Basics of probability

In this section, we introduce the basic language of probability. We use the framework of **experiments** and **outcomes** to give a general overview of probability distributions.

An **experiment** is an activity with an observable outcome. An example might be flipping a coin or rolling a six-sided die or taking a penalty shot in soccer.

The **Sample space** for an experiment is the list of all possible outcomes. We assume that no two of the outcomes on the list can happen at the same time.

In our above examples the sample spaces are

Experiment	**Sample Space**
Flip a fair coin	{Heads, Tails }
Roll a six-sided die	{1, 2, 3, 4, 5, 6 }
Take a penalty shot	{ Goal, No Goal }

We assign probabilities to each outcome in the sample space according to the following rules

1. Each probability assigned is a number between 0 and 1
2. A probability of 0 means an impossible outcome and a probability of 1 means certainty.
3. The sum of the probabilities of all of the outcomes in the sample space is 1.
4. The probabilities assigned should reflect the relative frequency of the outcome in reality.

A **Probability distribution** for an experiment is a table showing the outcomes in the sample space alongside their probabilities.

To assign probabilities for experiments such as rolling a die or flipping a coin, we can use logic to determine that the outcomes are equally likely. If there are n equally likely outcomes for an experiment, then each should get a probability of $\frac{1}{n}$, since the sum of the probabilities of the outcomes should be 1.

Flip a fair coin

Outcomes	**Probability**
Heads	$\frac{1}{2}$
Tails	$\frac{1}{2}$

Roll a die

Outcomes	**Probability**
1	$\frac{1}{6}$
2	$\frac{1}{6}$
3	$\frac{1}{6}$
4	$\frac{1}{6}$
5	$\frac{1}{6}$
6	$\frac{1}{6}$

Events Sometimes, it is necessary to calculate the probability of an **event** which is more complex than a single outcome. In this case, the event corresponds to a subset of the outcomes in the sample space and its probability is the sum of the probabilities of the outcomes in the event.

Example Suppose our experiment is to roll a fair six-sided die and we wish to calculate the probability of the event E, that the result is an even number. First we identify the subset of outcomes which are in this event (make this event happen);

$$E = \{2, 4, 6\}.$$

Then we add the probabilities of those outcomes;

$$P(E) = 1/6 + 1/6 + 1/6.$$

Chance as a Player Many games involve randomness. That may involve rolling a die or flipping a coin or being successful in shooting for a goal in soccer or a basket in basketball. It is thus often convenient to represent CHANCE as a player in our game so that we can incorporate this randomness into our game tree. We can incorporate CHANCE in our game tree by considering a single player playing against CHANCE or introduce a third player called CHANCE to our two player games. When incorporating chance in this way, in addition to the rules for constructing game trees given above we label each branch corresponding to a choice made by CHANCE by the probability that CHANCE will make that choice.

Example Simplified Poker; See Straffin. In a game of simplified poker, two players Rose and Colin, each player puts one dollar into the pot as "ante". Each is then dealt a hand, which consists of one card, from a large well shuffled deck which consists only of aces and kings (with equal numbers of both). Colin must then decide whether to *bet* \$2 or to *drop*. If he drops, Rose wins the pot. If Colin bets, Rose must decide whether to *call* by matching Colin's bet or to *fold*. If she folds, Colin wins the pot. If Rose calls, the players compare their hands and the higher card wins the pot, if the hands tie, the pot is split equally.

To model this game, the first choice is made by CHANCE in dealing the cards. The four possibilities are ace to both (A, A), ace to Rose and king to Colin (A, K), king to Rose and ace to Colin (K, A), and king to both (K, K). Since we are assuming that the cards are dealt randomly from a large pack, all four possibilities are equally likely and we assign equal probabilities of $\frac{1}{4}$ to each. In the game tree below, the players are denoted by Ch. (CHANCE), C (Colin) and R (Rose)

