

Red balls, Green Balls

2024-08-29

This work delves into a probability puzzle involving an urn with an unknown number of red and green balls. Initially, the number of red balls is randomly selected. After observing a red ball on the first draw, the likelihood of drawing another red ball is calculated using Bayesian probability, which adjusts the probability distribution over possible urn configurations. The puzzle is further explored by modifying the setup, introducing a scenario where the number of red balls is determined by repeated coin flips, resulting in a binomial distribution. Theoretical results suggest a $2/3$ probability of drawing a red ball again in the original setup, and a $1/2$ probability when using the binomial model. These outcomes are validated through simulations, highlighting the nuanced effects of initial conditions on probability outcomes and underscoring the importance of rigorous reasoning in probabilistic analysis.

blog: <https://tetraquark.vercel.app/posts/redballgreenball/>

email: quarktetra@gmail.com

Introduction

Probability is a fascinating yet often perplexing field because it deals with uncertainty and randomness, concepts that can be challenging to intuitively grasp. Many probability problems appear simple on the surface but reveal surprising and counterintuitive results upon closer inspection. For example, the Monty Hall problem, where switching doors in a game show leads to better odds of winning, defies most people's gut instincts. This highlights the importance of rigorous mathematical analysis in probability. Without careful calculation and logical reasoning, our intuitive judgments can lead us astray, underscoring the need for a systematic approach to understanding and applying probability in real-world situations. Today, we will tackle such a problem.

Problem Statement

Quoting from Quanta magazine, [1]

” Imagine, that you have an urn filled with 100 balls, some red and some green. You can’t see inside; all you know is that someone determined the number of red balls by picking a number between zero and 100 from a hat. You reach into the urn and pull out a ball. It’s red. If you now pull out a second ball, is it more likely to be red or green (or are the two colors equally likely)? “

— Daniel Litt

Let us define the number of balls in the urn as N . Originally, there are $N + 1$ different possible configurations for the urns with the number of red balls from 0 to N . Let us label the urns as u_i , with $i \in \{0, 1, 2, \dots, N\}$. Since the number of red balls is pulled from a uniform distribution, the probability of getting any u_i is simply $\frac{1}{N+1}$.

First, we will shut-up and apply the rigorous machinery of probability theory. Second, we will do a simulation to double check the results.

Shut up and calculate

Given the first ball is red, the probability that it came from urn u_i (i.e., an urn with i red balls) is

$$P(u_i|x_1 = R) = \frac{P(x_1 = R|u_i)P(u_i)}{P(x_1 = R)}, \quad (1)$$

which is the Bayesian formula. We can compute the denominator as

$$P(x_1 = R) = \sum_{i=0}^N P(x_1 = R|u_i)P(u_i) = \sum_{i=0}^N \frac{i}{N} \frac{1}{N+1} = \frac{1}{2}. \quad (2)$$

And the numerator as

$$P(x_1 = R|u_i)P(u_i) = \frac{i}{N} \frac{1}{N+1} = \frac{i}{N(N+1)}. \quad (3)$$

Putting Eqs. 2 and 3 back into Eq. 1 gives

$$P(u_i|x_1 = R) = \frac{2i}{N(N+1)}. \quad (4)$$

Note that this is very important. Originally, the probability of getting an urn with i red balls was flat at $\frac{1}{N+1}$. However, the fact that first ball is red implies that it is more likely

that it came from an urn with more red balls, see Figure 1. Now that we have a probability distribution for having u_i , we can compute the probability of getting a red ball in the second draw:

$$\begin{aligned}
P(x_2 = R) &= \sum_{i=0}^N P(x_2 = R|u_i)P(u_i|x_1 = R) = \sum_{i=0}^N \frac{i-1}{N-1} \frac{2i}{N(N+1)} \\
&= \frac{2}{(N-1)N(N+1)} \sum_{i=0}^N (i^2 - i) \\
&= \frac{2}{(N-1)N(N+1)} \left[\frac{N(N+1)(2N+1)}{6} - \frac{N(N+1)}{2} \right] \\
&= \frac{2N+1}{3(N-1)} - \frac{1}{N-1} = \frac{2}{3}.
\end{aligned} \tag{5}$$

Therefore, getting a red ball is twice as likely as getting a green ball, independent of how many balls there are in the urn.

