

Introduction to Probability
Stat 341 - Fall 2008

What is the probability of obtaining heads when flipping a coin? Almost everyone would answer 0.5. Why? What does the value of 0.5 really mean? When you flip a coin once, you get either heads or tails. So, when you flip a coin once, either the event occurs (heads) or doesn't occur (tails). Then, why do we assign a probability of 0.5 to the event of obtaining heads when flipping a coin.

The concept of probability is based on what occurs in the long run, in a large number of trials. So, if we flip a coin 100 times, we expect to obtain around 50 heads ($100 * 0.5$). The probability of an event is always defined in terms of this long run relative frequency. So, if we keep flipping the coin forever (an infinite number of trials), heads will occur in exactly one-half of the flips. This is the theoretical probability of the event, 0.5.

Clearly, we cannot flip a coin (or perform any other event) an infinite number of times. So, if we were repeating an event, or simulating an event on the computer, how many simulations would we need until the observed relative frequency would be close to this theoretical probability? When you flip a coin 100 times, you expect the relative frequency of heads to be 0.5 (50 flips out of 100). However, for every flip you are away from 50, the relative frequency of heads changes by 0.01. Get 45 flips or 55 flips out of 100 with heads and your relative frequency is a full 0.05 away from the theoretical probability of 0.5. Contrast this when you flip a coin 10,000 times. In this situation, you still expect the relative frequency of heads to be 0.5 (5,000 flips out of 10,000). This time however, for every flip you are away from 5,000, the relative frequency of heads only changes by 0.0001. Get 4995 flips or 5005 flips out of 10,000 with heads and your relative frequency is only 0.0005 away from the theoretical probability of 0.5.

To see how this works, we can simulate flipping a coin in R and look at the relative frequency of obtaining heads. To set up the coin, we will define a new variable **coin** to have the values either 0 = Tails or 1 = Heads. Here is the command in R.

```
coin<- c(0,1)
```

Flipping a coin is like sampling from the **coin** variable with replacement. To flip the coin 100 times, we will use the **sample** command in R and save the result to the variable **flips100**.

```
flips100<- sample(coin, 100, replace = T)
```

In my 100 flips, I obtained 51 out of 100 heads, for an empirical probability of 0.51. You can get the number of heads out of your 100 flips and the proportion of heads out of your 100 flips using the R commands

```
sum(flips100)  
sum(flips100)/100
```

Flipping the coin in R 100 more times will give you a different set of outcomes and most likely a different number and proportion of heads out of your 100 flips. If we repeat this process many times, we can see what would commonly happen for the number of heads and the proportion of heads out of 100 flips of the coin. Here is the code in R that will conduct 10,000 trials with each trial consisting of flipping the coin 100 times. From this simulation, we will save the number and proportion of heads for each trial.

```

numheads100<- rep(0, 10000)
propheads100<- rep(0, 10000)

for (i in 1:10000){
flips100<- sample(coin, 100, replace = T)
numheads100[i]<- sum(flips100)
propheads100[i]<- sum(flips100)/100
}

```

The histograms of both the number and proportion of heads occurring in 100 flips of the coin for these 10,000 trials can be found at the end of the document. For the number of heads occurring in 100 flips, the histogram shows that most trials contained between 40 and 60 flips, for a difference of ± 10 . However, in a few of the trials, there were as low as 30 or as high as 70 heads out of the 100 flips, for a difference of ± 20 . For the proportion of heads occurring in 100 flips, the histogram shows that most trials contained a proportion of heads between 0.4 and 0.6, for a difference of ± 0.1 . However, in a few of the trials, the proportion was as low as 0.3 or as high as 0.7, for a difference of ± 0.2 .

