

This is a draft; please send me corrections and/or suggestions.

Gambling

You may have heard of the following gambling strategy, called the *Martingale strategy*: if you win a particular gamble, you collect your earnings; if you lose, you double the bet you just made and try again.

In a more formal setup, assume you are betting on the outcome of a fair coin, which lands heads-up with probability $\frac{1}{2}$ and tails-up with probability $\frac{1}{2}$. You first bet on this coin is \$2; if you win, you are done. But if you lose, you double your bet (to \$4) and try again. If you lose that bet, you double your bet again (to \$8) and try again. You repeat this process until you win a coin toss.

Let X be a random variable representing when you finally win the coin toss.

(a) *What is the PMF of X ?*

Solution: this is somewhat more difficult than it looks; we'll start with an intuitive answer, then make things more robust.

The probability that you win on the first toss of the coin is $\frac{1}{2}$, so $f_X(1) = \frac{1}{2} = \frac{1}{2^1}$. To win on the second toss, you must have lost the first toss (probability $1 - \frac{1}{2} = \frac{1}{2}$) and then the coin must land heads-up (probability $\frac{1}{2}$); so the probability that you win on the second toss is $f_X(2) = \frac{1}{4} = \frac{1}{2^2}$. To win on the third toss, you must have lost the first *two* tosses (probability $1 - \frac{1}{2} - \frac{1}{4} = \frac{1}{4}$) and then the coin must land heads-up (probability $\frac{1}{2}$); so the probability that you win on the third toss is $f_X(3) = \frac{1}{8} = \frac{1}{2^3}$.

You can see a pattern emerging here; in general, we have

$$f_X(x) = \frac{1}{2^x}.$$

This is sufficient for an answer; if you are interested in showing more precisely how to get here, keep reading.

How can we show this rigorously? The probability that you win the n^{th} bet, conditional on having lost the previous $n - 1$ bets is $\frac{1}{2}$: losing all the earlier bets does not affect the probability that the coin lands heads-up the n^{th} time. By expansion of conditional probability, we know

$$P(\text{win } n^{\text{th}} \text{ bet}) = P(\text{win } n^{\text{th}} \text{ bet} | \text{lose } n - 1 \text{ bets})P(\text{lose } n - 1 \text{ bets}) \rightsquigarrow P(X = n) = P(X = n | X > n - 1)P(X > n - 1).$$

By the argument above, $P(X = n | X > n - 1) = \frac{1}{2}$; so we need only determine $P(X > n - 1)$. Using the law of complements, we know $P(X > n - 1) = 1 - P(X \leq n - 1)$; this in turn is

$$P(X \leq n - 1) = P(\{X = 1\} \cup \dots \cup \{X = n - 1\}).$$

Since the events $\{X = 1\}, \dots, \{X = n - 1\}$ are mutually exclusive, we know that the probability of their union is the sum of the probabilities,

$$P(X \leq n - 1) = \sum_{x=1}^{n-1} P(X = x).$$

It follows then that

$$P(X = n) = \frac{1}{2} \left(1 - \sum_{x=1}^{n-1} P(X = x) \right).$$

At this point, we're stuck given the mathematical tools at our disposal. We will guess that $f_X(x) = (\frac{1}{2})^x$ and verify that it is the case. Under this assumption, we have

$$\begin{aligned}
 P(X = n) &= \frac{1}{2} \left(1 - \sum_{x=1}^{n-1} \left(\frac{1}{2}\right)^x \right) \\
 &= \frac{1}{2} \left(1 - \sum_{x=1}^{\infty} \left(\frac{1}{2}\right)^x + \sum_{x=n}^{\infty} \left(\frac{1}{2}\right)^x \right) \\
 &= \frac{1}{2} \left(1 - \frac{1}{2} \sum_{x=0}^{\infty} \left(\frac{1}{2}\right)^x + \left(\frac{1}{2}\right)^n \sum_{x=0}^{\infty} \left(\frac{1}{2}\right)^x \right) \\
 &= \frac{1}{2} \left(1 - \frac{1}{2} \left(\frac{1}{1 - \frac{1}{2}} \right) + \left(\frac{1}{2}\right)^n \left(\frac{1}{1 - \frac{1}{2}} \right) \right) \\
 &= \frac{1}{2} \left(1 - 1 + 2 \left(\frac{1}{2}\right)^n \right) \\
 &= \left(\frac{1}{2}\right)^n.
 \end{aligned}$$

