1-sided vs. 2-sided tests

In hypothesis tests we want to determine whether our statistic falls in the tail of the null distribution. The null distribution is generated from the *null* hypothesis, but which tail we are interested in depends upon the *alterantive* hypothesis. If we have some reason to think that the actual parameter is *less than* the parameter used for the null distribution, then we are interested in the left-hand tail. If we think that the actual parameter is *greater than* the null parameter, then we are interested in the right hand tail. In both of these cases we are interested only in one tail. So we do a 1-sided test. However, if we don't really have a reason *a priori* to think that the actual parameter is greater than or less than the null parameter, then we need to take both tails into account. So we do a 2-sided test.

Here are all the possibilities based on the predicted parameter from the alternative hypothesis and the actual statistic value from the data. In these examples, I have plotted based on a particular null and statistic, but don't pay attention to the x axis. Instead just look at which area is shaded red compared to the entire distribution.

1) Alternative π > null hypothesis
Statistic > null hypothesis

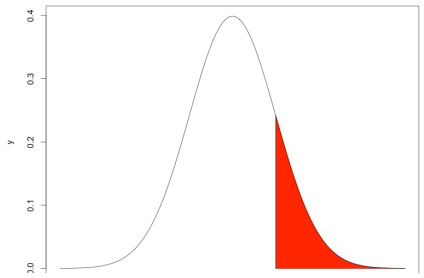

2) Alternative π < null hypothesis
Statistic < null hypothesis

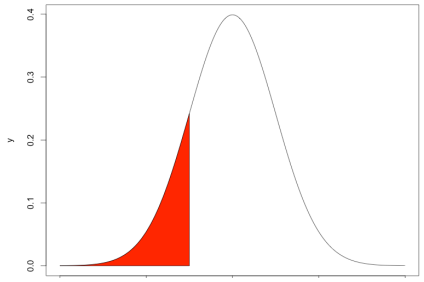

3) Alternative π < null hypothesis
Statistic > null hypothesis

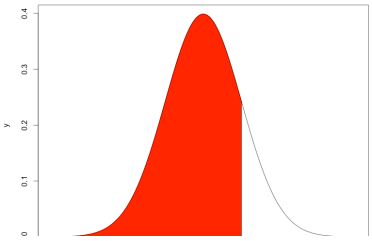

4) Alternative π > null hypothesis
Statistic < null hypothesis

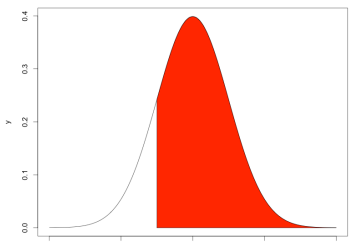

5) Alternative $\pi \neq$ null hypothesis
Statistic < null hypothesis
or Statistic > null hypothesis. Whichever it actually is, you just take the same amount in the tail on the other side