Chapter 8

The Simple Random Walk

In this chapter we consider a classic and fundamental problem in random processes; the simple random walk in one dimension. Suppose a 'walker' chooses a starting point on a line (one-dimensional motion) and then decides to take one step right or left at random, dependent on the toss of a coin. If the coin is HEADS, they move right, if it comes up TAILS then the move is to the left. The second step proceeds in the same manner, and so on for each subsequent step. Then the position of the walker at any time is uncertain (random), although the location at future times will be determined by a probability distribution.

We contrast this with a *deterministic* walk, for example, in which a walker takes steps in a fixed direction (left or right) at a uniform pace. In such a case, the position of the walker at any future time is completely predictable.

The random walk is called 'simple' if each and every step has the same length. Consider an example, in which the length of the step is one unit, and the walker starts at the location $x = 5$. Suppose that we use a fair coin so that the probability of heads or tails is the same. Then the position of the walker, according to the number of steps taken (that is as time progresses) will trace out a path. An example of such a process is shown in figure 8.1. We note that such a graph of position versus time is equivalent to recording the sequence of outcomes.

8.1 Unrestricted simple random walks

Suppose the walk started at $x = 0$ and was allowed to continue along this infinite line unhindered: an *unrestricted random walk*. We now calculate the probability distribution for such a process, that is, determine an expression for the probability mass function for the position of the walker.

We have noted that the walk (the graph of position versus time) is completely equivalent to a sequence of Bernoulli trials. And a sequence of Bernoulli trials gives rise to a *binomial distribution*. Indeed a famous experiment illustrating this is given by a *Galton machine*. In the experiment

Each trial/toss can be mapped to a discrete random variable. Let X_i be the distance moved on the ith step. Then, for any i, we have the probability mass function:

$$P(X_i = +1) = p \qquad , \qquad P(X_i = -1) = q = 1 - p \qquad . \tag{8.1}$$

Each step is thus an *independent identically-distributed* Bernoulli variable, and we see that:

$$\mathbb{E}\left(X_i\right) = p - q \qquad , \qquad \operatorname{var}\left(X_i\right) = pq \qquad . \tag{8.2}$$

So the position of the walker after n steps we can denote as:

$$S_n = X_1 + X_2 + \cdots + X_n \qquad , \tag{8.3}$$

given that the walker starts at the origin, $S_0 = 0$. One can immediately make some assertions about the walker's position S_n. The 'best' guess (in the usual sense of the expected least-squares error) will be:

$$\mathbb{E}\left(S_n\right) = \mathbb{E}\left(X_1 + X_2 + \cdots + X_n\right) = \mathbb{E}\left(X_1\right) + \mathbb{E}\left(X_2\right) + \cdots + \mathbb{E}\left(X_n\right) \qquad . \tag{8.4}$$

Since each step/toss has the same (identical) probability mass, this can be written as, using equations 8.2:

$$\boxed{\mathbb{E}\left(S_n\right) = n\mathbb{E}\left(X_1\right) = n(p-q)} \quad . \tag{8.5}$$

The uncertainty in this estimate is given by the mean square error (variance). Since the X_i are all mutually independent, it follows that:

$$\text{var}\left(S_n\right) = \text{var}\left(X_1 + X_2 + \cdots + X_n\right) = \text{var}\left(X_1\right) + \text{var}\left(X_2\right) + \cdots + \text{var}\left(X_n\right) \quad . \tag{8.6}$$

Thus:

$$\boxed{\text{var}\left(S_n\right) = n \ \text{var}\left(X_1\right) = npq} \quad . \tag{8.7}$$

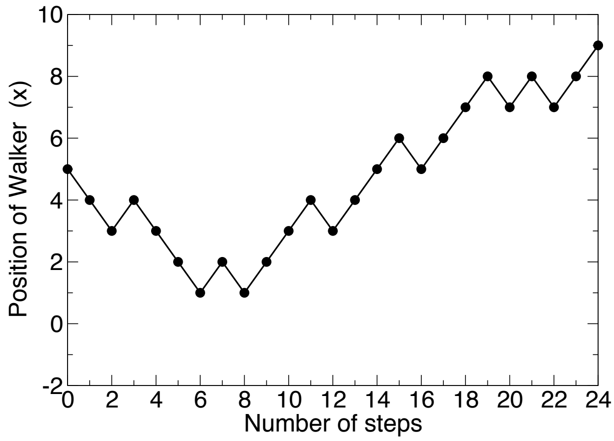

Figure 8.1: Example of a simple random walk in one dimension. In this example the walker begins at $x = 5$. This distance-time graph is equivalent to recording the sequence of the Bernoulli trials that determine each step. So an equivalent Bernoulli sequence would be: 001000101110111011101011, in which 0 indicates a step left (tails), and 1 a step right (heads).