So we have verified $P(X = n) = (\frac{1}{2})^n$. Then the PMF of X is

$$f_X(x) = \left(\frac{1}{2}\right)^x.$$

(b) *How much money have you won if $X = T$?*

Solution: consider this: if you win at time $T = 1$ on a bet of \$2, you have no losses, so your profit is \$2. If you win at time $T = 2$, you have lost a bet of \$2 but won a bet of \$4, so your profit is \$2. More generally, since 2^t is your bet at time t , if you finally win at time T your profit is

$$\pi_T = 2^T - \sum_{t=1}^{T-1} 2^t = 2^T - 2 \sum_{t=0}^{T-2} 2^t = 2^T - 2(2^{T-1} - 1) = 2.$$

(c) *What is your expected profit?*

Solution: expected profit can be computed as

$$\mathbb{E}[\pi_t] = \mathbb{E}[2] = 2.$$

One unrealistic feature of this setup is that you can continue doubling your bet if you lose. Imagine you have lost 10 flips in a row: you will have to place an extremely large bet on the 11th flip if you hope to recoup your lost earlier investments.¹ So now, let's assume that you start gambling with only \$100, and you cannot gamble away money you don't have (so if you have m left on your bankroll and your next bet would be $b > m$, you have to just walk away).

(d) *If you start with \$100, how many bets can you lose before you have to quit using the Martingale strategy?*

Solution: as in part (c), if you lose at time T you have lost

$$\sum_{t=1}^T 2^t = 2 \sum_{t=0}^{T-1} 2^t = 2(2^T - 1) = 2^{T+1} - 2.$$

¹Depending on what you consider a “large” bet. On a graduate student salary, this bet is a dealbreaker.

If you are budget-constrained, you cannot lose more money than you enter with; therefore if you enter with \$100, you must exit before $2^{T+1} - 2 > 100$. Solving through, we find

$$2^{T+1} - 2 \leq 100 \iff 2^T \leq 51 \iff 2 \leq 5.6724.$$

Thus you can lose 5 consecutive bets before you must stop following the Martingale strategy.

(e) *Let Y be a random variable representing the time at which you finally win or walk away when your budget is constrained. What is the PMF of Y ?*

Solution: let $Y = 0$ represent the outcome that you never win, and are forced to walk away from the table. When you don't "break the bank," your outcomes exactly map on to those of the random variable X from part (a); therefore, the probability that you win on the first toss is $\frac{1}{2}$, the probability you win on the second toss is $\frac{1}{4}$, etc. The probability that you never win is then the probability that none of the bets win, prior to you being forced to leave the table. Since you may lose at most 5 consecutive bets, the probability is

$$P(Y = 0) = 1 - (P(Y = 1) + P(Y = 2) + P(Y = 3) + P(Y = 4) + P(Y = 5)) = 1 - \left(\frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \frac{1}{16} + \frac{1}{32} \right) = \frac{1}{32}.$$

The PMF of Y is then

Y	f_Y	Y	f_Y
0	$\frac{1}{32}$	3	$\frac{1}{8}$
1	$\frac{1}{2}$	4	$\frac{1}{16}$
2	$\frac{1}{4}$	5	$\frac{1}{32}$

(f) *What is your expected profit?*

Solution: as in part (c), at all times prior to breaking the bank your profit is \$2. If you are forced to stop gambling due to your budget constraint, you will have lost

$$\sum_{t=1}^5 2^t = 62.$$

Your expected profit is therefore

$$\mathbb{E}[\pi_t] = \sum_{t \in \text{Support}(Y)} \pi_t f_Y(t) = \frac{1}{32}(-62) + \frac{31}{32}(2) = 0.$$

(g) *Thought experiment: what assumption have we made here that, while making the math friendly, is extremely unrealistic? If we change this assumption to make it better-match the real world, what happens to your expected profit in the previous question? Is the Martingale strategy smart? Bonus: can the strategy be adjusted to account for this (hint: think about what you could do if you made the math more difficult than, "double my previous bet.")?*

Investments

Here's a simple model of the stock market which actually follows some models that economists use.² There are two stocks: one can increase by 20% per year with probability $\frac{1}{2}$ or decrease by 10% per year with probability $\frac{1}{2}$; the other can increase by 60% per year with probability $\frac{1}{3}$ or decrease by 20% per year with probability $\frac{2}{3}$. You are concerned with investing over a period of two years; asset performance is independent between the first year and the second.

²The formal name for this, if we consider a slightly broader class of investments, is the *binomial options pricing model*.

(a) For each stock, what are the possible outcomes (increases or decreases in value) over two years?

