Randomized Trials: computing p-values

An alternative to computing Z-scores using the CLT is to explicitly simulate the distribution under the null hypothesis. Going back to coin flip examples, say we flip a coin 30 times and end up with 13 Heads. This is our *observed data*. We would like to ask the question: is this a fair coin?

1. To create the null distribution, say we perform many trials and during each trial we flip a *fair* coin 30 times. The results of these trials are shown below – in each one I record the number of Heads. What is T, the number of trials?

$$[19, 14, 12, 16, 12, 16, 13, 18, 18, 12, 17, 16, 11, 12, 12, 15, 15, 16, 15, 13]$$

2. What is N_e, the number of trials where we obtained a value *as or more* extreme than our observed (true) data? Since we are trying to determine if the coin is fair or not (as opposed to weighted towards Tails specifically), perform a two-sided computation.

3. Putting this all together, what is your estimate of the p-value? Based on your p-value, what do you conclude about the coin? (Note: this is a very small number of trials – just for illustration purposes!)

4. Roughly sketch the null distribution and shade the area(s) representing the p-value.

Acknowledgements: example from MIT OpenCourseWare 18.650