8.2 Restricted random walks

Suppose the walk is bounded, that is, there is some barrier (or barriers) that restrict the range of walker's movement. For example, we consider the case in which there are boundaries to the left and right, and the walk terminates whenever a boundary is reached. The walker is said to be *absorbed* by the boundary. It is not too difficult to calculate the probability of the walker reaching one boundary before the other, the expected duration of the walk until *absorption*, and many other quantities besides. Let's visualize this walk in terms of a game.

8.3 Gambler's ruin

A game is played between a gambler and a 'banker'. A coin is tossed repeatedly in a sequence of identical experiments. The probability of heads is, $0 \leq p \leq 1$, and if this occurs, the gambler wins 1 unit. The probability of tails is $q = 1 - p$; if this arises, the gambler loses 1 unit.

The player starts the game with k pounds. The game ends when either

(a) the gambler is bankrupt (has £0), and is *ruined*.

(b) the gambler reaches a total of N pounds, has won the game, and retires.

The aim of the following calculation is to determine the probability that the gambler ultimately loses (is ruined). Mathematically this problem is equivalent to a simple random walk with absorbing boundaries. As before, the walker steps left or right by one unit depending if the toss is TAILS or HEADS, respectively. However once the walker (gambler) reaches either boundary, $x = 0$ or $x = N$, the walker stops (is absorbed) there permanently and the game ends.

Before we get into the complications, let us consider a simple case to which the solution is obvious. Suppose that $N = 2$. Then beginning at the point $k = 1$ there is only one toss of the coin, the gambler wins or is ruined on the outcome of this single game. The probability of ruin in this case is q, and we should check that our final answer is agrees with this special case.

The walk is called asymmetric when $p \neq q$ (and the game is said to be biased or unfair). Conversely, when the walk is symmetric (and the equivalent game unbiased or fair) then $p = q$.

The game starts with the walker at $x = k$ (that is the gambler has £k).

Let A denote the event that the gambler is ruined, and let B be the event that the gambler wins the first game (toss). Conditioning on the 1st game: Then, by the partition theorem:

$$P(A|\text{starting at } x = k) \equiv P_k(A) = p_k \tag{8.8}$$

$$P_k(A) = P_k(A|B)P(B) + P_k(A|B^c)P(B^c) \tag{8.9}$$

We know that: $P(B) = p$, $P(B^c) = 1 - p = q$. Consider, $P_k(A|B)$. Given the 1st game is won, the walker moves from k to $k + 1$ and continues the game. That is:

$$P_k(A|B) = P_{k+1}(A) = p_{k+1} \tag{8.10}$$

Similarly;

$$P_k(A|B^c) = P_{k-1}(A) = p_{k-1} \tag{8.11}$$

Losing the 1st game, the walker moves from $x = k$ to $x = k - 1$.

Then the conditional probability (partition theorem), equation(8.9), gives:

$$p_k = p_{k+1}p + p_{k-1}q \qquad , \qquad 1 \leq k \leq N - 1 \qquad . \tag{8.12}$$

In addition to this, we have the boundary conditions, for $k = 0$ and $k = N$,

$$p_0 = 1 \qquad , \qquad p_N = 0 \quad . \tag{8.13}$$

This expresses the fact that, if the gambler has no money to begin with, he is certain to lose; if the gambler has N pounds at the beginning, there is no need to play - the gambler has won. The relation (8.12) is called a *difference equation*, and there are a number of different ways of solving such problems. In the following, I'll discuss just 3 of these.

Most simply, bthe difference equations can be expressed as a set of linear equations as follows. The unknowns (and knowns) $\{p_0, p_1, \ldots, p_k, \ldots, p_{N-1}, p_N\}$, are taken to be elements of a vector, y:

$$\mathsf{y} = \begin{pmatrix} p_0 \\ p_1 \\ \vdots \\ p_k \\ \vdots \\ p_{N-1} \\ p_N \end{pmatrix} \tag{8.14}$$

Then the set of difference equations, and boundary conditions, can be expressed as the linear equations:

$$\begin{pmatrix} 1 & 0 & 0 & \cdots & \cdots & 0 & 0 \\ -q & 1 & -p & 0 & \cdots & 0 & 0 \\ \vdots & & & \ddots & & & \vdots \\ 0 & \cdots & -q & 1 & -p & \cdots & 0 \\ \vdots & & & \ddots & & & \vdots \\ 0 & 0 & \cdots & 0 & -q & 1 & -p \\ 0 & 0 & \cdots & 0 & \cdots & 0 & 1 \end{pmatrix} \begin{pmatrix} p_0 \\ p_1 \\ \vdots \\ p_k \\ \vdots \\ p_{N-1} \\ p_N \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix} \tag{8.15}$$

This is a standard linear equation of the form:

$$\mathsf{Ay} = \mathsf{b} \tag{8.16}$$

with a known matrix, A, and given vector b. The unknown vector, y, can be found by Gaussian elimination. However, because of the special structure of the matrix (it's *tridiagonal*), there is a very simple and direct method of solution, described below.

