Maximum Likelihood Methods for the Social Sciences

POLS 510 · CSSS 510

# Monty Hall Monte Carlo

Political Science *and* CSSS
University of Washington, Seattle

Christopher Adolph

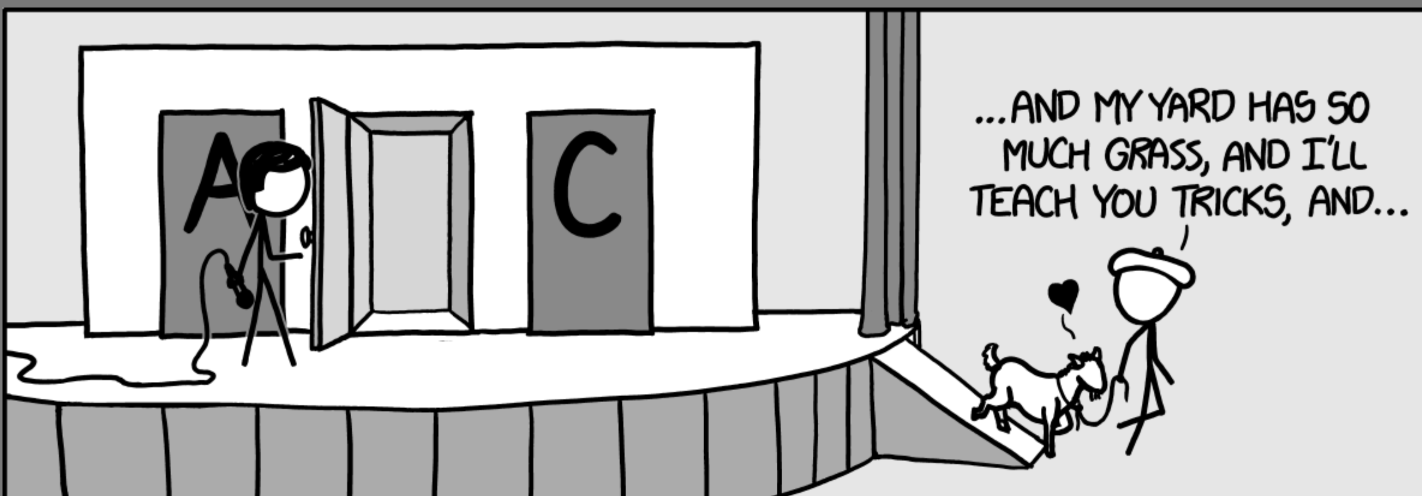


A few minutes later, the goat from behind door C drives away in the car

Randall Munroe xkcd.com/1282

# The Monty Hall Problem

On the 1970s version of the TV show *Let's Make a Deal*,
host Monty Hall presents contestants the following (possibly apocryphal) game:

1. There are 3 doors: $A$, $B$, $C$
2. Behind one door is a car; behind the other two are goats
3. The contestant choose a door (it stays closed for now)
4. Monty opens one of the two other doors to reveal a goat
5. Contestant's final choice: open the door chosen in step 3
   or switch to the third door

What should the contestant do to maximize the chance of winning a car?

## A longer example: Monte Hall

What is the probability problem here?

What is the probability of winning a car from staying?

What is the probability of winning a car from switching?

How can we solve the problem?

1. Probability theory: Bayes Rule (also known as "math")
2. Brute force: stochastic simulation (also known as Monte Carlo)

The first gets harder for hard problems, sometimes too hard to ever solve

The second never get any harder but always yields less general solutions

# A longer example: Monty Hall

Pseudo-code: A sketch of the solution

```
# Set up the doors, goats, and car

# Contestant picks a door

# Monty ``picks'' a remaining door

# Record where the car and goats were

# Do all of the above many many times

# Print the fraction of times a car was found
```

