ECON30580 Economics of Betting Markets
Repeated Gambles: The Binomial Theorem and Classic Problems

Karl Whelan

School of Economics, UCD

Spring 2025

Louis Bachelier on Gambling

It is almost always gambling that enables one to form a fairly clear idea of a manifestation of chance; it is gambling that gave birth to the calculus of probability; it is to gambling that this calculus owes its first faltering utterances and its most recent developments; it is gambling that allows us to conceive of this calculus in the most general way; it is, therefore, gambling that one must strive to understand, but one should understand it in a philosophic sense, free from all vulgar ideas."

- Louis Bachelier (founder of modern financial economics, 1914).

Repeated Gambles and Statistics

- We have presented lots of results on the economics of people placing once-off bets and bookmakers' strategies in accepting these bets.
- But most people don't just place one bet. Generally, they will place repeated bets over a long period.
- The story about people's thinking about what the outcomes will be from repeated gambles is a long and interesting one and, as Louis Bachelier noted, played a key role in the development of probability theory.
- In these lecture notes, we will discuss some ideas about repeated gambles incorporating a number of famous "fallacies".
- In the next set of notes, we will discuss what the famous Law of Large Numbers and Central Limit Theorem imply for repeated gambles.

Part I

The Gambler's Fallacy and Momentum

Repeated Gambles and the Gambler's Fallacy

- Suppose you place a bet that pays off when X happens, which you think has a probability p of happening.
- Let's say $p = 0.1$. Are you going to win or lose? You might win but you will probably lose.
- Now suppose you can repeat this bet over and over, with either X happening or not happening, but with separate events occurring independently from each other.
- You figure you should win this bet one time in 10.
- Suppose you lose the first 9 bets. Should that raise your hopes that you will win the 10th bet? You're due right? The averages suggest you should win.
- No. This is an example of "the gambler's fallacy" and it's a mistake lots of gambler's make: "I'm going to keep betting because I feel I'm due a win."
- The most famous example of the gambler's fallacy occurred in a game of roulette at the Monte Carlo Casino on August 18, 1913, when black came up 26 times in a row. Gamblers began to place large amounts of money of red, figuring the streak of blacks must end. In reality, the chances of black and red were always the same.

A Roulette Wheel

Uncorrelated Outcomes or Momentum?

- The Gambler's Fallacy comes from thinking independent events are negatively correlated: If I've tossed a lot of tails, I'm more likely to toss heads now.
- The mistake is thinking there is some correlation between events that are actually independent of each other.
- For example, you may know that your chance of winning a game is 20% but there is no magic force making 2 wins out of 10 more likely after we have played 9 times and won only once. The past events are irrelevant.
- An alternative possibility is that outcomes could be positively correlated.
- For example, what if I am betting on a specific team to win and they go on a streak of wins? Maybe this is just a random outcome—the team had a probability p of winning the games and these kinds of streaks just happen sometimes—or maybe the team has "**momentum**".
- Sports fans and bettors tend to believe more in momentum as a phenomenon than academics that study sports with statistical models.
- But figuring out whether there is momentum—or whether a team's chance of winning games is changing over time because they are getting bettor or worse—is an important part of the challenge facing bookmakers and bettors.

Part II

Gambling and the Binomial Distribution

Gambling and Probability

- Gambling may seem a frivolous thing to study.
- However, the study of gambling was central to the development of probability theory, which is used everywhere in the modern world.
- The earliest formulation of something that looked like probability theory was from Giralamo Cardano in his *Liber de Ludo Aleae (Book on Games of Chance)*, apparently written about 1525 but not published until after his death.
- Cardano was a compulsive gambler. He confessed to "*immoderate devotion to table games and dice ... During many years ... I played not off and on but, as I am ashamed to say, every day.*"
- Cardano's work had little practical influence but the principal promoters of modern probability, French mathematicians, Blaise Pascal and Pierre de Fermat were also inspired by questions related to games of chance.
- Many of the statistical methods used today, from the binomial distribution, to the Law of Large Numbers and the Central Limit Theorem had their origins in the exploration of gambling.

Repeated Gambles

- We have covered how to estimate probabilities from sports betting odds.
- But we have only discussed betting once on a single outcome.
- What if we repeatedly take bets with probability p of success? How often are we likely to win?

Question

- You are playing a game with a probability of success for each turn of $p = 0.2$.
- You are going to play 10 rounds of the game.
- What is the probability that you will win at least 2 rounds?
  1. Greater than 0.5.
  2. Exactly 0.5
  3. Less than 0.5

Repeated Gambles

- We have covered how to estimate probabilities from sports betting odds.
- But we have only covered single outcomes.
- What if we repeatedly take bets with probability p of success? How often are we likely to win?