Solution: let U represent the stock going up, and D represent the stock going down. The possible outcomes for a particular stock are $\{UU, UD, DU, DD\}$. For the first stock, the rates of return represented by these outcomes are

$$\begin{aligned} UU : \quad (1.2)(1.2) &= \frac{36}{25} \\ UD : \quad (1.2)(0.9) &= \frac{54}{50} = \frac{27}{25} \\ DU : \quad (0.9)(1.2) &= \frac{27}{25} \\ DD : \quad (0.9)(0.9) &= \frac{81}{100}. \end{aligned}$$

For the second stock, the rates of return represented by these outcomes are

$$\begin{aligned} UU : \quad (1.6)(1.6) &= \frac{64}{25} \\ UD : \quad (1.6)(0.8) &= \frac{32}{25} \\ DU : \quad (0.8)(1.6) &= \frac{32}{25} \\ DD : \quad (0.8)(0.8) &= \frac{16}{25}. \end{aligned}$$

(b) Let X_i denote the relative return of stock i over two years. What is the PMF of X_1 ? X_2 ?

Solution: this is a matter of translating the above outcomes into probabilities. We know

$$\text{Support}(X_1) = \left\{ \frac{36}{25}, \frac{27}{25}, \frac{81}{100} \right\}, \quad \text{Support}(X_2) = \left\{ \frac{64}{25}, \frac{32}{25}, \frac{16}{25} \right\}.$$

Then for the first stock we can see

$$\begin{aligned} f_{X_1} \left(\frac{36}{25} \right) &= P(U_1 U_1) = P(U_1)P(U_1) = \frac{1}{4}, \\ f_{X_1} \left(\frac{27}{25} \right) &= P(U_1 D_1) + P(D_1 U_1) = P(U_1)P(D_1) + P(D_1)P(U_1) = \frac{1}{2}, \\ f_{X_1} \left(\frac{81}{100} \right) &= P(D_1 D_1) = P(D_1)P(D_1) = \frac{1}{4}. \end{aligned}$$

For the second stock, we have

$$\begin{aligned} f_{X_2} \left(\frac{64}{25} \right) &= P(U_2 U_2) = P(U_2)P(U_2) = \frac{1}{9}, \\ f_{X_2} \left(\frac{32}{25} \right) &= P(U_2 D_2) + P(D_2 U_2) = P(U_2)P(D_2) + P(D_2)P(U_2) = \frac{4}{9}, \\ f_{X_2} \left(\frac{16}{25} \right) &= P(D_2 D_2) = P(D_2)P(D_2) = \frac{4}{9}. \end{aligned}$$

(c) What is $\mathbb{E}[X_1]$? $\mathbb{E}[X_2]$?

Solution: this is an algebraic exercise:

$$\mathbb{E}[X_1] = \sum_{x \in \text{Support}(X_1)} x f_{X_1}(x) = \frac{1}{4} \left(\frac{36}{25} \right) + \frac{1}{2} \left(\frac{27}{25} \right) + \frac{1}{4} \left(\frac{81}{100} \right) = \frac{441}{400} = 1.1025;$$

$$\mathbb{E}[X_2] = \sum_{x \in \text{Support}(X_2)} x f_{X_2}(x) = \frac{1}{9} \left(\frac{64}{25} \right) + \frac{4}{9} \left(\frac{32}{25} \right) + \frac{4}{9} \left(\frac{16}{25} \right) = \frac{256}{225} \approx 1.1378.$$

(d) *Based solely on expected return, which stock would you rather purchase?*

Solution: if you're in this to make money, it would be wise to choose the stock with the higher expected return; in this case, this is the second stock.

(e) *What is $\text{Var}(X_1)$? $\text{Var}(X_2)$?*

Solution: we may compute

$$\begin{aligned} \text{Var}(X_1) &= \sum_{x \in \text{Support}(X_1)} (x - \mathbb{E}[X_1])^2 f_{X_1}(x) \\ &= \sum_{x \in \text{Support}(X_1)} x^2 f_{X_1}(x) - \mathbb{E}[X_1]^2 \\ &= \frac{1}{4} \left(\frac{36}{25} \right)^2 + \frac{1}{2} \left(\frac{27}{25} \right)^2 + \frac{1}{4} \left(\frac{81}{100} \right)^2 - \left(\frac{441}{400} \right)^2 \end{aligned}$$

$$\text{Var}(X_1) \approx 0.05012,$$

$$\begin{aligned} \text{Var}(X_2) &= \sum_{x \in \text{Support}(X_2)} (x - \mathbb{E}[X_2])^2 f_{X_2}(x) \\ &= \sum_{x \in \text{Support}(X_2)} x^2 f_{X_2}(x) - \mathbb{E}[X_2]^2 \\ &= \frac{1}{9} \left(\frac{64}{25} \right)^2 + \frac{4}{9} \left(\frac{32}{25} \right)^2 + \frac{4}{9} \left(\frac{16}{25} \right)^2 - \left(\frac{256}{225} \right)^2 \end{aligned}$$

$$\text{Var}(X_2) \approx 0.3439.$$

(f) *Does knowing the variance change your decision regarding which stock you would rather purchase? Why or why not?*

