

Lab 4 - Probability

OpenIntro Lab

This includes lab materials available in the Probability lab available through OpenIntro, which was adapted for OpenIntro by Andrew Bray and Mine Cetinkaya-Rundel from a lab written by Mark Hansen of UCLA Statistics. These materials are available at OpenIntro: Probability.

Probability

Hot Hands

Basketball players who make several baskets in succession are described as having a hot hand. Fans and players have long believed in the hot hand phenomenon, which refutes the assumption that each shot is independent of the next. However, a 1985 paper by Gilovich, Vallone, and Tversky collected evidence that contradicted this belief and showed that successive shots are independent events (<http://www.cs.colorado.edu/~mozer/Teaching/syllabi/7782/readings/gilovich%20vallone%20tversky.pdf>). This paper started a great controversy that continues to this day, as you can see by Googling *hot hand basketball*.

We do not expect to resolve this controversy today. However, in this lab we'll apply one approach to answering questions like this. The goals for this lab are to (1) think about the effects of independent and dependent events, (2) learn how to simulate shooting streaks in R, and (3) to compare a simulation to actual data in order to determine if the hot hand phenomenon appears to be real.

Getting Started

Our investigation will focus on the performance of one player: Kobe Bryant of the Los Angeles Lakers. His performance against the Orlando Magic in the 2009 NBA finals earned him the title *Most Valuable Player* and many spectators commented on how he appeared to show a hot hand. Let's load some data from those games and look at the first several rows. Run this chunk - no need to change anything.

```
## load libraries
library(ggplot2)

## load data
download.file("http://www.openintro.org/stat/data/kobe.RData", destfile = "kobe.RData")
load("kobe.RData")
head(kobe)
```

```
##      vs game quarter time
## 1  ORL      1         1 9:47
## 2  ORL      1         1 9:07
## 3  ORL      1         1 8:11
## 4  ORL      1         1 7:41
## 5  ORL      1         1 7:03
## 6  ORL      1         1 6:01
##
##                                description basket
## 1                Kobe Bryant makes 4-foot two point shot      H
## 2                Kobe Bryant misses jumper                    M
```

```
## 3                Kobe Bryant misses 7-foot jumper      M
## 4 Kobe Bryant makes 16-foot jumper (Derek Fisher assists) H
## 5                Kobe Bryant makes driving layup      H
## 6                Kobe Bryant misses jumper            M
```

In this data frame, every row records a shot taken by Kobe Bryant. If he hit the shot (made a basket), a hit, H, is recorded in the column named basket, otherwise a miss, M, is recorded.

Just looking at the string of hits and misses, it can be difficult to gauge whether or not it seems like Kobe was shooting with a hot hand. One way we can approach this is by considering the belief that hot hand shooters tend to go on shooting streaks. For this lab, we define the length of a shooting streak to be the *number of consecutive baskets made until a miss occurs*.

For example, in Game 1 Kobe had the following sequence of hits and misses from his nine shot attempts in the first quarter: H M | M | H H M | M | M | M

To verify this, look at just the first nine entries in the column ‘basket’. (There are multiple ways to do this, but the code you write should only produce a vector of length 9 that shows the same sequence of H and M as shown above). Note, the vertical lines will not show up in your R output.

Answer:

Within the nine shot attempts, there are six streaks, which are separated by a “|” above. Their lengths are one, two, zero, zero, zero (in order of occurrence).

What does a streak length of 1 mean, i.e. how many hits and misses are in a streak of 1? What about a streak length of 0?

Answer:

DELETE THIS LINE AND ANSWER THIS QUESTION USING TEXT. NO R REQUIRED.

It is not always the case that functions are built into R or loaded through packages. It is possible to write functions that complete a specific task that you care about. Here, the custom function `calc_streak`, which was loaded in with the data, may be used to calculate the lengths of all shooting streaks and then look at this distribution.

Run the following code to use the `calc_streak` function to calculate the lengths of all shooting streaks and return a dataframe.

