
Hard Probability Questions With Solutions

*Hard
Probability
Questions
With
Solutions*

*Downloaded from
everybodystreet.com
by Heath & Tessa*

MASTERING HARD PROBABILITY QUESTIONS WITH SOLUTIONS: A COMPREHENSIVE GUIDE

Probability is the mathematical backbone of uncertainty, governing everything from weather forecasts to quantum mechanics. While many students breeze through basic coin tosses and dice rolls, the real challenge

emerges when problems become layered with conditional dependencies, combinatorics, and intricate sampling strategies. Hard probability questions with solutions are not just exam hurdles—they sharpen logical reasoning and prepare you for real-world decision-making under uncertainty. In this guide, we dissect several notoriously difficult probability problems, offering clear, step-by-step solutions. Whether you

are preparing for a contest, an interview, or simply want to upgrade your problem-solving toolkit, these examples will teach you to think probabilistically even in the most tangled scenarios.

What Makes a Probability Question “Hard”?

Before diving into specific problems, it helps to understand the common features that elevate a probability question from routine to

challenging:

- **Complex sample spaces**
– large or non-uniform outcomes that require combinatorial counting or generating functions.
- **Conditional probability traps** – hidden dependencies that trick even experienced students (e.g., the Monty Hall problem).
- **Non-standard distributions** – problems that go beyond binomial or normal approximations.
- **Multiple layers of conditioning**
– scenarios where you must apply Bayes’ theorem or chain

rule more than once.

- **Counterintuitive results** – answers that defy common sense, making verification crucial.

Mastering hard probability questions with solutions involves not only knowing formulas but also learning how to restate a problem in terms of simpler events. Below we present several problems that exemplify these challenges.

HARD PROBABILITY PROBLEM 1: THE BIRTHDAY PROBLEM EXTENDED

The Question

In a group of 30 people, what is the probability that at least two people share a birthday? That classic version is well known ($\approx 70.6\%$). Now consider a harder twist: **In a group of 40 people, what is the probability that exactly two pairs share a birthday?** (Assume 365 days, equally likely birthdays, no leap year.)

Solution Approach

This is a multinomial selection problem. We want the event that there are exactly two distinct birthday dates

each shared by exactly two people, and the remaining 36 people all have distinct birthdays (with no overlaps among the pairs or with the uniques).

Step 1: Choose the two dates that will be “pair dates.”

Number of ways: $\binom{365}{2}$.

Step 2: Choose which 4 people (out of 40) form the two pairs. Number of ways:

$\binom{40}{4}$.

However, we must then partition these four people into two unordered pairs. The number of ways to split 4 distinct individuals into 2 unlabeled pairs is

$\frac{4!}{2!2!} = 3$

(or $\frac{\binom{4}{2}}{2} = 3$). So total assignments = $\binom{40}{4} \times 3$.

Step 3: Assign the two chosen dates to the two pairs. Since the dates are distinct, we can assign the two sets of two people to the two dates in $2!$ ways.

Step 4: The remaining 36 people must all have distinct birthdays from each other and also distinct from the two pair dates.

So we choose 36 distinct dates from the remaining 363 days. Number of ways: $\binom{363}{36}$. Then assign these 36 dates to the 36 people in $36!$ orders.

Step 5: Calculate total possible assignments of 40 people to 365 birthdays. That is simply 365^{40} (each person independently chooses a day).

Thus the probability is:

$$\begin{aligned} P = & \frac{\binom{365}{2}}{\binom{40}{4}} \times 3 \times 2! \\ & \times \binom{363}{36} \\ & \times 36! \times 365^{40} \end{aligned}$$

We can simplify:
 $\frac{\binom{363}{36} \times 36!}{\binom{363}{36} \times 36!} = \frac{36!}{(363-36)!}$
 $= \frac{36!}{327!}$.
 Combining with other factorials yields a numerical value of approximately 0.00076, i.e., about 0.076%.

Insight: The probability is extremely small because requiring exactly two pairs (and no other shared birthdays) in a group of 40 is a very specific outcome compared to the “at least one birthday match” which is around 89%.

HARD PROBABILITY PROBLEM 2: THE MARTINGALE BETTING SYSTEM UNDER A CASINO LIMIT

The Question

A gambler uses a double-up betting strategy (Martingale) on a fair roulette wheel (European, no zero, for simplicity). He starts with a \$1 bet. If he loses, he doubles his bet for the next spin; if he wins, he stops for the session (he wins exactly the original \$1 profit). The casino imposes a maximum bet of \$256. What is the probability that the gambler will eventually win \$1 in a session, assuming he has an

initial bankroll that can cover up to 8 losses in a row (so his ninth bet would exceed the maximum)?

Note: Because the wheel is fair, the probability of winning or losing each spin is 0.5. The gambler stops as soon as he wins.

