may be tempted to pick the one that gives them the results they want. This is not overtly dishonest, but it does lead to a bias towards positive results.

Practical and Statistical Significance

Statistical significance is not equivalent to practical significance. The larger the sample, the smaller your p-values will be, so do not confuse p-values with an important predictor effect. With large datasets it will be very easy to get statistically significant results, but the actual effects may be unimportant. Would we really care that test scores were 0.1% higher in one state than another or that some medication reduced pain by 2%? CIs on the parameter estimates are a better way of assessing the size of an effect. They are useful even when the null hypothesis is not rejected, because they tell us how confident we are that the true effect or value is close to the null.

It is also important to remember that a model is usually only an approximation of underlying reality which makes the exact meaning of the parameters debatable at the very least. The precision of the statement that $\beta_1 = 0$ exactly is at odds with the acknowledged approximate nature of the model. Furthermore, it is highly unlikely that a predictor that one has taken the trouble to measure and analyze has exactly zero effect on the response. It may be small but it will not be zero.

This means that in many cases, we know the point null hypothesis is false without even looking at the data. Furthermore, we know that the more data we have, the greater the power of our tests. Even small differences from zero will be detected with a large sample. Now if we fail to reject the null hypothesis, we might simply conclude that we did not have enough data to get a significant result. According to this view, the hypothesis test just becomes a test of sample size. For this reason, we prefer CIs.

Exercises

1. For the `prostate` data, fit a model with `lpsa` as the response and the other variables as predictors.

   (a) Compute 90 and 95% CIs for the parameter associated with `age`. Using just these intervals, what could we have deduced about the p-value for `age` in the regression summary?

   (b) Compute and display a 95% joint confidence region for the parameters associated with `age` and `lbph`. Plot the origin on this display. The location of the origin on the display tells us the outcome of a certain hypothesis test. State that test and its outcome.

   (c) Suppose a new patient with the following values arrives:

   ```
    lcavol  lweight       age      lbph       svi       lcp
   1.44692  3.62301  65.00000  0.30010  0.00000  -0.79851
    gleason    pgg45
    7.00000 15.00000
   ```

   Predict the `lpsa` for this patient along with an appropriate 95% CI.

(d) Repeat the last question for a patient with the same values except that he or she is age 20. Explain why the CI is wider.

(e) In the text, we made a permutation test corresponding to the F-test for the significance of all the predictors. Execute the permutation test corresponding to the t-test for age in this model. (Hint: `{summary(g)$coef[4,3]` gets you the t-statistic you need if the model is called g.)

2. For the model of the previous question, remove all the predictors that are not significant at the 5% level.

   (a) Recompute the predictions of the previous question. Are the CIs wider or narrower? Which predictions would you prefer? Explain.

   (b) Test this model against that of the previous question. Which model is preferred?

3. Using the `teengamb` data, fit a model with `gamble` as the response and the other variables as predictors.

   (a) Which variables are statistically significant?

   (b) What interpretation should be given to the coefficient for `sex`?

   (c) Predict the amount that a male with average (given these data) status, income and verbal score would gamble along with an appropriate 95% CI. Repeat the prediction for a male with maximal values (for this data) of status, income and verbal score. Which CI is wider and why is this result expected?

   (d) Fit a model with just `income` as a predictor and use an F-test to compare it to the full model.

4. Using the `sat` data:

   (a) Fit a model with `total` sat score as the response and `expend`, `ratio` and `salary` as predictors. Test the hypothesis that $\beta_{salary} = 0$. Test the hypothesis that $\beta_{salary} = \beta_{ratio} = \beta_{expend} = 0$. Do any of these predictors have an effect on the response?

   (b) Now add `takers` to the model. Test the hypothesis that $\beta_{salary} = 0$. Compare this model to the previous one using an F-test. Demonstrate that the F-test and t-test here are equivalent.

5. Find a formula relating R^2 and the F-test for the regression.