Solution: this is a matter of opinion. While the second stock offers slightly higher returns, its variance is significantly larger; we refer to it as have *greater risk*. Depending on your attitudes towards loss, you may wish to purchase the first stock with a lower expected return, but a much better “worst case.”

Class quizzes

The week-3 quiz has 10 questions. A student has a 90% chance of getting any one question right, independent of all others. Let X represent the number of correct answers on the quiz (the grade).

(a) *What is the distribution of X ?*

Solution: we can regard each individual question on the quiz as a Bernoulli experiment: with probability 0.9 the student gets the answer correct (success), and with probability 0.1 the student gets the answer incorrect (failure). The grade of the quiz, X , is then the sum of 10 Bernoulli experiments; this defines a binomial distribution. X is then distributed $\text{binomial}(10, 0.9)$.

(b) *What is $P(\text{all answers are correct})$? $P(\text{only one answer is wrong})$?*

Solution: first, we must convert these concepts to outcomes of the random variable X . When all answers are correct, the score on the exam is $X = 10$; when only one answer is wrong, the score on

the exam is $X = 9$. Determining these probabilities is then a matter of evaluating the PMF at the particular X values. Since the binomial is a fairly common distribution, you can almost obtain these probabilities from the tables in the back of the text. However, we can just as well plug into the PMF ourselves to obtain the answer.

$$f_X(10) = \binom{10}{10} 0.9^{10} 0.1^0 = 0.9^{10},$$

$$f_X(9) = \binom{10}{9} 0.9^9 0.1^1 = 0.9^9.$$

Then we see that $P(\text{all answers are correct}) = 0.9^{10}$ and $P(\text{only one answer is wrong}) = 0.9^9$.

Now, what if we had tried to look these values up in the table in the back of the book? We run into trouble: there are no tables for the binomial distribution with $p > 0.5$! The book does this to save space, as it turns out the extra tables would be redundant. Notice that instead of looking at getting a question correct with probability 0.9, we could be looking at getting a question wrong with probability 0.1.

Let Y be a random variable denoting the total number of incorrect answers on the quiz; from the discussion above, we should have $Y \sim \text{binomial}(10, 0.1)$. Since each of the 10 questions on the quiz must be either correct or incorrect, we should have $X + Y = 10$. Then we know

$$f_X(10) = f_Y(10 - 10) = f_Y(0) \approx 0.3487, \quad f_X(9) = f_Y(10 - 9) = f_Y(1) \approx 0.7361 - 0.3487 = 0.3874.$$

These are exactly the probabilities we would find above by direct computation (although expressed as decimals rather than exponents).

(c) *Is {all answers are correct} independent of {only one answer is wrong}?*

Solution: this addresses the linguistic implication of “independence,” which differs in an important way from the notion of statistical independence. Statistical independence says that knowing information about one event does not give us any information on the probability of another event. In particular, the definition of independence of events A and B is

$$P(A \cap B) = P(A)P(B) \iff \frac{P(A \cap B)}{P(B)} = P(A) \iff P(A|B) = P(A).$$

That is, knowing that B has happened does not change our assessment of the probability that A will happen.

In the case of this question, knowing that the student “has only missed one question” informs us perfectly that he has not gotten “all answers correct.” We have

$$P(\text{all answers are correct} | \text{only one answer is wrong}) = 0 \neq 0.9^{10} = P(\text{all answers are correct});$$

it follows then that these events are not independent.