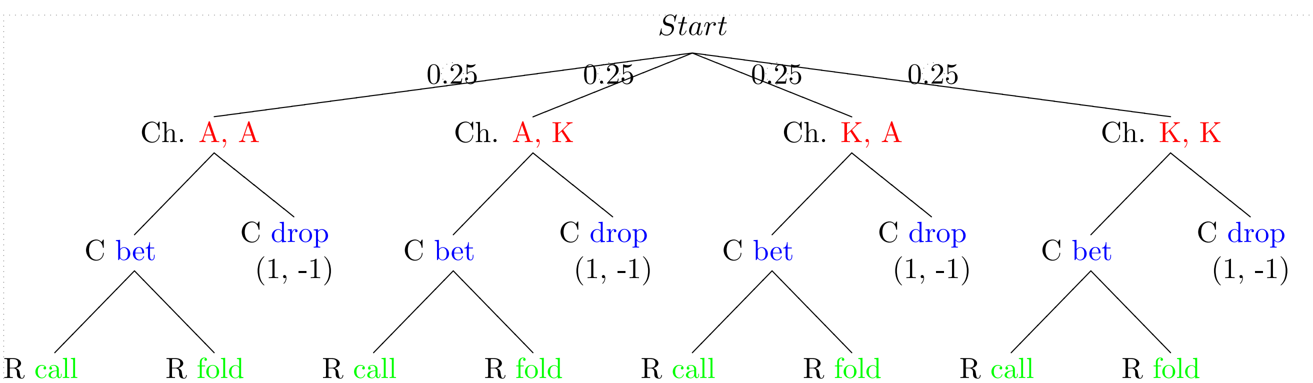

Note the pay-offs are given for the pathways ending at "C drop". Here (1, -1) means that Rose wins \$1 and Colin loses \$1.

Exercise (a) Fill in the pay-offs for the paths ending with a move by Rose on the diagram above.

(b)If Colin always bets and Rose always calls, what is the probability that Rose will win \$3 next time they play?

Using Relative Frequency to estimate chance

To assign probabilities in experiments such as taking a penalty shot, we can use statistical data to derive relative frequencies.
We let $P(O)$ denote the probability that the outcome O will happen. We can estimate this using the **relative frequency of the outcome**, O,

$$P(O) \approx \text{Relative Frequency of Outcome O} = \frac{\#\text{ Times we observe outcome O}}{\#\text{ Times we try the experiment}}$$

The **Law of large numbers** says that the relative frequency should get closer to the true probability as the number of trials of the experiment we use in the calculation gets very large.

Therefore in assigning probabilities using relative frequency, we may be able to achieve more accuracy by either

- Taking a larger number of experiments into account or
- Making sure the conditions of the experiments used in the calculation match the given conditions.

The latter point will be discussed under conditional probability.

Example Suppose David is taking a penalty shot against Raul. If we look at David's past history, we might see that of all of the 50 penalty shots that he has taken in official games, 25 resulted in goals. This would give us a relative frequency of $25/50 = 1/2$ and an estimate of the probability of a goal in the upcoming penalty shot at $1/2$. The probability distribution for the upcoming penalty shot would look like

Penalty General

Outcomes	Probability
Goal	$\frac{1}{2}$
No Goal	$\frac{1}{2}$

Penalty Against Raul

Outcomes	Probability
Goal	$\frac{1}{5}$
No Goal	$\frac{4}{5}$

On the other hand, we may have a little more information on our hands, 5 of the penalty shots that David might have taken in official games may have been against Raul and only one of those shots may have been successful. This might indicate that David does not have a fifty-fifty chance of scoring a goal against Raul. On the other hand, the sample size is rather small and the underperformance against Raul may have been due to random variation as opposed to Raul's expertise. There are statistical methods to measure variation and make recommendations about minimum sample sizes that should be used to calculate relative frequencies.

Number of Trials With more sophisticated statistical methods, we can show that our approximation of probability using relative frequency works best under the condition given below. Let $\hat{p}$ denote the estimate of probability that we derive for an outcome, O, using n trials of an experiment. This approximation works best when n is large enough so the the following inequalities hold true:

$$0 \leq \hat{p} - 3\sqrt{\frac{\hat{p}(1-\hat{p})}{n}} < \hat{p} + 3\sqrt{\frac{\hat{p}(1-\hat{p})}{n}} < 1.$$

We can also show that for such an n and $\hat{p}$ there is roughly a 95% chance that the true probability of the outcome lies between $\hat{p} - 2\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$ and $\hat{p} + 2\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$.

Example (a) If $n = 5$ and $\hat{p} = 1/5$ as with the goal shots that David took against Raul, do the above inequalities hold?

(b) If $n = 50$ and $\hat{p} = 1/2$ as with the overall success rate for David on Goal shots, do the above inequalities hold?

Although the sample size is too small to make a relatively sound estimate of the probability that David will score a goal against Raul, we may want to use a more conservative estimate than 1/2 for the probability that David will score a goal against Raul on his next penalty.

CHANCE as the second player

We can also consider CHANCE as the second player in a one-person game or a game of solitaire as shown in the example below.

