Programming, Probability, and the Modern Mathematics Classroom Exercises — Part 2

Manan Shah
Mathematician-At-Large

June 3, 2013

Please make sure to have read the blog post with the same topic name on the Math Misery website, otherwise this will be out of context.

Some reminders

- Have students to comment the code, line by line.
- Let students tweak, modify, and otherwise alter the code.
- Don't worry about programming formalities.
- Discuss probability theory to put the code in context.
- Have fun, relax, and learn.

Understanding Independent Events

One of the gambler's fallacies is to assert that a sequence of independent events still carries with it some level of memory. For example, if a fair coin were flipped four times and four heads were observed, the probability that the fifth flip would be heads would still be 0.5. The outcome of the last N coin flips are irrelevant when determining the probability that heads will be observed on the $(N+1)$st coin flip.

The notion of conditional probability can be confusing and counter-intuitive, especially when combined with independent trials. If "H" represents heads and "T" represents tails, then the probability that HHHH when a fair coin is flipped four times, is $\frac{1}{16}$.

The following sets of exercises will focus on demonstrating, empirically, some properties of independent trials as well as highlighting what "conditions" are in conditional probability. These exercises will use `for` and `if-else` statements.

We will delay discussing "expected value" and "confidence intervals" to avoid complicating matters. We will focus on a frequentist approach to empirical probability.

Flipping a coin

What is the probability that in a group of N coin flips, the kth coin flip is heads? $k \in \{1, 2, \ldots, N\}$**.**

- How does the code below work towards this end?

- What does `raise` do? (Try it! Give a "bad" value for a probability!)
- What does `[0 for i in range(numflips)]` produce? (Try it in the command line to see what the output is!)
- What do the nested `for` loops do?
- Why is the statement `if flip < headsprobability:` a way to determine if the `flip` was heads? Why can't `if flip > headsprobability:` be correct? Would there be a difference if `if flip < headsprobability:` were `if flip <= headsprobability:`?
- What can be said about the quality of the results as `numtrials` is increased?

```
>>> import random
In 10000 trials of 5 per trial, I observed 4963 heads in position 1
In 10000 trials of 5 per trial, I observed 5049 heads in position 2
In 10000 trials of 5 per trial, I observed 5070 heads in position 3
In 10000 trials of 5 per trial, I observed 4984 heads in position 4
In 10000 trials of 5 per trial, I observed 4981 heads in position 5
>>> def coinflip(headsprobability, numflips, numtrials):
        if headsprobability > 1 or headsprobability < 0:
                raise ValueError("That's not a valid value for a probability!")
        headscounter = [0 for i in range(numflips)]
        for trialindex in range(numtrials):
                for flipindex in range(numflips):
                        flip = random.random()
                        if flip < headsprobability:
                                headscounter[flipindex] += 1
        for index, count in enumerate(headscounter):
                print("In " + str(numtrials) + " trials of ", end = "")
                print(str(numflips) + " flips per trial, I observed ", end = "")
                print(str(count) + " (" + str(count/numtrials*100) + "%)", end = "")
                print(" heads in position " + str(index + 1))


>>> coinflip(0.5,5,10000)
In 10000 trials of 5 flips per trial, I observed 4983 (49.83%) heads in position 1
In 10000 trials of 5 flips per trial, I observed 4977 (49.77%) heads in position 2
In 10000 trials of 5 flips per trial, I observed 5062 (50.62%) heads in position 3
In 10000 trials of 5 flips per trial, I observed 4982 (49.82%) heads in position 4
In 10000 trials of 5 flips per trial, I observed 4957 (49.57%) heads in position 5
>>> coinflip(0.25,5,10000)
In 10000 trials of 5 flips per trial, I observed 2560 (25.6%) heads in position 1
In 10000 trials of 5 flips per trial, I observed 2521 (25.21%) heads in position 2
In 10000 trials of 5 flips per trial, I observed 2489 (24.89%) heads in position 3
In 10000 trials of 5 flips per trial, I observed 2545 (25.45%) heads in position 4
In 10000 trials of 5 flips per trial, I observed 2497 (24.97%) heads in position 5
>>> coinflip(0.75,5,10000)
In 10000 trials of 5 flips per trial, I observed 7554 (75.54%) heads in position 1
In 10000 trials of 5 flips per trial, I observed 7443 (74.43%) heads in position 2
In 10000 trials of 5 flips per trial, I observed 7564 (75.64%) heads in position 3
In 10000 trials of 5 flips per trial, I observed 7487 (74.87%) heads in position 4
In 10000 trials of 5 flips per trial, I observed 7542 (75.42%) heads in position 5
```

```
>>> coinflip(0.75,8,100000)
In 100000 trials of 8 flips per trial, I observed 74980 (74.98%) heads in position 1
In 100000 trials of 8 flips per trial, I observed 75013 (75.013%) heads in position 2
In 100000 trials of 8 flips per trial, I observed 74995 (74.995%) heads in position 3
In 100000 trials of 8 flips per trial, I observed 75201 (75.201%) heads in position 4
In 100000 trials of 8 flips per trial, I observed 75137 (75.137%) heads in position 5
In 100000 trials of 8 flips per trial, I observed 74912 (74.912%) heads in position 6
In 100000 trials of 8 flips per trial, I observed 75058 (75.058%) heads in position 7
In 100000 trials of 8 flips per trial, I observed 74723 (74.723%) heads in position 8
```