(d) *A friend tells you she missed **no more than one** question on the quiz. What is the probability that she achieved a perfect score?*

Solution: we are asked to assess

$$P(\text{all answers are correct} | \text{no more than one answer is wrong}).$$

In terms of our random variable X , this is

$$P(X = 10 | X \geq 9).$$

By the usual expansion of conditional probability, we have

$$P(X = 10|X \geq 9) = \frac{P(X = 10, X \geq 9)}{P(X \geq 9)}.$$

Consider the numerator: when $X = 10$ we know that $X \geq 9$, so requiring both is equivalent to only requiring $X = 10$; then $P(X = 10, X \geq 9) = P(X = 10)$. Now consider the denominator: when $X \geq 9$, we may have either $X = 9$ or $X = 10$. Since $X = 9$ and $X = 10$ are mutually exclusive³, we have

$$P(X \geq 9) = P(\{X = 9\} \cup \{X = 10\}) = P(X = 9) + P(X = 10).$$

We are now set to plug in for quantities we know. We obtain

$$P(\text{all answers are correct}|\text{no more than one answer is wrong}) = \frac{P(X = 10)}{P(X = 9) + P(X = 10)} = \frac{0.9^{10}}{0.9^9 + 0.9^{10}} = \frac{1}{1.9} \approx 0.4737.$$

(e) Suppose that the TAs messed up, and two of the questions are identical; missing one of these questions means missing both. If X is the grade on the quiz, what is the PMF of X ?

Solution: this question is nontrivial. Consider this: to receive a score of $X = 8$ on the quiz, the student could either get all the “okay” questions correct and miss the pair of questions which were messed up, or he could get two of the “okay” questions wrong and get the pair of messed-up questions correct. We need to find a way to disentangle these effects.

Let Y_1 be the score on the eight questions which are not messed up, and Y_2 be the score on the two questions which are. Since Y_1 represents the outcome of eight Bernoulli trials, each with probability of success 0.9, we know $Y_1 \sim \text{binomial}(8, 0.9)$. The support of Y_2 is $\{0, 2\}$: either the student gets both answers correct, or neither; with probability 0.9 he gets both answers correct, and with probability 0.1 he gets both wrong.

Now notice that $X = Y_1 + Y_2$; that is, the score on the overall quiz is the sum of the scores of these “subquizzes.” Consider the ways that X can take particular values:

X	X
0	$Y_1 = 0, Y_2 = 0$
1	$Y_1 = 1, Y_2 = 0$
2	$Y_1 = 2, Y_2 = 0$ or $Y_1 = 0, Y_2 = 2$
3	$Y_1 = 3, Y_2 = 0$ or $Y_1 = 1, Y_2 = 2$
4	$Y_1 = 4, Y_2 = 0$ or $Y_1 = 2, Y_2 = 2$
5	$Y_1 = 5, Y_2 = 0$ or $Y_1 = 3, Y_2 = 2$
6	$Y_1 = 6, Y_2 = 0$ or $Y_1 = 4, Y_2 = 2$
7	$Y_1 = 7, Y_2 = 0$ or $Y_1 = 5, Y_2 = 2$
8	$Y_1 = 8, Y_2 = 0$ or $Y_1 = 6, Y_2 = 2$
9	$Y_1 = 7, Y_2 = 2$
10	$Y_1 = 8, Y_2 = 2$

Then we know, for example, that

$$\begin{aligned} P(X = 4) &= P(Y_1 = 4, Y_2 = 0) + P(Y_1 = 2, Y_2 = 2) \\ &= P(Y_1 = 4)P(Y_2 = 0) + P(Y_2 = 2)P(Y_1 = 2) \\ &= 0.1f_{Y_1}(4) + 0.9f_{Y_1}(2). \end{aligned}$$

Applying this expansion generally, we find

x	$P(X = x) = f_X(x)$	x	$P(X = x) = f_X(x)$
0	$0.1f_{Y_1}(0)$	5	$0.1f_{Y_1}(5) + 0.9f_{Y_1}(3)$
1	$0.1f_{Y_1}(1)$	6	$0.1f_{Y_1}(6) + 0.9f_{Y_1}(4)$
2	$0.1f_{Y_1}(2) + 0.9f_{Y_1}(0)$	7	$0.1f_{Y_1}(7) + 0.9f_{Y_1}(5)$
3	$0.1f_{Y_1}(3) + 0.9f_{Y_1}(1)$	8	$0.1f_{Y_1}(8) + 0.9f_{Y_1}(6)$
4	$0.1f_{Y_1}(4) + 0.9f_{Y_1}(2)$	9	$0.9f_{Y_1}(7)$
		10	$0.9f_{Y_1}(8)$

³We can even ignore mutual exclusion:

$P(X \geq 9) = P(\{X = 9\} \cup \{X = 10\}) = P(X = 9) + P(X = 10) - P(\{X = 9\} \cap \{X = 10\}) = P(X = 9) + P(X = 10).$