17.1: Simple Example- Coin-flipping (Section 16.3.5.1)

Let's say that we flipped 100 coins and observed 70 heads. We would like to use these data to test the hypothesis that the true probability is 0.5.

First let's generate our data, simulating 100,000 sets of 100 flips. We use such a large number because it turns out that it's very rare to get 70 heads, so we need many attempts in order to get a reliable estimate of these probabilties. This will take a couple of minutes to complete.

```
# simulate tossing of 100,000 flips of 100 coins to identify
# empirical probability of 70 or more heads out of 100 flips

nRuns <- 100000

# create function to toss coins
tossCoins <- function() {
0.5  flips <- runif(100) >
  return(tibble(nHeads=sum(flips)))
}

# create an input data frame for do()
input_df <- tibble(id=seq(nRuns)) %>%
  group_by(id)

# use do() to perform the coin flips
flip_results <- input_df %>%
  do(tossCoins()) %>%
  ungroup()

p_ge_70_sim <-
  flip_results %>%
  summarise(p_gt_70 = mean(nHeads >= 70)) %>%
  pull()

p_ge_70_sim
```

```
## [1] 3e-05
```

For comparison, we can also compute the p-value for 70 or more heads based on a null hypothesis of $P_{heads} = 0.5$, using the binomial distribution.

```
# compute the probability of 69 or fewer heads,
# when P(heads)=0.5
0.5p_lt_70 <- pbinom(69, 100, )

# the probability of 70 or more heads is simply
# the complement of p_lt_70
p_ge_70 <- 1 - p_lt_70

p_ge_70
```

```
## [1] 3.9e-05
```