely than men to vote for one of two candidates in an election. If we say "the *p*-value is 8.2%," the listener says: "What does that mean?" Explaining the statistical meaning to a lay person may not help. So we choose a significance level and say: "yes" or "no." The listener may not realize that we could change "yes" to "no" by changing the significance level.

Jacob: For actuaries, is the *p*-value a good measure?

Rachel: It is a better measure than a significant test, but it suffers from the same problems. Listeners thinks we are testing the observed relation between the X and Y variables, but we are only testing the null hypothesis. In many regression analyses, we are confident that β is not zero, but we don't know its true value.

Illustration: Suppose we are determining the inflation rate, the interest rate, or a loss cost trend. We know that the trend is not 0%, but we don't know its true value, such as 8%, 9%, or 10%. A *p*-value is no help. If the observed trend is 8.7%, the *p*-value may be 0.01%. This doesn't tell us that the trend is 8.7%; it says that the trend is not 0%, which we know.

Jacob: Is a confidence interval better?

Rachel: It is better to say that we are P% confident that the true trend is between 8.7% – z and 8.7% + z.

Jacob: This seems like a good statement; it answers our concerns about the true trend.

Rachel: Not necessarily. We want to know the current trend. But the statistical statement says the following: "If the trend has been stable over the experience period, and any observed differences over the years stem solely from sampling error, then the true trend is between 8.7% – z and 8.7% + z." Our listeners respond: "We do not assume the trend is the same every year. It may change from year to year. We want to know the best estimate of the current trend."

Jacob: Isn't the ordinary least squares estimator the best estimate of the current trend?

Rachel: Suppose we have 11 years with trends of 8.0%, 8.2%, 8.4%, …, 9.8%, and 10.0%. The standard trend analysis gives an ordinary least squares estimator of 9.0%. Our listeners are likely to reject this in favor of a 10.0% current trend.

Jacob: For trend analyses, should we should use the most recent value?

Rachel: Suppose we examine 11 years, and we find trends of

9.0%, 8.2%, 8.8%, 9.8%, 8.4%, 9.6%, 8.6%, 9.4%, 8.0%, and 10.0%.

We ascribe the differences to sampling error, and we choose a trend of 9%, not 10%.

Jacob: How do we choose between these two scenarios?

Rachel: The time series course deals with this choice. The first scenario is a random walk, and the second scenario is white noise. The time series question is "How much of the observed annual differences is the drift of a random walk and how much is sampling error of white noise?"