8.4 Solution of the difference equation

A useful technique for treating difference equations is to consider the *trial solution*, $p_k = \theta^k$. In essence this is a *guess*! It's a good guess if we can find a value of θ compatible with the equations, and thus solve the problem.

Using the trial solution, the difference equation (8.12) can be written:

$$\theta^k = p\theta^{k+1} + q\theta^{k-1} \quad , \quad 1 \le k \le N-1 \quad . \tag{8.17}$$

that is,

$$\theta^{k-1}[-p\theta^2 + \theta - q] = 0 \quad . \tag{8.18}$$

The non-trivial solution ($\theta \neq 0$) is given by the solution(s) of the quadratic equation

$$-p\theta^2 + \theta - q = 0 \quad . \tag{8.19}$$

This quadratic equation can be solved by factorization, noting that $\theta = 1$ is a solution since, $-p+1-q = 0$. This gives:

$$-p\theta^2 + \theta - q = (\theta - 1)(-p\theta + q) = 0 \tag{8.20}$$

Hence, the pair of solutions is, $\theta_1 = 1$, $\theta_2 = q/p$.
Since the difference equation is *linear*, the general solution is any *linear combination* of these two solutions:

$$p_k = a_1\theta_1^k + a_2\theta_2^k = a_1 + a_2(q/p)^k \tag{8.21}$$

where $a_{1,2}$ are arbitrary constants that must be determined from the boundary conditions. Given that we have the two boundary conditions: $p_0 = 1$ and $p_N = 0$, these constants are determined uniquely. The boundary conditions for $k = 0$ and $k = n$ give the respective equations:

$$1 = a_1 + a_2(q/p)^0 \quad , \quad 0 = a_1 + a_2(q/p)^N \quad . \tag{8.22}$$

By elimination (subtracting the equations) we have: $1 = a_2(1 - (q/p)^N)$, and therefore, for $q/p \neq 1$,

$$a_2 = \frac{1}{1 - (q/p)^N} \quad . \tag{8.23}$$

Then, it follows that:

$$\boxed{p_k = \frac{(q/p)^k - (q/p)^N}{1 - (q/p)^N} \quad , \quad (q/p) \neq 1} \quad . \tag{8.24}$$

is the solution to the problem. That is the probability of eventually finishing at $x = 0$ having started at $x = k$. It follows that, given the walk must terminate at one end or the other, the probability of leaving the game with the desired fortune of £N will be:

$$1 - p_k = \frac{1 - (q/p)^k}{1 - (q/p)^N} \tag{8.25}$$

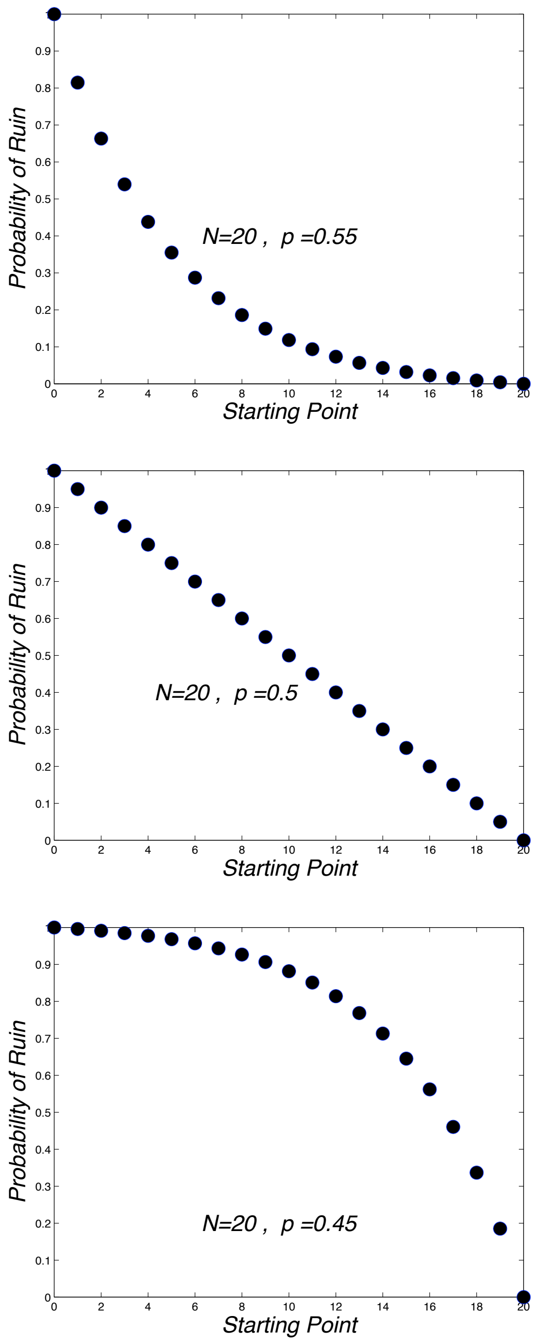

Figure 8.2: . Probability of ruin (termination of the walk at $x = 0$) p_k as a function of starting point, k, for $N = 20$. Top; $p = 0.55$, the steps are biased in favour of the gambler winning each mini-game, Middle $p = 0.50$, each mini-game is fair and thus, if the gambler starts in the middle there is a 50% change of ending up ruined. Bottom $p = 0.45$, there is bias against the gamble winning, and thus an increased probability of being ruined compared with the fair game. Note that even this small degree of bias $p = 0.45$ for each game, means a very high probability of ruin in the long run (over many games). For example starting at $k = 10$ midway between ruin ($x = 0$) and fortune ($x = 20$), there is an 88% probability of ruin. This is a manifestation of the law of large numbers.