We act this out with slips of paper on a table and get the right answer

But computers is much faster, so let's look at some `R` code

## Monty Hall: Easy-to-read solution

```
# Monty Hall Problem
# Chris Adolph
# 1/6/2005

sims <- 10000                 # Simulations run
doors <- c(1,0,0)             # The car (1) and the goats (0)
cars.chosen <- 0              # Save cars from first choice here
cars.rejected <- 0            # Save cars from switching here

for (i in 1:sims) {           # Loop through simulations

  # First, contestant picks a door
  first.choice <- sample(doors, 3, replace=FALSE)

  # Choosing a door means rejecting the other two
  chosen <- first.choice[1]
  rejected <- first.choice[2:3]

  # Monty Hall removes a goat from the rejected doors
  rejected <- sort(rejected)
```

```
  if (rejected[1]==0)
    rejected <- rejected[2]

  # Record keeping:  where was the car?
  cars.chosen <- cars.chosen + chosen
  cars.rejected <- cars.rejected + rejected

}

cat("Probability of a car from staying with 1st door",
    cars.chosen/sims,"\n")

cat("Probability of a car from switching to 2nd door",
    cars.rejected/sims,"\n")
```

## Monty Hall: Less code solution

```
# Monty Hall Problem
# Chris Adolph
# 1/6/2005

sims <- 10000                   # Simulations run
doors <- c(1,0,0)               # The car (1) and the goats (0)
cars.chosen <- 0                # Save cars from first choice here
cars.rejected <- 0              # Save cars from switching here

for (i in 1:sims) {
  first.choice <- sample(doors, 3, replace=FALSE)
  cars.chosen <- cars.chosen + first.choice[1]
  cars.rejected <- cars.rejected + sort(first.choice[2:3])[2]
}

cat("Probability of a car from staying with 1st door",
    cars.chosen/sims,"\n")

cat("Probability of a car from switching to 2nd door",
    cars.rejected/sims,"\n")
```

## Monty Hall: Faster runtime solution

```
# Monty Hall Problem
# Chris Adolph
# 10/8/2013

sims <- 10000                    # Simulations run
doors <- c(1,0,0)                # The car (1) and the goats (0)

# Faster: avoiding loops with lapply()
result <- lapply (1:sims,
                  function (x, doors) {
                         pick <- sample(doors, 3, replace=FALSE);
                         c(pick[1],max(pick[2:3])) },
                  doors)

# Combine the list of results into a matrix
result <- do.call(rbind, result)

# Take the average of each column
result <- apply(result, 2, mean)
```

# Monty Hall: Faster runtime solution

```
cat("Probability of a car from staying with 1st door",
    result[1],"\n")

cat("Probability of a car from switching to 2nd door",
    result[2],"\n")
```

## Monty Hall: Intuitive solution

Can we explain the Monty Hall problem intuitively?

## Monty Hall: Intuitive solution

Can we explain the Monty Hall problem intuitively?

Without loss of generality, suppose the contestant selected door $A$ initially

The key is to notice Monty is *filtering out a goat* from the two remaining doors:

- Assume probability of a car behind each door is $1/3$ ex ante
- Collectively, the total probability of car behind door $B$ *or* $C$ is $2/3$
- By revealing a goat, Monty shows us where car must be if it is behind $B$ or $C$
- In effect, Monty is giving us a choice to take *any* car behind door $A$ *or* any car behind doors $B$ or $C$

## Monty Hall & Bayes Rule

We solved Monty Hall by brute force [in `R`] – could we solve it with math?

Recall that for two events $a$ and $b$,

$$\text{conditional probability} \quad = \quad \frac{\text{joint probability}}{\text{marginal probability}}$$

## Monty Hall & Bayes Rule

We solved Monty Hall by brute force [in `R`] – could we solve it with math?

Recall that for two events $a$ and $b$,

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

formally $$\mathrm{P}(a|b) = \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(b)}$$

# Monty Hall & Bayes Rule

We solved Monty Hall by brute force [in `R`] – could we solve it with math?

Recall that for two events $a$ and $b$,

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

$$\text{formally} \qquad \mathrm{P}(a|b) = \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(b)}$$

$$\text{which implies} \qquad \mathrm{P}(b|a) = \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(a)}$$

# Monty Hall & Bayes Rule

We solved Monty Hall by brute force [in `R`] – could we solve it with math?

Recall that for two events $a$ and $b$,

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

$$\begin{aligned}
\text{formally} \qquad \mathrm{P}(a|b) &= \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(b)} \\
\text{which implies} \qquad \mathrm{P}(b|a) &= \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(a)} \\
\mathrm{P}(b|a)\mathrm{P}(a) &= \mathrm{P}(a \cap b)
\end{aligned}$$

## Monty Hall & Bayes Rule

We solved Monty Hall by brute force [in `R`] – could we solve it with math?