From these results, we can see that 100 flips would not be enough to get a relative frequency of the event close to the theoretical probability. We need more flips. Let's try 10,000 flips. The code to produce the number of heads and the proportion of heads out of 10,000 flips of a coin is given below.

```

flips10000<- sample(coin, 10000, replace = T)

```

In my 10,000 flips, I obtained 4,993 heads for an empirical probability of 0.4993. Just like before, flipping the coin another 10,000 times will give you a different set of outcomes and most likely a different number and proportion of heads out of your 10,000 flips. And just like before, if we repeat this process many times, we can see what would commonly happen for the number of heads and the proportion of heads out of 10,000 flips of the coin. Here is the code in R that will conduct 10,000 trials with each trial consisting of flipping the coin 10,000 times. From this simulation, we will save the number and proportion of heads for each trial.

```

numheads10000<- rep(0, 10000)
propheads10000<- rep(0, 10000)

for (i in 1:10000){
flips10000<- sample(coin, 10000, replace = T)
numheads10000[i]<- sum(flips10000)
propheads10000[i]<- sum(flips10000)/10000
}

```

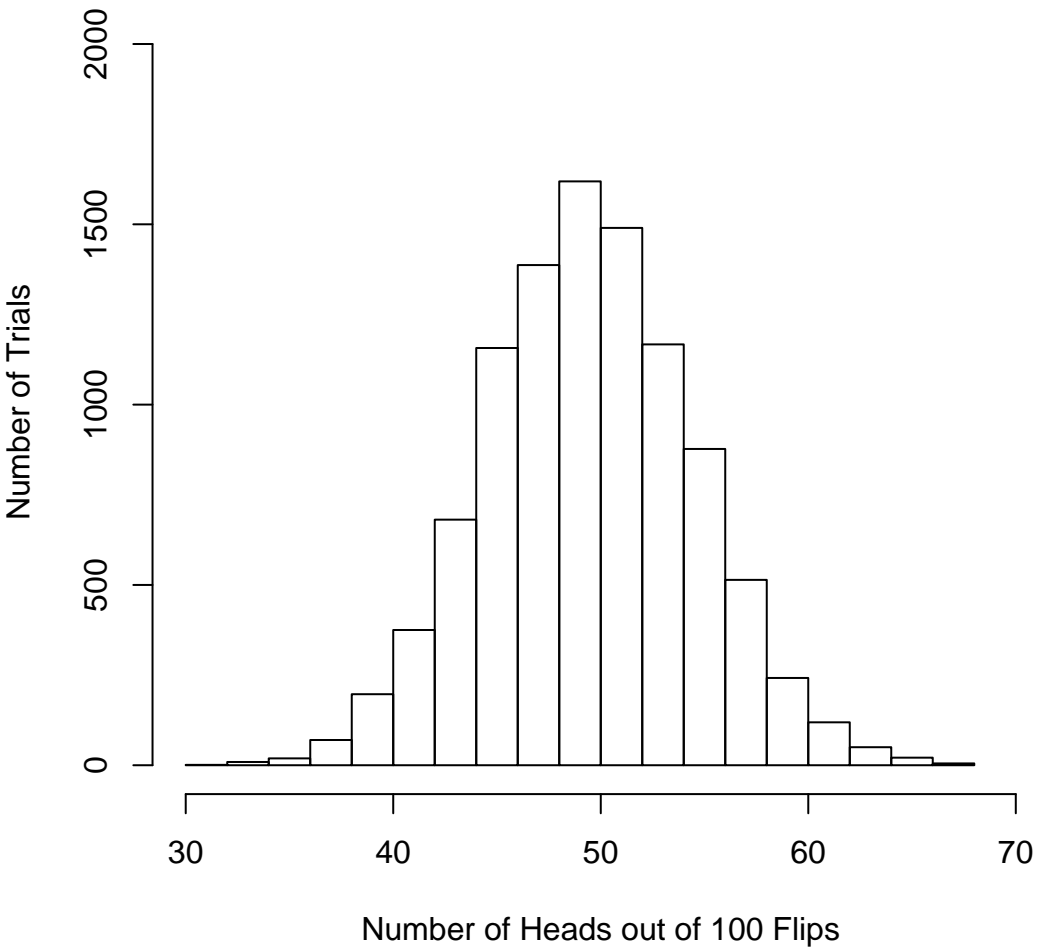
The histograms of both the number and proportion of heads occurring in 10,000 flips of the coin for these 10,000 trials can be found at the end of the document. For the number of heads occurring in 10,000 flips, the histogram shows that most trials contained between 4,900 and 5,100 flips, for a difference of ± 100 . However, in a few of the trials, there were as low as 4,800 or as high as 5,200 heads out of the 10,000 flips, for a difference of ± 200 . For the proportion of heads occurring in 10,000 flips, the histogram shows that most trials contained a proportion of heads between 0.49 and 0.51, for a difference of ± 0.01 . However, in a few of the trials, the proportion was as low as