8.5 The biased game in the long run

Consider the case $(q/p) > 1$, that is the game is biased *against* the gambler (and this is usually the case). Furthermore, suppose that, $N \to \infty$. In other words, the gambler is aiming to win an unlimited amount of money, and will not stop playing until he is ruined. Not surprisingly,

$$\lim_{N\to\infty} p_k = \lim_{N\to\infty} \frac{(q/p)^k - (q/p)^N}{1-(q/p)^N} = \lim_{N\to\infty} \frac{1-(p/q)^{N-k}}{1-(p/q)^N} = 1 \qquad , p < q \qquad . \tag{8.26}$$

That is the gambler, in the attempt to gain an infinite fortune, is *certain* to lose all his/her money.

In contrast, suppose $(q/p) < 1 \Rightarrow q < p$ that the game is now biased *in favour* of the gambler. Knowing this the gambler aims again to win an unlimited amount of money. That is, he will not quit until he is infinitely rich or bankrupt.

Clearly $\lim_{N\to\infty}(q/p)^N = 0$, and after some simplification, we arrive at the result:

$$\lim_{N\to\infty} p_k = (q/p)^k \quad , \quad k = 0, 1, 2, \ldots \qquad (p > q) \quad . \tag{8.27}$$

That is, even with a game biased in his favour, there is still a non-zero possibility of losing everything if luck runs against him. However, and not surprisingly, this possibility diminishes as k increases as the gambler begins with more money, or equivalently, the walker starts further from the absorbing barrier at $x = 0$.

8.6 Bold play

Suppose the gambler can change the strategy. For example, each game can be played for £2 (or £0.50) instead of £1 pound. In mathematical terms, doubling the bet is the same as doubling the step-size, and this is equivalent to halving the length of the walk. That is, the distance between the boundaries (in terms of the number of steps) is reduced by a factor 2. Similarly, the distance to the boundaries, from the starting point is now half the number of steps. Doubling the stake, £1 $\to$ £2 per game is equivalent to modifying the problem as follows:$k \to \frac{1}{2}k$ and $N \to \frac{1}{2}N$.

Let's consider a concrete case and show how this works in practice. Suppose we have the example of the following biased game: $N = 20$, $k = 10$ and $p = 0.4$, $q = 0.6$. As before, the gambler plays for £1 per game and so the probability of ruin is

$$p_{10} = \frac{(1.5)^{10} - (1.5)^{20}}{1-(1.5)^{20}} = 0.983 \qquad . \tag{8.28}$$

If instead the gambler decides to bet £2 per game, this will shorten the game. Since this is equivalent to $N = 10$ and $k = 5$ with $p = 0.4$, $q = 0.6$, we see that the barriers are now closer, but is this good or bad?

We can calculate this, since raising the stakes in this way changes the probability of loss to:

$$P_{\text{ruin}} = \frac{(1.5)^5 - (1.5)^{10}}{1-(1.5)^{10}} \simeq 0.884 \tag{8.29}$$

That is although the gambler is still likely to lose because the game is inherently unfair, his odds are slightly improved by risking more. Faced with a game biased against the player, is is better to be 'bold' rather than timid with the stakes.

The ultimate way to shorten the game is to reduce it to a single game, and set the stake at £10, the length of the walk is shortened by a factor of 10 so that, $k = 1$ $N = 2$, then:

$$P_{\text{ruin}} = \frac{(q/p)^1 - (q/p)^2}{1-(q/p)^2} = \frac{(q/p)(1-(q/p))}{(1+q/p)(1-q/p)} = \frac{q}{p+q} = q = 0.60 \tag{8.30}$$

and this improves the odds even further. In fact, this single game is the optimal strategy for a game biased against the player. So if Rory McIlroy challenges you to a game of golf according to *match play* rules, you should play only 1 hole. That way you at least have a fighting chance.

When *forced* to play a game in which the odds are not in your favour, a strategy of 'bold play' (or 'aggressive play') maximizes the chances of winning. An even better strategy is *not* to play a game under such conditions, unless you wish to lose! When faced with a game in which the odds are against you, and you are **forced** to play- make the game as short as possible. Play *boldly* or *aggressively* to maximize your chances.

