

The Monty Hall Problem

A Bayesian Solution, with a Monte Carlo Check

Hedibert Lopes (Assisted by Claude for Mac)

May 11, 2026

The Setup

Three doors. A car sits behind one; goats behind the other two.

1. You pick a door — say **Door 1**.
2. Monty (who knows where the car is) opens another door revealing a goat — say **Door 3**.
3. He offers you the chance to switch to **Door 2**.

Should you switch?

The Bayesian answer falls out mechanically once we write down the right conditional probabilities.

Setting Up the Problem

Let C_i denote “the car is behind door i .”

Before any door is opened, the prior is flat:

$$P(C_1) = P(C_2) = P(C_3) = \frac{1}{3}.$$

You have chosen Door 1. Let M_3 denote “Monty opens Door 3.”

The crucial step is the **likelihood** $P(M_3 | C_i)$, governed by Monty’s rules:

- He never opens *your* door.
- He never opens the door hiding the car.

The Likelihoods

Car location	Monty's options	$P(M_3 C_i)$
C_1 (your pick)	free choice between Doors 2 and 3	$\frac{1}{2}$
C_2	forced to open Door 3	1
C_3	cannot open the car door	0

This asymmetry — Monty has a free choice under C_1 but a forced move under C_2 — is the whole engine of the puzzle.

Applying Bayes' Theorem

The marginal probability of the observation:

$$P(M_3) = \sum_{i=1}^3 P(M_3 | C_i) P(C_i) = \frac{1}{2} \cdot \frac{1}{3} + 1 \cdot \frac{1}{3} + 0 \cdot \frac{1}{3} = \frac{1}{2}.$$

Posterior for sticking with Door 1:

$$P(C_1 | M_3) = \frac{P(M_3 | C_1) P(C_1)}{P(M_3)} = \frac{\frac{1}{2} \cdot \frac{1}{3}}{\frac{1}{2}} = \boxed{\frac{1}{3}}.$$

Posterior for switching to Door 2:

$$P(C_2 | M_3) = \frac{P(M_3 | C_2) P(C_2)}{P(M_3)} = \frac{1 \cdot \frac{1}{3}}{\frac{1}{2}} = \boxed{\frac{2}{3}}.$$

Switch.

Why Intuition Fails

The trap: once Door 3 is eliminated, the remaining doors *look* symmetric, so people guess $\frac{1}{2}$ each.

Bayes reveals exactly what's wrong:

$$P(M_3 | C_2) = 1 \quad \text{but} \quad P(M_3 | C_1) = \frac{1}{2}.$$

Observing M_3 is **twice as informative** about C_2 as about C_1 , and the posterior inherits the asymmetry.

Diagnostic. Replace Monty with a “Random Monty” who picks uniformly between the two doors you didn't pick — sometimes revealing the car. Conditioning on the event that he *happens* to reveal a goat behind Door 3:

$$P(C_1 | \text{goat at 3}) = P(C_2 | \text{goat at 3}) = \frac{1}{2}.$$

Switching no longer helps. The whole puzzle hinges on Monty's *information*.

Monte Carlo Check

A Monte Carlo Exercise

We will simulate N rounds of the game, always choosing Door 1 and always conditioning on the event M_3 (Monty opens Door 3). For each round we record where the car actually was. Empirical frequencies should approach

$$P(C_1 | M_3) \approx \frac{1}{3}, \quad P(C_2 | M_3) \approx \frac{2}{3}.$$

Equivalently: the win rate of the “always switch” strategy converges to $\frac{2}{3}$ and the “always stay” strategy to $\frac{1}{3}$.

Python Simulation

```
import numpy as np

rng = np.random.default_rng(42)
N = 1_000_000

# Car location uniformly at random
car = rng.integers(1, 4, size=N)          # 1, 2, or 3

# Contestant always picks Door 1
pick = 1

# Monty's choice: never the pick, never the car
def monty_open(c):
    if c == 1:          # car behind pick -> uniform over {2, 3}
        return rng.choice([2, 3])
    if c == 2:          # forced
        return 3
    return 2            # c == 3, forced

monty = np.array([monty_open(c) for c in car])

# Condition on the event "Monty opened Door 3"
mask = (monty == 3)
car_given_M3 = car[mask]
```

Vectorized Version (Faster)

```
import numpy as np
rng = np.random.default_rng(0)
N = 10_000_000

car = rng.integers(1, 4, size=N)

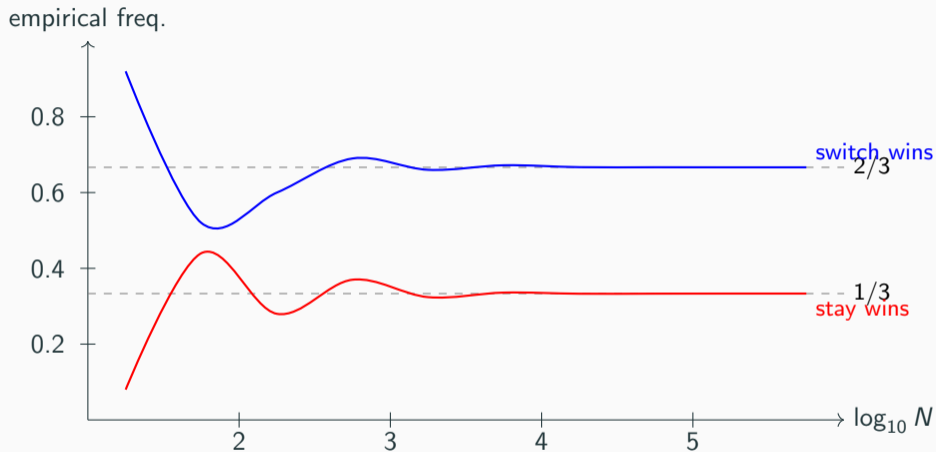
# When car==1, Monty picks uniformly from {2,3}; else forced
coin = rng.integers(0, 2, size=N)      # 0 or 1
monty = np.where(car == 1, 2 + coin,   # 2 or 3 with prob 1/2
                 np.where(car == 2, 3, 2)) # forced otherwise

mask = (monty == 3)
print("P(C1 | M3) ~", (car[mask] == 1).mean())
print("P(C2 | M3) ~", (car[mask] == 2).mean())
print("P(C3 | M3) ~", (car[mask] == 3).mean())
```

Typical output (10M draws):

$P(C1 | M3) \approx 0.3333$ $P(C2 | M3) \approx 0.6667$ $P(C3 | M3) = 0.0000$

Convergence Plot (Schematic)



Empirical frequencies tighten around $2/3$ and $1/3$ as N grows.

Takeaways

- The Monty Hall puzzle is **not** about counting doors; it's about conditioning on an observation whose likelihood depends on the hidden state.
- Bayes' theorem makes the asymmetry visible: $P(M_3 | C_2) = 1$ versus $P(M_3 | C_1) = \frac{1}{2}$.
- The posterior $2/3$ vs. $1/3$ is *not* a quirk — it is confirmed by simulation to arbitrary precision.
- Change Monty's protocol (e.g., random opening) and the likelihoods change — and so does the posterior. Information is everything.

Always switch.