What proportion of N trials of k coin flips have k consecutive heads?

- How does `foundallconsecutiveheads` inside the `for` blocks keep track if all the flips in the given trial are heads?
- Why does `if flip > headsprobability` make sense? Compare against the previous example.
- What is the difference between `=` and `==`?

```
>>> def consecutiveheads(headsprobability, numflips, numtrials):
        if headsprobability > 1 or headsprobability < 0:
                raise ValueError("That's not a valid value for a probability!")
        foundallconsecutiveheads = True
        consecutiveheadscounter = 0
        for trialindex in range(numtrials):
                for flipindex in range(numflips):
                        flip = random.random()
                        if flip > headsprobability:
                                foundallconsecutiveheads = False
                                break
                if foundallconsecutiveheads == True:
                        consecutiveheadscounter += 1
                foundallconsecutiveheads = True
        print("In " + str(numtrials) + " trials of ", end = "")
        print(str(numflips) + " flips per trial, I observed ", end = "")
        print(str(consecutiveheadscounter), end = "")
        print(" (" + str(consecutiveheadscounter/numtrials*100) + "%)", end = "")
        print(" of consecutive heads")


>>> consecutiveheads(0.5,4,10000)
In 10000 trials of 4 flips per trial, I observed 627 (6.27%) of consecutive heads
>>> consecutiveheads(0.5,10,10000)
In 10000 trials of 10 flips per trial, I observed 9 (0.09%) of consecutive heads
>>> consecutiveheads(0.5,10,100000)
In 100000 trials of 10 flips per trial, I observed 114 (0.114%) of consecutive heads
>>> consecutiveheads(0.5,10,1000000)
In 1000000 trials of 10 flips per trial, I observed 987 (0.0987%) of consecutive heads
>>> consecutiveheads(0.75,10,1000000)
In 1000000 trials of 10 flips per trial, I observed 56388 (5.6388%) of consecutive heads
>>> consecutiveheads(0.25,2,1000000)
In 1000000 trials of 2 flips per trial, I observed 62387 (6.2387%) of consecutive heads
>>> consecutiveheads(0.5,2,1000000)
```

```
In 1000000 trials of 2 flips per trial, I observed 251228 (25.1228%) of consecutive heads
```

Summary

In this part, only two exercises were given. But there is a lot of depth here. Instructors are encouraged to explore this further. Some suggestions:

- Discuss how changing the probability affects the observations. Does it make sense?
- Slowly introduce the notion of a Bernoulli trial and a binomial distribution. Show how the theoretical and observed probabilities match up.
- Ask: How would anything change if instead of coin flips, these were dice rolls? Since we are looking at binary events (Bernoulli trials) nothing should change in the code!
- What other patterns can be checked? Suppose instead of consecutive heads, what about alternating heads and tails? What is the probability of that event? How does the code change?

If you need help, have questions, or would like to set up a workshop at your school get in touch with me at help@mathmisery.com.