Conversely, when a game is biased in your favour, and the games are independent you should aim to make the game last as long as possible. Suppose that the probabilities were, $p = 0.55$ and $q = 0.45$, and the gambler starts with with £10, with a target of £20. Playing each game for £1, gives the probability of ruin as

$$P_{\text{ruin}} = \frac{(0.45/0.55)^{10} - (0.45/0.55)^{20}}{1 - (0.45/0.55)^{20}} \approx 0.119 \tag{8.31}$$

which is very favourable. But suppose the gambler adjusts the stake to £0.5 per game, then $k \to 20$ and $N \to 40$ so that:

$$P_{\text{ruin}} = \frac{(0.45/0.55)^{20} - (0.45/0.55)^{40}}{1 - (0.45/0.55)^{40}} \approx 0.018 \tag{8.32}$$

which shows that *timid play* is even more effective when the 'odds' are in one's favour. Thus a casino, with bias in its favour, prefers gamblers who are regular visitors who bet frequently and in small amounts.

Before leaving the topic, one needs to careful in developing a betting strategy on these rules alone. If one can vary the step-size from game to game, there is an *optimal strategy* for choosing the step-size.

8.7 The fair game

We speak of a *fair* (unbiased) game if the expected profit for a player is zero. In terms of a fair coin, the probability that is HEADS or TAILS is the same. And for the random walker starting at $X_0 = k$, then the position after the first step ,X_1 has an expectation:

$$\mathbb{E}\left(X_1\right) = (k+1)p + (k-1)q \qquad . \tag{8.33}$$

So in the case of a fair game:$p = q = \frac{1}{2}$.

$$\mathbb{E}\left(X_1\right) = k = X_0 \qquad . \tag{8.34}$$

In this case ($p = q$) the formula for p_k needs to be modified. One can use L'Hôpital's rule to find the expression as $q/p \to 1$. Let $x = q/p$ and consider the limit

$$p_k = \lim_{x\to 1} \frac{x^k - x^N}{1 - x^N} = \lim_{x\to 1} \frac{kx^{k-1} - Nx^{N-1}}{-Nx^{N-1}} = 1 - \left(\frac{k}{N}\right). \tag{8.35}$$

The special case, $k = \frac{1}{2}N$, gives the result $p_{N/2} = \frac{1}{2}$. This makes sense, since, given the walk is symmetric, and the walker starts half-way between the boundaries, there is an equal chance of reaching $x = 0$ before $x = N$.

8.8 Martingales

An elegant solution to the gambler's ruin problem was provided by De Moivre. The trick he suggested is to convert the *biased game* into an equivalent *fair game*, termed a *Martingale*.

In general, the technique of solving problems indirectly, by transforming the problem in a different (and simpler) form is a very powerful method in mathematics. The simplest versions include, changing the variable, integration by substitution etc. while more advanced version include Laplace transforms, Wiener-Hopf methods, and so on.

For a win $X_i = +1$, for a loss $X_i = -1$, and when at the boundary, $X_i = 0$.

$$S_n = S_0 + X_1 + X_2 + \cdots + X_n \quad . \tag{8.36}$$

where $S_0 = k$ is the starting point and: $S_n = l$ is the point after n steps. In the long run ($n \to \infty$), as shown above, the game terminates either with loss or win with the probability:

$$P(S_\infty = 0) \equiv p_k \quad , \quad P(S_\infty = N) \equiv 1 - p_k \tag{8.37}$$

where p_k is given by 8.24.

Consider a (fictitious) mathematical game running in parallel to the real game. In this fictitious game, the gambler plays with *toy money*. The rules of this game are as follows, if the gambler wins a game they get a return of £q/p (of toy money) for every £1 they wager, and a return of £p/q for every game they lose. We can summarise this as follows, if Z_0 is the (fictitious) fortune at the beginning of the game, and the gambler bets that entire amount, then after the first (fictitious) game the fortune is:

$$Z_1 = Z_0(q/p)^{X_1} \tag{8.38}$$

Then the expected value of the fortune after the first game is

$$\mathbb{E}\left(Z_1\right) = Z_0\left[p(q/p) + q(p/q)\right] = Z_0\left[q + p\right] = Z_0 \quad , \tag{8.39}$$

that is, on average we make neither a profit nor a loss. Such a game is said to be *fair* or *unbiased*. Now consider the following betting strategy: for each (and every) subsequent game, the gambler bets the entire amount (of toy money) in their possession, then we have:

$$Z_n = Z_{n-1}(q/p)^{X_n} \tag{8.40}$$

Now since, $\mathbb{E}\left((q/p)^{X_n}\right) = 1$, for any n, then we have:

$$\boxed{\mathbb{E}\left(Z_n | Z_{n-1}\right) = Z_{n-1}} \quad . \tag{8.41}$$

That is, because of the rules of this toy game, the expected value of toy money is the same as that which we started with. Any stochastic process, Z, in which such a relation exists, is called a *martingale*. Clearly:

$$Z_n = Z_{n-1}(q/p)^{X_n} = Z_0(q/p)^{X_1+\cdots+X_n} \quad . \tag{8.42}$$

So if the walk takes us to $x = l$ after n steps, that is: $k + X_1 + X_2 + \cdots + X_n = l$ then:

$$Z_n = Z_0(q/p)^{l-k} \quad . \tag{8.43}$$

Since each game is identical and independent, then according to (8.42)

$$\mathbb{E}\left(Z_n\right) = Z_0 \mathbb{E}\left((q/p)^{X_1}\right) \cdots \mathbb{E}\left((q/p)^{X_n}\right) = Z_0 \quad . \tag{8.44}$$