A tweak

The statement “All you know is that someone determined the number of red balls by picking a number between zero and 100 from a hat” is so critical. In order to show why that is, let’s revise the preparation of the urn as follows: “The urn is prepared by adding balls to the urn one by one by flipping a coin. If the coin is heads, you add a red ball, a green one otherwise.” This will completely transform the question. The number of red balls in the urn will have the binomial probability density:

$$f_b(i, N) = \binom{N}{i} p^i (1-p)^{N-i}. \tag{6}$$

Now, we just need to redo the math. Given the first ball is red, the probability that it came from urn u_i (i.e., an urn with i red balls) is

$$P(u_i|x_1 = R) = \frac{P(x_1 = R|u_i)P(u_i)}{P(x_1 = R)}, \tag{7}$$

which is still the Bayesian formula. We can compute the denominator as

$$P(x_1 = R) = \sum_{i=0}^N P(x_1 = R|u_i)P(u_i) = \sum_{i=0}^N f_b(i, N) \frac{i}{N} = \frac{1}{N} \langle i \rangle = p. \tag{8}$$

For a fair coin used in the preparation $p = 1/2$. The numerator reads

$$P(x_1 = R|u_i)P(u_i) = \frac{i}{N} f_b(i, N). \tag{9}$$

Putting Eqs. 8 and 9 back into Eq. 7 gives

$$P(u_i|x_1 = R) = \frac{1}{Np} i f_b(i, N). \quad (10)$$

Now that we have a probability distribution for having u_i , we can compute the probability of getting a red ball in the second draw:

$$\begin{aligned} P(x_2 = R) &= \sum_{i=0}^N P(x_2 = R|u_i)P(u_i|x_1 = R) \\ &= \frac{1}{pN(N-1)} \sum_{i=0}^N (i-1) i f_b(i, N) = \frac{1}{pN(N-1)} (\langle i^2 \rangle - \langle i \rangle) \\ &= \frac{1}{pN(N-1)} (\langle i \rangle^2 + \sigma^2 - Np) = \frac{1}{pN(N-1)} (N^2 p^2 + Np(1-p) - Np) \\ &= \frac{1}{pN(N-1)} N(N-1)p^2 \\ &= p, \end{aligned} \quad (11)$$

which is the probability of the coin that was used to build the urn! Therefore, getting a red ball in the second draw is as probable as getting a green one provided that the coin was unbiased.

Simulation

Here is a simulation code written in R. If you are interested in playing with the simulation, you can copy it below.

```
# A code to simulate the problem:
# https://www.quantamagazine.org/perplexing-the-web-one-probability-puzzle-at-a-time-2024082
# Find the math at
#https://tetraquark.vercel.app/post/redballgreenball/redballgreenball/index.html?src=githubbr
library(plotly)
colorize <- function(x) {
  if (x < 0) {
    "R"
  } else {
    "G"
  }
} # will use this to map +1,-1 to colors
probBalls <- 0.5 # Set this as a global variable.
```

```

Nballs <- 100
Nsim <- 40000
# set the number of balls in an urn, and the number of simulations
t2 <- list(size = 20)
wbinomProb <- function(x) {
  x / Nballs * dbinom(x, size = Nballs, prob = probBalls)
} # will use this to map +1,-1 to colors
binomProb <- function(x) {
  dbinom(x, size = Nballs, prob = probBalls)
} # will use this to map +1,-1 to colors

firstBall <- c()
secondBall <- c()
gBalls <- c()
rBalls <- c()
BallsOneTwo <- c() # initialize arrays to store various values

simulator <- function(drawMode) {
  for (s in c(1:Nsim)) {
    if (drawMode == "uniform") {
      randomRedCount <- sample.int(Nballs, 1) # this is how the urn is set up: number of red
      balls <- c(rep(-1, randomRedCount), rep(1, Nballs - randomRedCount)) # repeat -1 random
    }
    if (drawMode == "binomial") {
      balls <- 2 * rbinom(Nballs, 1, probBalls) - 1
    }
    ballsC <- sapply(balls, colorize) # map numbers to colors
    NumOfGreens <- round(Nballs / 2 + sum(balls) / 2) # compute the number of Green balls
    gBalls <- c(gBalls, NumOfGreens) # log the value for the green balls in this urn
    rBalls <- c(rBalls, Nballs - NumOfGreens) # log the value for the red balls in this urn

    randomIndex <- sample.int(length(balls), 1) # draw a random index, this will select a ball
    thisfirstBall <- ballsC[randomIndex] # color of the first ball
    firstBall <- c(firstBall, thisfirstBall) # log the first ball color
    ballsC <- ballsC[-randomIndex] # remove this ball from the urn.

    randomIndex <- sample.int(length(ballsC), 1) # draw another random index for the second ball
    thissecondBall <- ballsC[randomIndex] # color of the second ball
    secondBall <- c(secondBall, thissecondBall) # log the second ball
    BallsOneTwo <- c(BallsOneTwo, paste0(thisfirstBall, "_", thissecondBall)) # Create a pair
  }
}
dtBall <- data.frame("gBallsC" = gBalls, "rBallsC" = rBalls, "firstBall" = firstBall, "secondBall" = secondBall)