Question

- You are playing a game with a probability of success for each turn of $p = 0.2$.
- You are going to play 10 rounds of the game.
- What is the probability that you will win at least 2 rounds?
  1. Greater than 0.5.
  2. Exactly 0.5
  3. Less than 0.5

- The answer is 0.62. You have almost a two-thirds chance of winning at least twice in 10 rounds.
- We will show how we got the answer but start with some easier stuff.

Two Rounds

- Labelling a win as W and a loss as L, there are 4 possible outcomes
  1. WW with probability p^2.
  2. WL with probability $p(1-p)$.
  3. LW with probability $(1-p)p$.
  4. WW with probability $(1-p)^2$.
- Check out the "tree" below which illustrates this.

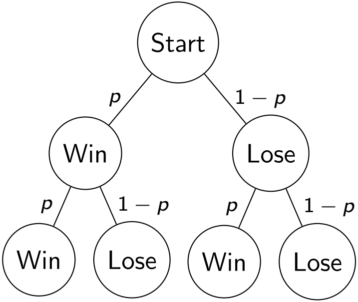

Three Rounds

- Now let's play 3 rounds of the game. What is the probability that you will win exactly twice?
- Describe the outcome where you win all 3 rounds as WWW. There are 8 possible outcomes
  1. WWW which happens with probability p^3
  2. WWL which happens with probability $p^2(1-p)$
  3. WLW which happens with probability $p^2(1-p)$
  4. WLL which happens with probability $p(1-p)^2$
  5. LWW which happens with probability $p^2(1-p)$
  6. LWL which happens with probability $p(1-p)^2$
  7. LLW which happens with probability $p(1-p)^2$
  8. LLL which happens with probability $(1-p)^3$
- There are 3 different ways you could win 2 times out of 3 (WWL, WLW, WLW) each of them occurring with probability $p^2(1-p)$.
- For $p = 0.2$, the probability of winning exactly twice is $3(0.2)^2 0.8 = 0.096$

Combinations: n-choose-k

- We just worked out that there were 3 different ways we could win 2 out of 3 rounds.
- To calculate the probabilities for these kind of problems, we need to know the general formula for winning k out of n rounds.
- The formula for that is usually written

$$\binom{n}{k} = \frac{n!}{k!(n-k)!}$$

  where $n! = (1)(2)(3)....(n-1)(n)$ (this is called "n factorial") and $0! = 1$.
- The term we use to describe $\binom{n}{k}$ is "n choose k"
- In our example, we were calculating

$$\binom{3}{2} = \frac{3!}{2!1!} = \frac{(1)(2)(3)}{2} = 3$$

- We now have the tools to solve the problem we posed before.

Probability of At Least 2 Wins in 10 Rounds

- Let P_k be the probability of winning exactly k times in 10 rounds. We now know that
$$P_k = \binom{10}{k} p^k (1-p)^{10-k}$$
- On the next page we can see the full probability distribution, showing P_k for each of the possible values of k.
- We can figure out the probability of winning at least twice by adding $P_2 + P_3 + ... + P_{10}$ or, somewhat easier, calculating one minus the probability of winning either one or zero
$$\begin{aligned} 1 - P_0 - P_1 &= 1 - \binom{10}{0}(0.8)^{10} - \binom{10}{1}(0.8)^9\, 0.2 \\ &= 1 - (0.8)^{10} - 10\,(0.8)^9\, 0.2 \\ &= 0.6242 \end{aligned}$$
- This set of probabilities—the binomial distribution—was first discovered by Jacob Bernoulli, published posthumously in his book *Ars Conjectandi* (1713).

Probabilities of Winning k Times in 10 Rounds

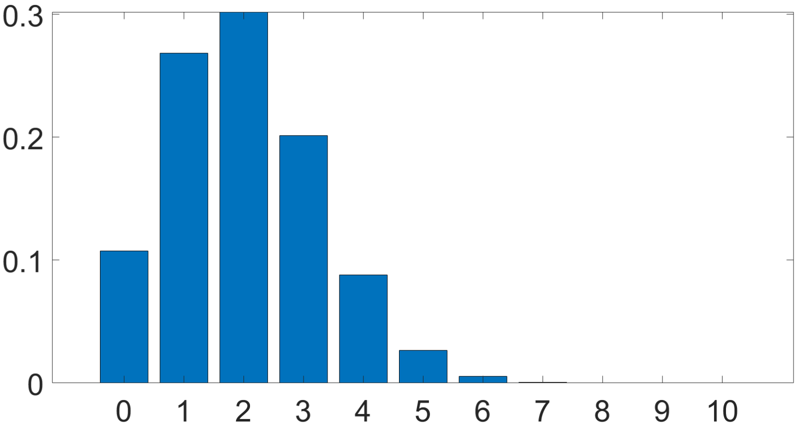

On Binomial Distribution Outcomes

- When I said we were playing 10 rounds of a game with a probability of success for each turn of $p = 0.2$ what did you think would be the likely outcome?
- You might have thought "I think there will be 2 wins"
- And indeed, 2 is the mean of the distribution just shown. Putting it technically