So according to these rules and this betting strategy, in expectation, we have the same (toy) money at the beginning and ending of the game ($n \to \infty$). The end of the game corresponds to the absorption (stopping) of the walker.

$$\mathbb{E}\left(Z_\infty\right) = Z_0 \quad . \tag{8.45}$$

Now, how does this equivalent *martingale game* help us solve the *real* game. The real game has two outcomes, absorption at $x = 0$ (loss) with probability p_k (as yet unknown) or absorption at $x = N$ (win) with probability $1 - p_k$. In other words, the probability that $x = 0$ after an infinite number of steps, that is $Z_\infty = Z_0(q/p)^{0-k}$, is p_k. Similarly, the probability that $x = N$, after an infinite number of steps, that is $Z_\infty = Z_0(q/p)^{N-k}$, is $1 - p_k$. Therefore, the expected value is:

$$\mathbb{E}\left(Z_\infty\right) = p_k Z_0(q/p)^{-k} + (1 - p_k)Z_0(q/p)^{N-k} \quad . \tag{8.46}$$

Comparing this expression with (8.45), we have

$$Z_0 = p_k Z_0(q/p)^{-k} + (1 - p_k)Z_0(q/p)^{N-k} \quad . \tag{8.47}$$

and thus:

$$p_k = \frac{(q/p)^k - (q/p)^N}{1 - (q/p)^N} \quad . \tag{8.48}$$

In (primary) financial capital markets, one often assumes that asset values (share prices, for example), change according to a stochastic process, though not one as simple as that discussed above. Here Martingale methods are extremely useful in pricing financial instruments derived from the assets: so called *derivatives*. In this case the corresponding *artificial game* involves changing the *probabilities* rather than the rewards (prices). The *fair game* in this case corresponds to making *risk-free* investments such as *bonds*, or eliminating *arbitrage* strategies that allow the trade of risky assets in a risk-free manner. The artificial probability (or *measure*) under these conditions is called the *risk-neutral measure*.

8.9 Mathematics of games

Casino games have relatively simple rules and thus are amenable to mathematical analysis. It is fair to say that these ‘games of chance’ stimulated the study of probability theory in the 18th century. Most of the elementary problems were solved by Laplace, De Moivre, Euler, Cramer, the Bernoullis, and their contemporaries.

Mathematics is extremely useful in analysing these simple problems when our intuition fails us. For example the coin tossing game known as the St. Petersburg paradox (Bernoulli, 1713) describes a game in which a fair coin is tossed repeatedly. The player pays a fee F to take part in the game. The rules of the game are that, if the first heads occurs at the nth toss, the player receives a payment of: $w_n = £2^n$, and the game ends. The game has a maximum duration of N tosses. The question is, what would be a *fair* value for the fee, F? A fair fee would seem to be the expected value of the money won by the player. Given that the probability that the first heads occurs on the nth toss is 2^{-n}, then the expected value for the winnings is given by:

$$\mathbb{E}\left(W_N\right) = \sum_{n=1}^{N} 2^n \times \frac{1}{2^n} = N \quad .$$

So $F = N$ would seem to be fair price for such a game. Suppose the rules of the game where that it continues (indefinitely) until a heads occurs. If this case, $N \to \infty$, then the conclusion is that the fee for such a game would be, $F = \infty$! This is the St. Petersburg paradox. The paradox being that this conclusion conflicts with our intuition that one would consider paying a very large amount of money to play a risky game. In principle one would say that $F = £100,000$ would be a bargain for such an unlimited game, but it flies against our intuitive ideas. The risk of losing such a large amount is too great. In fact, people are naturally skeptical when presented with such an answer (with good reason). Furthermore, people tend to be more conservative (risk averse) when betting large amount of money, such as £100,000. Moreover, they have finite resources and are not able to repeatedly bet successive large amounts of money.

8.10 Is there such a thing as optimal play?

In the previous section, the question of how to minimise the probability of ruin was addressed. That is, with a game biased in favour of the player, a playing strategy that minimises risk when the game is biased in favour of the player would be timid play in which the minimal amount is staked in each game. However such a strategy would take a long time to yield your fortune. Consider an alternative strategy which aims to to maximise the *expected return* on a series of (biased) games while being less concerned about minimising risk. This is called an *optimal betting strategy.* We use the simple example of betting on a series of independent, identically-distributed Bernoulli trials (coin tosses). However, the same ideas can be applied to any investment strategy. A portfolio of investments is essentially a group of bets on the future values of assets in the portfolio.