Recall that for two events $a$ and $b$,

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

$$\begin{aligned}
\text{formally} \qquad \mathrm{P}(a|b) &= \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(b)} \\
\text{which implies} \qquad \mathrm{P}(b|a) &= \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(a)} \\
\mathrm{P}(b|a)\mathrm{P}(a) &= \mathrm{P}(a \cap b) \\
\mathrm{P}(a|b)\mathrm{P}(b) &= \mathrm{P}(a \cap b)
\end{aligned}$$

## Monty Hall & Bayes Rule

We solved Monty Hall by brute force [in R] – could we solve it with math?

Recall that for two events $a$ and $b$,

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

formally
$$P(a|b) = \frac{P(a \cap b)}{P(b)}$$

which implies
$$P(b|a) = \frac{P(a \cap b)}{P(a)}$$

$$P(b|a)P(a) = P(a \cap b)$$

$$P(a|b)P(b) = P(a \cap b)$$

$$P(a|b)P(b) = P(b|a)P(a)$$

# Monty Hall & Bayes Rule

We solved Monty Hall by brute force [in R] – could we solve it with math?

Recall that for two events $a$ and $b$,

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

$$
\begin{aligned}
\text{formally} \qquad \mathrm{P}(a|b) &= \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(b)} \\
\text{which implies} \qquad \mathrm{P}(b|a) &= \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(a)} \\
\mathrm{P}(b|a)\mathrm{P}(a) &= \mathrm{P}(a \cap b) \\
\mathrm{P}(a|b)\mathrm{P}(b) &= \mathrm{P}(a \cap b) \\
\mathrm{P}(a|b)\mathrm{P}(b) &= \mathrm{P}(b|a)\mathrm{P}(a) \\
\text{Bayes Rule} \qquad \mathrm{P}(a|b) &= \frac{\mathrm{P}(b|a)\mathrm{P}(a)}{\mathrm{P}(b)}
\end{aligned}
$$

## Monty Hall & Bayes Rule (1)

Ex ante, the probability the car is behind each door is just

$$\Pr(A) = \Pr(B) = \Pr(C) = \frac{1}{3}$$

Because the contestant picks a door at random,

$$\Pr(\text{contestant's initial door has a car}) = \frac{1}{3}$$

For the sake of argument and to simplify notation,
suppose the contestant's initial selection is door $A$

# Monty Hall & Bayes Rule (2)

Now, Monty Hall opens a door (either $B$ or $C$) to reveal a goat

Let's assume (without loss of generality) that this is door $B$

We want to know the probability that the remaining door, $C$, hides a car *given Monty's exposure of door $B$*

Formally, this is the conditional probability $\Pr(C \text{ is a car} | \text{Monty opened } B)$

We have enough information to solve this by applying Bayes Rule:

$$\Pr(a|b) = \frac{\Pr(b|a) \times \Pr(a)}{\Pr(b)}$$

which in our case is:

$$\Pr(C \text{ is a car} | \text{Monty opened } B) = \frac{\Pr(\text{Monty opened } B | C \text{ is a car}) \times \Pr(C \text{ is a car})}{\Pr(\text{Monty opened } B)}$$

# Monty Hall & Bayes Rule (3)

$$\Pr(C \text{ is a car}|\text{Monty opened } B) = \frac{\Pr(\text{Monty opened } B|C \text{ is a car}) \times \Pr(C \text{ is a car})}{\Pr(\text{Monty opened } B)}$$

The conditional probability Monty opens door $B$ if remaining door $C$ holds the car:
*by the rules of the game, this must be* $1$*; by design, Monty never shows the car*

The marginal probability that $C$ holds the car:
*this is* $1/3$*, because the car is randomly assigned a door*

The marginal probability that Monty would open $B$ rather than $C$:
*without conditioning on* $C$*'s contents, this, of course, is* $1/2$

Substituting into Bayes Rule, we find

$$\Pr(C \text{ is a car}|\text{Monty opened } B) = \frac{1 \times \frac{1}{3}}{\frac{1}{2}} = \frac{2}{3}$$