```

```

dtBallS <- dtBall[dtBall$firstBall == "R", ] # we are told that the first ball is red, so,
redRatio <- nrow(dtBallS[dtBallS$secondBall == "R", ]) / nrow(dtBallS) # simply compute the
xv <- c(0:Nballs)
if (drawMode == "uniform") {
  titleText <- "Uniform Preparation"

  yv <- 100*2 * xv / (Nballs * (Nballs + 1))
  rAVGtheory <- 100*2 / 3
  yv0 <- 100*rep(1 / (1 + Nballs), length(xv))
} # theoretical formula; https://tetraquark.vercel.app/post/redballgreenball/redballgreenball
if (drawMode == "binomial") {
  titleText <- "Binomial Preparation"
  yv <- sapply(xv, wbinomProb)
  yv0 <- sapply(xv, binomProb)
  yv0 <- 100*yv0 / sum(yv0)

  yv <- 100*yv / sum(yv)
  rAVGtheory <- probBalls
} # theoretical formula; https://tetraquark.vercel.app/post/redballgreenball/redballgreenball

# Pre-calculate histogram data
hist_data <- hist(dtBallS$rBalls, breaks = seq(-0.5, Nballs + 0.5, by = 1), plot = FALSE)
hist_counts <- 100 * hist_data$counts / sum(hist_data$counts) # Convert to percentages using
hist_mids <- hist_data$mids

m <- list(l = 70, r = 10, b = 70, t = 20, pad = 8)
figOut <- plot_ly(alpha = 0.5) %>%
  add_bars(x = ~hist_mids,
           y = ~hist_counts,
           name = "<span style='color:blue;'>Simulation</span>",
           marker = list(color = 'blue', opacity = 0.5)) %>%
  add_trace(x = ~xv,
            y = ~( yv),
            name = "<span style='color:orange;'>Revised</span>",
            type = "scatter",
            mode = "markers+lines") %>%
  layout(
    hovermode = "x",
    xaxis = list(
      hoverformat = ".0f",
      title = list(text = "Number of Red Balls")
    ),
  ),

```

```

yaxis = list(
  hoverformat = ".2r",
  title = list(text = "Probability Density (%)")
),
plot_bgcolor = "rgba(0, 0, 0, 0)",
paper_bgcolor = "rgba(0, 0, 0, 0)",
bargap = 0
) %>%
layout(title = list(y = 0.95, x = 0.2, text = paste0(titleText), font = t2)) %>%
layout(legend = list(x = 0.1, y = 0.85, orientation = "v", font = t2), margin = m) %>%
config(mathjax = "cdn") # Changed from TRUE to "cdn"
# figOut%>%config(mathjax = "cdn"). # enable this to render MathJax
figOut <- figOut %>% add_trace(x = xv, y = yv0, name = "<span style='color:green;'>Original")

output <- list(figOut, redRatio, rAVGtheory)
return(output)
}

if (FALSE) {
  drawMode <- "uniform"
  # drawMode="binomial";
  returns <- simulator(drawMode)
  redRatioFormatted <- formatC(signif(returns[[2]], digits = 3), digits = 3, format = "fg",
  redRatioThFormatted <- formatC(signif(returns[[3]], digits = 3), digits = 3, format = "fg"
  print(paste0("Nballs:", Nballs, ";", drawMode, "-->probability of second ball being red is
  returns[[1]]
}

```

The code simulates the original problem with 100 balls and repeats it 40000 times.

Figure 1: Green: original density, Orange: The probability distribution of having an urn with i red balls given that the first ball was red, Blue: simulation results.

We can now simply count the cases of second ball being red, and the simulation result is 0.673 ± 0.005 with 95% confidence level which is very close to the 66.7 value we predicted in Eq. 5.

Below is the tweaked problem where the balls are decided with coin flip.

Figure 2: Green: original density, Orange: The probability distribution of having an urn with i red balls inside that the first ball was red, Blue: simulation results.

We can now simply count the cases of second ball being red, and the simulation result is 0.501 ± 0.005 with 95% confidence level which is very close to the 0.5 value we predicted in Eq. 11. The color of the first ball isn't really a useful information, it barely moves the distribution. In fact, the shift towards the higher values of red ball count exactly cancels the fact that we threw away a red ball in the first draw. It is a perfect cancellation!

- [1] E. Klarreich, "Perplexing the web, one probability puzzle at a time," *Quanta Magazine*, 2024 [Online]. Available: <https://www.quantamagazine.org/perplexing-the-web-one-probability-puzzle-at-a-time-20240829/>