Example Consider a mountain climber who wishes to climb a mountain with two main paths to the top. The first path(Path 1) takes longer to ascend but the attempt to reach the summit will not have to be abandoned in the event of a rainstorm. The second path (Path 2) is steeper and has a scree slope near the summit. Any attempt to reach the summit along the second path will have to be abandoned in the event of a rainstorm. A rainstorm on Path 1 will not have any effect on the climber's progress. Any attempt to reach the summit must be abandoned after 4 hours of climbing, since the climber must leave enough time to descend before darkness falls. The climber may experience delays on path 1, due to delays at a bottleneck on the route (causing the climber to fail to reach the summit about 30% of the time) because of the popularity of the route among inexperienced climbers. On the other hand, if there is no rainstorm, the climber is more or less certain to reach the summit on time on path 2. Complete the game tree for this game shown below with CHANCE as the second player assuming that there is a 50% chance of a rainstorm the day of the climb.

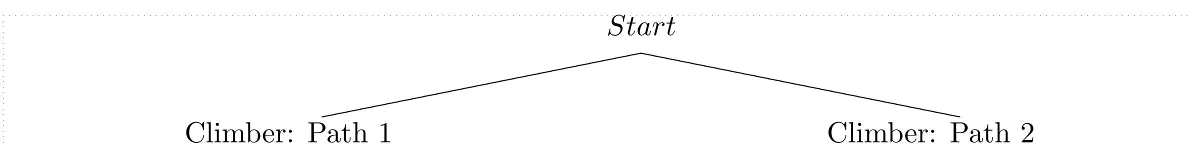

Which of the strategies "Choose Path 1" or "Choose Path 2" will maximize the climber's chances of reaching the summit?

Conditional Probability

We saw above that probabilities and relative frequencies often depend on underlying conditions. For example, the chances of success in a penalty shot may depend on the who the shooter is or who the goalie is or both. When we take into account these special given conditions, we call our probability **conditional probability** and indicate which conditions we have taken into account. We denote the conditional probability of an outcome, O, given the underlying condition C, by the symbol

$P(O|C)$, to be read as "The probability that O will happen given the underlying condition C".

In general to calculate a relative frequency estimate for the probability that O will occur given the underlying condition C, we use

$$P(O|C) \approx \frac{\#\text{ Times we observe outcome O}}{\#\text{ Times we try the experiment with underlying condition C in effect.}}$$

Example Suppose that David is about to take a penalty shot against Peter. We have statistics on David's record available and 15 of the 50 penalties he took in official games were against Peter. The table below shows the success rate of David's penalty shots against Peter and against others.

	Goal	**No Goal**
vs. Peter	5	10
vs. Other	20	15

Here, as noted above the overall probability of a goal for David is estimated by the relative frequency of goals scored from all of his penalty kicks:

$$P(G) = \frac{25}{50} = \frac{1}{2}.$$

On the other hand, the probability that David will score a goal given that he is facing Peter is $P(\text{Goal}|\text{Facing Peter})$ and is best approximated by the relative frequency of goals scored when facing Peter.

$$P(G|P) \approx \frac{10}{15} = \frac{2}{3}.$$

Note that in this case, the number of shots against Peter is large enough to ensure that the statistical inequalities given above to hold.

Sometimes prior information about an underlying condition, C, changes the probability of an outcome or event , O, in other words sometimes $P(O|C) \neq P(O)$. The condition, C, can be regarded as an event that has happened or will happen when we are performing our experiment. If the probability is not changed by the underlying condition, we say the events are **independent**.

Example In the example given above where David is about to take a penalty shot against Peter. Based on our estimates of probabilities using relative frequencies, we see that the events "David will score a goal" and "Peter is the goalie" are not independent, since our estimates for $P(G)$ and $PG|P)$ are not equal.

Probability of two events happening simultaneously or sequentially If I have two events or outcomes A and B, then the probability that both will happen is denoted by $P(A \cap B)$. The symbol $\cap$ here corresponds to intersection and is read as "and". For two events, A and B, we can calculate $P(A \cap B)$ using the following formula:

$$P(A \cap B) = P(A)P(B|A).$$

If the events A and B are **independent**, then we can multiply their probabilities to get $P(A \cap B)$.

$$P(A \cap B) = P(A)P(B) \qquad \text{if } A \text{and } B \text{ are independent.}$$

Example Suppose you flip a fair coin and roll a fair six-sided die, what is the probability that you will get a head on the coin and a six on the die?

Example If Peter (the goalie) has saved 70% of the penalties taken against him, estimate the chances that he will save all five of the penalties taken against him in a penalty shoot out, assuming that he has the strength of mind not to let the outcome of one of the shots influence what happens on the others.

Example A basketball player gets a basket on 90% of his free throws. What are the chances that he will get all three baskets if awarded a such a free throw?
(assuming that his shots are independent).

Example You are planning on running a marathon in Dublin. The chances that it will rain during the race are 50%. The chances that you will win depend on whether it is raining or not. Your chance of winning in the rain are 40% and your chances of winning when it is not raining are 60%.