The player starts, as before, with an *initial capital*, X_0, but this time the opponent allows one to wager any amount on each game. The rules of the game are, if we wager W_n on the nth game and lose, then we lose all our money to our opponent, so the change in our capital is $-W_n$. If we win the game, having wagered W_n, we get bW_n as winnings ($b > 0$), and our original investment back: that is, the opponent returns to us $(1 + b)W_n$. So the change in our capital is bW_n.

The probability of winning/losing the nth game is the same (for any n), $0 < p < 1$. So, if the nth toss is the discrete random variable $T_n \in \{-1, b\}$, where $+b$ is a win, and -1 is a loss then:

$$P(T_n = b) = p \quad , \quad P(T_n = -1) = q = 1 - p \quad . \tag{8.49}$$

In principle one could adjust W_n from game to game. Let’s simplify the strategy to bet a fixed fraction $0 \leq f \leq 1$ of our capital (at the time) on each game. What is the optimal value of f that one can use?

After the nth game our capital, X_n will be worth:

$$X_n = X_{n-1} + W_n \quad . \tag{8.50}$$

where W_n is the increase in capital (the *winnings*):

$$W_n = fX_{n-1}T_n \quad , \tag{8.51}$$

and thus,

$$X_n = (1 + fT_n)X_{n-1} \quad . \tag{8.52}$$

This is a simple coin toss, already discussed in detail, and the calculations are straightforward. Since the outcome of each game is independent of the amount wagered, then T_n and X_{n-1} are independent, our expected winnings are:

$$\mathbb{E}\left(W_n\right) = \mathbb{E}\left(fX_{n-1}T_n\right) = f\mathbb{E}\left(X_{n-1}\right)\mathbb{E}\left(T_n\right) = f\mathbb{E}\left(X_{n-1}\right)(pb - q) \quad . \tag{8.53}$$

So we see that the bias of the game (towards the player) is proportional to: $pb-q$. Thus, $pb-q > 0$ is our mathematical definition of a favourable game. If $pb - q < 0$ then the game is biased *against* the player (the expected winnings are negative) and, with certainty in the long run, the player will be bankrupt.

Then for our capital we have:

$$\mathbb{E}\left(X_n\right) = \mathbb{E}\left(X_{n-1}\right)\left[1 + f(pb - q)\right] = X_0\left[1 + f(pb - q)\right]^n \quad . \tag{8.54}$$

This answers our initial question over how does the expected value of our final capital depend on f. In general the variance in the player's capital after n games can then be derived from:

$$\mathbb{E}\left(X_n^2\right) = X_0^2\left[p(1 + fb)^2 + q(1 - f)^2\right]^n \tag{8.55}$$

and this gives us an estimate of the risk in this strategy. The best strategy for a sequence of identical games will be the best strategy for a single game. After all once we work out the best strategy for the first game, for a Markov process this will be the optimal f for the second and third games, and so on. The variance will (naturally) be proportional to f^2:

$$\mathrm{var}(X_1) = f^2X_0^2pq(b + 1)^2 \tag{8.56}$$

We posed the problem above, what is the best value of $0 < f < 1$ to use? Recall that f is the one variable we can control, the rest are random. The answer appears obvious from the expression (8.54). Clearly, to maximise (8.54), given $n > 0$ and $(pb - q)$, one should choose f to have its largest possible value, $f = 1$. That is wager our entire amount on every game!

Then we have:

$$\mathbb{E}\left(X_n\right) = p^n(b + 1)^nX_0 \quad . \tag{8.57}$$

We can only escape from the game in profit if we win every single game. But the probability of leaving penniless, after n games is: $P(\mathsf{ruin}) = 1 - p^n$ and since $p < 1$: $\lim_{n\to\infty} P(\mathsf{ruin}) = 1$.

So, this strategy is optimal in one sense, maximising the *expected* winnings, but it is against our intuition in the same way as the St. Petersburg paradox. An idea was introduced by Bernoulli (1748), called a *utility transformation*, to reflect the idea that people hate losing much more than they love winning!

We define a function that transforms our *real* money to *utility* money. Let's consider

$$U(X_n) = X_n^\alpha \quad , \quad 0 < \alpha \quad . \tag{8.58}$$

Figure (8.3) shows some typical forms of this function. Clearly:

$$U'(X_n) = \alpha X_n^{\alpha-1} \quad , \quad U''(X_n) = \alpha(\alpha - 1)X_n^{\alpha-2} \quad . \tag{8.59}$$

