

2.1 A warm-up on coin tossing

The scores of the coin-tossing game are really a random walk; your score at each stage is the sum of the previous gains or losses. After each toss your score goes up or down by one. Although your average score after N tosses is zero (the game is fair) the scatter in the scores is $\sim \sqrt{N}$. The score at any point is $\sum X_i$, where each X_i can be plus or minus one; the root-mean-square of the score is $\sqrt{\sum X_i^2}$. This means that one player can get a *long* way ahead, and consequently changes of lead are likely to happen near the beginning of the game. By its time symmetry, they are also likely at the end. It also follows that the number of changes of lead is likely to be small, and in fact the most common number of changes of lead is just one.

These are quite easy results to investigate by simulation; here is an example, a histogram of the number of changes of lead happening at each point in a gain 100 tosses long.

These remarkable results are described more in a non-technical way Haigh's book, and in mathematical detail in Feller (An Introduction to Probability Theory and its Applications).