Answer:

```
## Calculate streak lengths
kobe_streak <- data.frame(streak = calc_streak(kobe$basket))
```

Note that we opted to make a barplot instead of histogram for the streak data. This is preferable because we are dealing with small counts here.

Make a barplot to look at the distribution. Then, in words, describe the distribution of Kobe’s streak lengths from the 2009 NBA finals. What was his typical streak length? How long was his longest streak of baskets?

Answer:

```
## Make barplot using ggplot
```

Compared to What?

We've shown that Kobe had some long shooting streaks, but are they long enough to support the belief that he had hot hands? What can we compare them to?

To answer these questions, let's return to the idea of *independence*. Two processes are independent if the outcome of one process doesn't effect the outcome of the second. If each shot that a player takes is an independent process, having made or missed your first shot will not affect the probability that you will make or miss your second shot.

A shooter with a hot hand will have shots that are *not* independent of one another. Specifically, if the shooter makes his first shot, the hot hand model says he will have a *higher* probability of making his second shot.

Let's suppose for a moment that the hot hand model is valid for Kobe. During his career, the percentage of time Kobe makes a basket (i.e. his shooting percentage) is about 45%, or in probability notation,

$$P(\text{shot } 1 = H) = 0.45.$$

If he makes the first shot and has a hot hand (*not* independent shots), then the probability that he makes his second shot would go up to, let's say, 60%,

$$P(\text{shot } 2 = H | \text{shot } 1 = H) = 0.60.$$

As a result of these increased probabilities, you'd expect Kobe to have longer streaks. Compare this to the skeptical perspective where Kobe does *not* have a hot hand, where each shot is independent of the next. If he hit his first shot, the probability that he makes the second is still 0.45.

$$P(\text{shot } 2 = H | \text{shot } 1 = H) = P(\text{shot } 2 = H) = 0.45.$$

In other words, making the first shot did nothing to effect the probability that he'd make his second shot. If Kobe's shots are independent, then he'd have the same probability of hitting every shot regardless of his past shots: 45%.

Now that we've phrased the situation in terms of independent shots, let's return to the question: how do we tell if Kobe's shooting streaks are long enough to indicate that he has hot hands? We can compare his streak lengths to someone without hot hands: an independent shooter.

Simulations in R

While we don't have any data from a shooter we know to have independent shots, that sort of data is very easy to simulate in R. In a simulation, you set the ground rules of a random process and then the computer uses random numbers to generate an outcome that adheres to those rules. As a simple example, you can simulate flipping a fair coin with the following:

```
outcomes <- c("heads", "tails")
sample(outcomes, size = 1, replace = TRUE)

## [1] "tails"
```

The vector `outcomes` can be thought of as a hat with two slips of paper in it: one slip says heads and the other says tails. The function `sample` draws one slip from the hat and tells us if it was a head or a tail.

Run the second command listed above several times. Just like when flipping a coin, sometimes you'll get a heads, sometimes you'll get a tails, but in the long run, you'd expect to get roughly equal numbers of each.

If you wanted to simulate flipping a fair coin 100 times, you could either run the function 100 times or, more simply, adjust the size argument (within the sample function), which governs how many samples to draw (the `replace = TRUE` argument indicates we put the slip of paper back in the hat before drawing again). Save the resulting vector of heads and tails in a new object called `sim_fair_coin`.

Answer:

To view the results of this simulation, type the name of the object and then use `table` to count up the number of heads and tails (Use `table()` with the argument `sim_fair_coin`).

Answer:

Since there are only two elements in outcomes ("heads" and "tails"), the probability that we "flip" a coin and it lands heads is 0.5, because the sample function assigns equal probability to outcomes by default. Say we're trying to simulate an unfair coin that we know only lands heads 20% of the time. We can adjust for this by adding an argument called `prob=c(0.2,0.8)`, which provides a vector of two probability weights. This is what you will do in your Homework 4 assignment.