Solution Approach

The gambler will win if a win occurs within the first 8 bets. If he loses eight times in a row, on the ninth spin he would need to bet \$256 (since after 8 losses the bet would be $(2^8=256)$), but the casino maximum prevents him from betting \$512 (which would be needed after a ninth loss). Actually, careful: Starting bet

\$1. After 1 loss, bet \$2; after 2 losses, bet \$4; ... after 7 losses, bet \$128; after 8 losses, bet \$256 (still within limit). After 8 losses, if he loses again (9th loss), he would need to bet \$512, which is not allowed. So the session ends in a total loss if he loses 9 times in a row? Let's think: The strategy says he stops after a win. He can continue as long as his next bet does not exceed the limit. After 8 consecutive losses, he has lost $(1+2+4+\dots+256 = 511)$ dollars. Now for the 9th spin, the required bet is \$512, but the casino maximum is \$256, so he cannot double. The problem states he has an initial bankroll to cover up to 8 losses in a row – so he has at least \$511. However, if

he loses 8 times, he cannot make the required 9th bet because the bet size would exceed the limit. The session effectively ends with a net loss of \$511 (since he cannot recover). Therefore, the gambler wins \$1 if he wins at any point during the first 8 bets (i.e., before having lost 8 times).

Probability of winning on the first spin: 0.5

Probability of winning on the second spin (must lose first, then win): $(0.5 \times 0.5 = 0.25)$

And so on up to the 8th spin: probability of losing first 7 then winning 8th: $(0.5^7 \times 0.5 = 0.5^8 = 1/256)$.

Sum of these probabilities: $(0.5 + 0.5^2 + 0.5^3 + \dots + 0.5^8)$ is a geometric

series with first term 0.5, ratio 0.5, 8 terms. $\text{Sum} = \left(\frac{0.5(1-0.5^8)}{1-0.5} \right) = 1 - 0.5^8 = 1 - \frac{1}{256} = \frac{255}{256}$.

Hence the probability of winning \$1 is $255/256 \approx 0.9961$, or 99.6%.

But wait: if he loses 8 times in a row (probability $(0.5^8 = 1/256)$), he neither wins nor can continue, so he loses his \$511 bankroll. So the probability of a loss is $1/256$.

Lesson: Hard probability questions with solutions often expose the fallacy of “sure thing” systems – here the high probability of small gain is offset by a rare but catastrophic loss.

PROBABILITY PROBLEM 3: THE TWO-ENVELOPE PARADOX

RESOLVED

The Question

Two envelopes each contain a positive integer amount of money. One envelope contains twice as much as the other. You pick an envelope at random and open it to find \$100. You are then given the option to switch to the other envelope. Should you switch? The paradox arises because it seems that the expected value of the other envelope is either \$50 or \$200, with an average of \$125, so you should

switch. Yet if you had not opened the envelope, you would have no reason to switch. Resolve the paradox by computing the actual expected gain from switching, assuming that the pair of amounts $(x, 2x)$ is chosen by some prior distribution (e.g., x is equally likely to be any power of 2 from 1 to 2^n).

Solution Approach

Without an assumed prior distribution over the possible pairs, the problem is ill-posed. A common resolution is to observe that when you see \$100, you might be more likely to have the smaller envelope if 100 is a number that appears often as the smaller of

a pair. Let's assume a specific prior: The smaller amount X is drawn from a geometric distribution: $\mathbb{P}(X = 2^k) = 1/2^{k+1}$ for $k=0,1,2,\dots$ (so X can be $1,2,4,8,\dots$). This is a proper distribution summing to 1. Then the larger amount is $2X$. You pick an envelope uniformly at random.

Now, given you see \$100. The number 100 is not a power of two, so it cannot be the smaller amount in this prior? Actually, 100 is not of the form 2^k , so under this prior the probability of seeing \$100 is zero - a problem. Resolve by allowing continuous amounts or a different prior. A more standard approach: Suppose the pair is $(x, 2x)$ and x is drawn from a

continuous uniform distribution on $[1, A]$ for some large A . Then the probability of seeing an exact value y is zero, but we condition on " y belongs to some interval". For simplicity, we can treat an approximate argument: When you see \$100, it is equally likely to be the smaller (if the pair was $(100, 200)$) or the larger (if the pair was $(50, 100)$). But the prior probabilities of those two pairs are not equal because the underlying x distribution gives higher density to smaller x if it is uniform. Actually, if $x \sim \text{Uniform}[1, A]$, then the probability that x falls in $[50, 50+\epsilon]$ is about $\epsilon/(A-1)$, and the probability that x falls in $[100, 100+\epsilon]$ is also $\epsilon/(A-1)$. So the two pairs are equally likely

given the continuous uniform prior? Wait: The event “you see \$100” can arise either from the case where the smaller envelope contains 100 (and you picked the smaller) or the larger envelope contains 100 (and you picked the larger). The probability of the first scenario: prior density of x being near 100 times probability you pick the smaller (0.5). The second scenario: prior density of x being near 50 times probability you pick the larger (0.5). Since the density at 50 equals density at 100 under a uniform prior, the two possibilities have equal likelihood. Therefore, conditional on seeing 100, the probability that it is the smaller envelope is 0.5 (and similarly 0.5 that it is the larger). Then the expected value of the other envelope: if you have the smaller (probability 0.5), the other envelope is 200; if you have the larger (probability 0.5), the other envelope is 50. Expected value = $0.5 \cdot 200 + 0.5 \cdot 50 = 125$. So paradox remains! The resolution is that the prior cannot be uniform over an infinite range because it would be improper. With a proper prior (e.g., x