0.48 or as high as 0.52, for a difference of ± 0.02 .

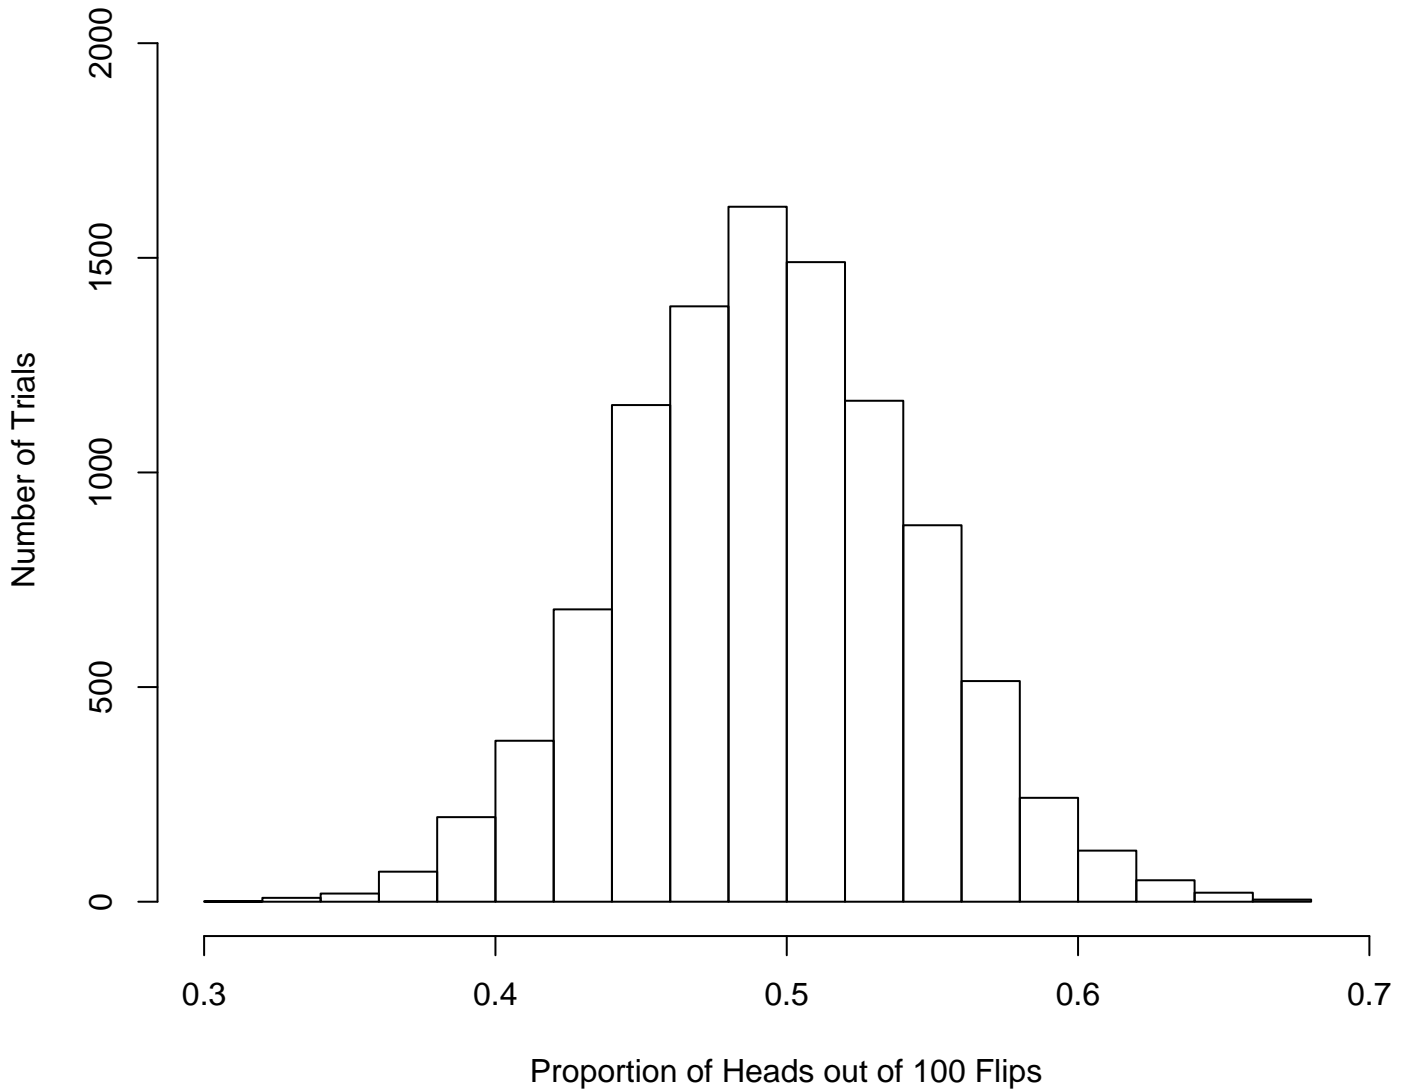
For 100 flips, the difference in the expected number of heads (50) and the observed number of heads in the 10,000 trials was fairly small, no more than ± 20 . However, since there were relatively few flips to begin with, the relative frequency of heads varied greatly (from 0.3 to 0.7) from the expected value of 0.5. For 10,000 flips, the difference in the expected number of heads (5,000) and the observed number of heads in the 10,000 trials was fairly large, with the largest being ± 200 . However, since there were many flips, each trial produced an empirical probability of heads that was within 0.02 of the theoretical probability of 0.5.