Simulate flipping this unfair coin 100 times (that has a 20% chance of heads, 80% of tails). You can modify your code that you used to create `sim_fair_coin` by adding the `prob` argument as described. Save the results in a new object called `sim_unfair_coin`.

Answer:

`prob=c(0.2, 0.8)` indicates that for the two elements in the outcomes vector, we want to select the first one, heads, with probability 0.2 and the second one, tails with probability 0.8. Another way of thinking about this is to think of the outcome space as a bag of 10 chips, where 2 chips are labeled "head" and 8 chips "tail". Therefore at each draw, the probability of drawing a chip that says "head" is 20%, and "tail" is 80%.

In your simulation of flipping the unfair coin 100 times, how many flips came up heads?

Answer:

DELETE THIS LINE AND ANSWER THIS QUESTION USING TEXT. NO R REQUIRED.

In a sense, we've shrunken the size of the slip of paper that says "heads", making it less likely to be drawn and we've increased the size of the slip of paper saying "tails", making it more likely to be drawn. When we simulated the fair coin, both slips of paper were the same size. This happens by default if you don't provide a `prob` argument; all elements in the outcomes vector have an equal probability of being drawn.

If you want to learn more about `sample` or any other function, recall that you can always check out its help file (`?sample`).

Simulating the Independent Shooter

Simulating a basketball player who has independent shots uses the same mechanism that we use to simulate a coin flip. To simulate a single shot from an independent shooter with a shooting percentage of 50% we type:

```
outcomes <- c("H", "M")
sim_basket <- sample(outcomes, size = 1, replace = TRUE)
```

To make a valid comparison between Kobe and our simulated independent shooter, we need to align both their shooting percentage and the number of attempted shots.

What change needs to be made to the `sample` function so that it reflects a shooting percentage of 45%? Make this adjustment, then run a simulation to sample 133 shots. Assign the output of this simulation to a new object called `sim_basket`.

Answer:

Note that we've named the new vector `sim_basket`, the same name that we gave to the previous vector reflecting a shooting percentage of 50%. *In this situation, R overwrites the old object with the new one, so always make sure that you don't need the information in an old vector before reassigning its name.*

With the results of the simulation saved as `sim_basket`, we have the data necessary to compare Kobe to our independent shooter. We can look at Kobe's data alongside our simulated data.

With the results of the simulation saved as `sim_basket`, we have the data necessary to compare Kobe to our independent shooter. We can look at Kobe's data alongside our simulated data. Compare the distributions of `sim_basket` and `kobe$basket`. You can do this by creating tables or barplots.

Answer:

Both data sets represent the results of 133 shot attempts, each with the same shooting percentage of 45%. We know that our simulated data is from a shooter that has independent shots. That is, we know the simulated shooter does not have a hot hand.

Comparing Kobe Bryant to the Independent Shooter

Using `calc_streak`, compute the streak lengths of `sim_basket`. Save the result in a data frame (as you did with `kobe_streak`), but call this new data frame `indep_streak`.

Answer:

Describe the distribution of streak lengths. What is the typical streak length for this simulated independent shooter with a 45% shooting percentage? How long is the player's longest streak of baskets in 133 shots? Use R to find typical streak length and longest streak. You can use a table or a barchart for this.

Answer:

DELETE THIS LINE AND ANSWER THIS QUESTION USING TEXT. NO R REQUIRED.

If you were to run the simulation of the independent shooter a second time, how would you expect its streak distribution to compare to the distribution from the question above? Exactly the same? Somewhat similar? Totally different? Explain your reasoning.

Answer:

DELETE THIS LINE AND ANSWER THIS QUESTION USING TEXT. NO R REQUIRED.

How does Kobe Bryant's distribution of streak lengths, which you looked at in the last part of the previous subsection, compare to the distribution of streak lengths for the simulated shooter? Using this comparison, do you have evidence that the hot hand model fits Kobe's shooting patterns? Explain.

Answer:

DELETE THIS LINE AND ANSWER THIS QUESTION USING TEXT. NO R REQUIRED.