So U is monotonically increasing with increasing X, but convex for $\alpha > 1$ and concave for $\alpha < 1$. An alternative function, one proposed by Bernoulli (Daniel), which has the concave property would be:

$$U(X_n) = \ln(X_n/X_0) \quad , \tag{8.60}$$

that is,

$$U(X_n) = n\ln(1 + fT_n) \quad . \tag{8.61}$$

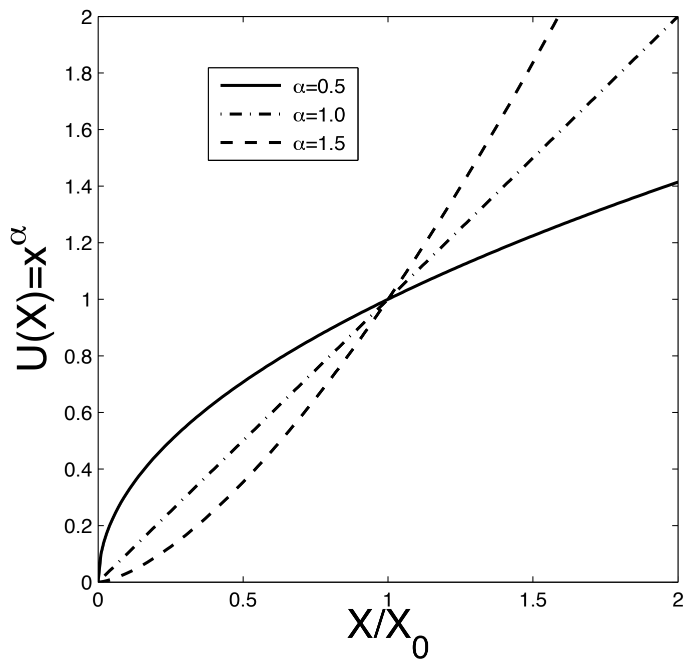

Figure 8.3: The utility function: $U(X) = (X/X_0)^\alpha$. X_0 is our starting value of our money and X its value at a later time. The utility function transforms X (real money) to U an equivalent *perceived value* of money. When $\alpha < 1$ the function is concave and U is less sensitive to changes in X, while for $\alpha > 1$, U is convex and more sensitive to changes in X.

For the log function (8.61) a single win ($n = 1$) would be worth (in utility terms) $\ln(1 + fb)$ but a loss would be $\ln(1 - f)$, and hence would be a strong penalty if $f \to 1$, whatever the bias of the game. The expected value of the utility function is then:

$$H(f) \equiv \mathbb{E}\left(U(X_n)\right) = n\mathbb{E}\left(\ln(1 + fT_n)\right) = np\ln(1 + fb) + nq\ln(1 - f) \quad . \tag{8.62}$$

The value of f that maximizes H, the solution of $H'(f) = 0$, is simply given by:

$$f_K = \frac{pb - q}{b} \tag{8.63}$$

and since $H''(f_K) < 0$ this is indeed a maximum turning point.

As anticipated our answer for f won't depend on n, since this is a Markov process (we apply the same strategy to each game). The strategy of maximizing the expected value of the log of the capital is called the *Kelly strategy* and the value f_K called the *Kelly value* or *Kelly criterion* [1]. The value for f_K is more intuitive, it is simply proportional to the bias of the game (8.53). The more favourable the game, the more one should wager, that is play should be bolder the more the advantage. Compare this to our timid play for the gambler's ruin problem which we recommended for a game in our favour. There is no contradiction between the conclusions. Timid play is aimed at optimising the chances of success as defined as minimising ruin, whereas the Kelly approach aims to optimise the expected value of the utility of the return.

On the game show *Deal or no deal*, a player chooses a box at random from a set of 24. Each box is closed but contains an amount of prize money. Each 'game' involves the player opening three boxes (not his/her own) revealing amounts of money in each of the three boxes. After each game the 'banker' offers to 'buy' the player's box for a price that, naturally enough, depends on the unrevealed prizes. If the player accepts the 'offer', the game ends and the prize is the amount of money for which the box was sold. The offer made to the player by the banker is always lower than the *fair* value (expected value), never higher. Usually the offers begin with derisory amounts to encourage the player to keep playing. Gradually the

[1] J L Kelly Jr. (1956) Bell System Technical Journal **35** p917926

offers become more closely related to the unknown/unrevealed amounts in the boxes. Often the player will accept an offer which is well below the expected value and, mathematically, this is a bad decision. However, from the player's pragmatic point of view, the expected value is relevant to the law of large numbers and not to a single game. Again, refer back to the St. Petersburg paradox. So, for the player, it makes sense to accept the unfair offer if the prize is a significant amount of money. By significant I mean that the sum of money is such that they are prepared to lose it in the attempt to get more.

Of course the Kelly strategy only works if the player can find a game which is inherently biased in his/her favour. In fact, according to the central-limit theorem, any reasonable (sensible) strategy will work under conditions in which the game is in favour of the player. The Kelly factor is just an efficient way to play. The problem is that nearly every opponent (who also knows mathematics) will offer to play such a game only if the bias is against the player.

Of course there are many *gambling strategies* that are non-Markovian (all misguided) in which the gambler tries to vary the amount in a sequence of bets. For example, if the coin comes up 'heads' four times in a row, the naive gambler might triple the bet on the next toss being 'tails'. The gambler might even invoke the 'law of averages' to justify this choice. However, this is an example of where intuition defies logic as well as mathematics.