The more flips of the coin, the closer we expect the relative frequency or the empirical probability to be to the theoretical probability. For 100 flips, the empirical probabilities are not always very close to the theoretical probability of 0.5. However, with 10,000 flips, the empirical probabilities are typically within 1-2 percentage points of the theoretical probability.

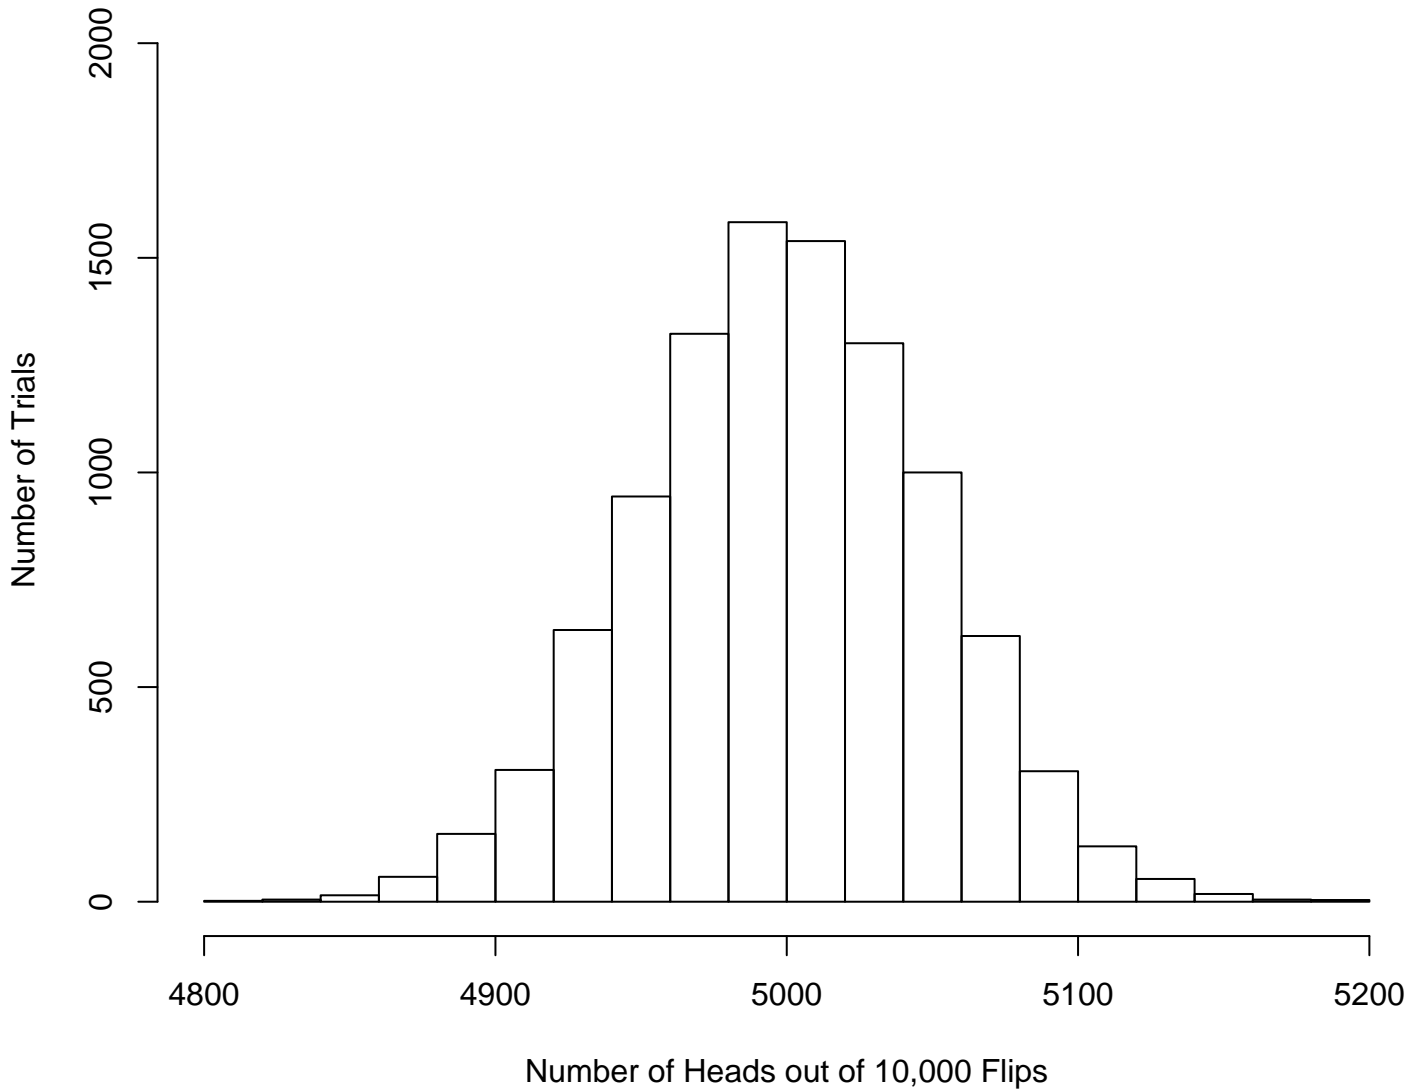
Results of 10,000 Trials: Number of Heads out of 100 Flips



Results of 10,000 Trials: Proportion of Heads out of 100 Flips



Results of 10,000 Trials: Number of Heads out of 10,000 Flips



Results of 10,000 Trials: Proportion of Heads out of 10,000 Flips