$$\sum_{k=0}^{10} kP_k = 2$$

- 2 is also the most likely outcome, i.e. the mode of the distribution.
- But $P_2 = 0.3$, so in fact most of the time there will be either more or less than 2 wins.
- Given there is a chance $P_0 = 0.11$ that you don't win at all, for the mean to be 2 you are going to have win at least 2 more than half the time, hence our 62% figure.
- Note also, for those of you used to looking at the nice bell-shaped Normal probability distributions, the distributions of outcomes in the binomial distribution can be quite skewed.

Part III

The Martingale Betting Strategy

The “Can’t Lose” Strategy?

- For centuries, people have been hyping the following “can’t lose” betting strategy, known as the **Martingale strategy**, supposedly named after an 18th century casino owner who encouraged his customers to use it.
- There is a double-or-nothing gamble where I either win or lose 1.
- There is a probability p that I lose the gamble and a probability $1 - p$ that I win.
- Consider the following strategy:
  - I start by staking 1. If I win, I stop and have won 1.
  - If I lose the first bet, I stake 2. If I win, I stop and have won 1.
  - If I also lose the second bet, I stake 4. If I win, I stop and have won 1.
  - ...
  - If I also lose the Nth bet, I stake 2^N. If I win, I stop and have won 1.
  - And so on
- Stop when you win and you are guaranteed to come out ahead no matter what the value of p is, right?
- No, not right.

The Problem? Finite Wealth

- You only have a finite amount of wealth and this restricts the application of the strategy. Even if you had infinite wealth, you are not going to find some who will keep offering you double-or-nothing bets at ever-higher stakes.
- Suppose your initial wealth is $W_0 = 2^n - 1$.
- If you play n rounds in a row and lose, you will lose precisely this amount (note $1 + 2 = 3 = 2^2$, $1 + 2 + 4 = 7 = 2^3$ and so on).
- So there is a probability p^n that you will lose all your wealth and a probability $1 - p^n$ that you will win one.
- Your expected gain is

$$E\left(\sum_{i=1}^{n} X_i\right) = 1 - p^n - p^n\left(2^n - 1\right) = 1 - (2p)^n$$

- If $p > 0.5$ so you are more likely to lose, then this is negative.
- See the next page for a roulette-inspired example. The probability of red coming up is $1 - p = \frac{18}{37}$. You are more likely to win than lose but when you win, you don't win much, and when you lose, you lose big.

The Martingale Strategy with $p = \frac{19}{37} = 0.5135$ and $W_0 = 2^{10} - 1 = 1023$

Round	Outcome	Probability	Bet	Cumulative Bet	Gain	Probability*Gain
1	R	0.4865	1	1	1	0.4865
2	BR	0.2498	2	3	1	0.2498
3	BBR	0.1283	4	7	1	0.1283
4	BBBR	0.0659	8	15	1	0.0659
5	BBBBR	0.0338	16	31	1	0.0338
6	BBBBBR	0.0174	32	63	1	0.0174
7	BBBBBBR	0.0089	64	127	1	0.0089
8	BBBBBBBR	0.0046	128	255	1	0.0046
9	BBBBBBBBR	0.0024	256	511	1	0.0024
10	BBBBBBBBBR	0.0012	512	1023	1	0.0012
10	BBBBBBBBBB	0.0013	512	1023	-1023	-1.3044
				Expected Gain		-0.3056

Part IV

The Gambler's Ruin

Blaise Pascal (1623-1662) and Pierre de Fermat (1607-1665)

(a) Blaise Pascal

(b) Pierre de Fermat

Pascal, Fermat and a Gambling Problem

- The key figures in inventing modern probability theory were French mathematicians, Blaise Pascal and Pierre de Fermat.
- Like Cardano, they were also inspired by questions related to games of chance.
- One problem that relates to our interest was posed in a letter from Pascal to Fermat.
- It was roughly as follows: "*Two gamblers are playing repeated games of chance against each other in which they can either win one unit or lose one unit with equal probability. Gambler 1 has n units available for gambling and Gambler 2 has m units. The game continues until one of the gamblers has zero and the other has $n+m$. What is the probability that each Gambler wins?*"
- Both solved the problem, though using different methods.
- Here I will provide a simple way to solve this problem and then discuss a more general version of it.