(a) Fill in the probabilities on the following tree diagram showing the possible outcomes (as paths).

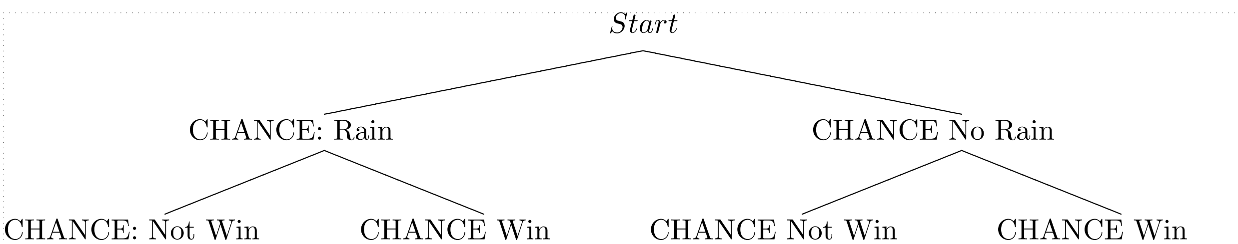

(b) Note that there are four possible outcomes here, R/W = it rains and you win, R/NW = it rains and you do not win, NR/W = it does not rain and you win, NR/NW= it does not rain and you do not win. Note that each outcome corresponds to a path on the tree diagram. Use the probabilities given above to find the probability of each outcome and derive the probability distribution.

Outcome	**probability**
R/W	
R/NW	
NR/W	
NR/NW	

(c) Now suppose there is a pair of shoes, the Rainikes, you could wear that would increase your chances of winning if it were raining. With these shoes your chances of winning if it rains increase to 60% but your chances of winning if does not rain decrease to 50%. Can you decide using a game tree which strategy "wear the Rainikes" or "do not wear the Rainikes" gives you the best chance of winning the race?

Using Expected Value to decide on a Strategy

Expected Value The expected value of a set of values $x_1, x_2, x_3, \ldots, x_n$ that occur with probabilities $p_1, p_2, p_3, \ldots, p_n$(which sum to 1) is

$$p_1x_1 + p_2x_2 + p_3x_3 + \cdots + p_nx_n.$$

If CHANCE is a player, we can use the expected value to calculate the average or expected payoff for a given scenario.

Example Simplified Poker; See Straffin. Recall the game of simplified poker;

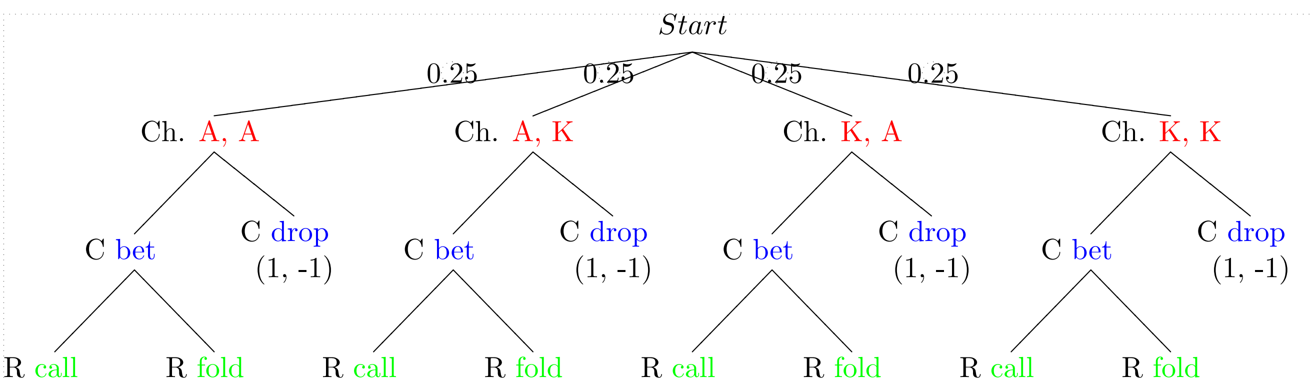

What is the expected value of Rose's payoff if Colin always bets and Rose always calls?

References

Mathematics beyond The Numbers: Gilbert, G, Hatcher, R.
Game Theory and Strategy, Phillip D. Straffin, The Mathematical Association of America, New Mathematical Library.
Mathletics, Wayne L. Winston, Princeton University Press.
Thinking Strategically, Avinash K. Dixit, Barry J. Nalebuff, W.W.Norton and Company.
Strategies and Games, Prajit K. Dutta, The MIT Press.
Sports Economics, R. D. Fort, Pearson.