Random Walks and Martingales

- Consider a gambler that starts with wealth of W_0, gambles \$1 per game obtaining profit X_i on the i-th round with expected value of zero $E(X_i) = 0$.
- Their wealth at time t can be written as

$$W_t = W_{t-1} + X_t$$

- This kind of process is known as a **random walk**: It takes a new step each period but the step is unpredictable and, on average, you expect next period's value to equal this period's value.
- Random walks have featured heavily in financial economics since Louis Bachelier's *The Theory of Speculation* in 1900 put forward the hypothesis that they could be considered a good model for stock prices. The term was popularized in a best-selling 1973 book by Princeton professor Burton Malkiel called *A Random Walk Down Wall Street*.
- Random walks are an example of a wider class of processes called **martingales**. A series X_t is a martingale if

$$E(X_{t+1} \mid X_1, X_2, ..., X_t) = X_t$$

The Optional Stopping Theorem

- Random walks, if they last long enough, can end up taking any possible value.
- Indeed, if let them run for a long enough time, they will eventually reach every possible value.
- But it is unlikely, however, that these extreme values will work in our gambling applications. There will always be a limit to how much anyone would be able to borrow to fund gambling losses and if someone is highly successful, they will start to spend it rather than let the wealth accumulate to ever higher levels.
- Let's assume there is a **stopping rule** so that the random walk has to stop when it reaches a specific high value or a specific low value.
- A result in statistics known as **the Optional Stopping Theorem** states that for any martingale process X_t and any stopping rule that keeps the process from going to plus or minus infinity, the expected value of the process when stopped equals the current value.
- Let's see what this means for our gamblers.

The Probability of Reaching Your Target

- Again consider the gambler playing a game where they can either win or lose one unit with equal probability

$$X_i = \begin{cases} 1 & \text{with probability } 0.5 \\ -1 & \text{with probability } 0.5 \end{cases}$$

- They start with a gambling fund of $W_0 = n$ and have a target of getting this up to T.
- This might be because they are playing against someone with a gambling fund of $m = T - n$ or because they are playing against someone with "deep pockets" (e.g. a bookmaker) who will let them pick their own target.
- Let P_n represent the probability that someone with starting wealth of n reaches the target T.
- The optional stopping theorem tells us that

$$P_n T + (1 - P_n)(0) = n \Rightarrow P_n = \frac{n}{T}$$

The Gambler's Ruin

- We know now that the probability of the gambler with wealth of n reaching their target T is
$$P_n = \frac{n}{T}$$
- Intuition for this result: If n is close to T, then the random walk doesn't have very far to go to reach the target. But if n is closer to zero than to T, then the random walk is more likely to reach zero before it reaches T.
- The bigger the target, T, the more likely the gambler is to lose all their money.
- As T goes to infinity (because you have an ambitious target and are playing against someone with lots of wealth), your probability of losing all your money goes to one. This is the result typically known as the Gambler's Ruin.
- Note you are not losing all your money because you are at a disadvantage in the game—it is a fair game with equal chances of winning or losing. You lose all your money because eventually, with probability one, you will go on a bad enough losing streak to drive your gambling fund to zero.

Non-Fair Games (For Completeness – Not on the Exam)

- What if the probability of winning is $p \neq 0.5$? In this case,

$$X_i = \begin{cases} 1 & \text{with probability} \quad p \\ -1 & \text{with probability} \quad 1-p \end{cases}$$

- The probability of reaching T starting with wealth of n satisfies the following so-called difference equation

$$P_n = pP_{n+1} + (1-p)\,P_{n-1}$$

- I won't go through how to solve these equations but the solution is

$$P_n = \frac{\left(\frac{1-p}{p}\right)^n - 1}{\left(\frac{1-p}{p}\right)^T - 1}$$

The Gambler has an Advantage (Not on the Exam)

- The probability of reaching T starting from n is

$$P_n = \frac{\left(\frac{1-p}{p}\right)^n - 1}{\left(\frac{1-p}{p}\right)^T - 1}$$

- If $p > 0.5$ then

$$\frac{1-p}{p} < 1$$

- As T gets bigger,

$$P_n \to 1 - \left(\frac{1-p}{p}\right)^n$$

- For example, if $n = 100$ and $p = 0.505$ (implying an expected profit on a one unit stake of 0.01) then for large T we have

$$P_n = 1 - \left(\frac{0.495}{0.505}\right)^{100} = 0.865.$$

- Even with an advantage, there is still a 13.5% chance of being ruined when staking 1% of your initial wealth each time, independent of the target wealth.