1 Introduction

The fundamental problem in gambling is to find positive expectation betting opportunities. The analogous problem in investing is to find investments with excess risk-adjusted expected rates of return. Once these favorable opportunities have been identified, the gambler or investor must decide how much of his capital to bet. This is the problem which we consider here. It has been of interest at least since the eighteenth century discussion of the St. Petersburg Paradox (Feller, 1966) by Daniel Bernoulli.

One approach is to choose a goal, such as to minimize the probability of total loss within a specified number of trials, N. Another example would be to maximize the probability of reaching a fixed goal on or before N trials (Browne, 1996).

A different approach, much studied by economists and others, is to value money using a utility function. These are typically defined for all non-negative real numbers, have extended real number values, and are non-decreasing (more money is at least as good as less money). Some examples are $U(x) = x^a$, $0 \leq a < \infty$ and $U(x) = \log x$, where log means $\log_e$, and $\log 0 = -\infty$. Once a utility function is specified, the object is to maximize the expected value of the utility of wealth.

Daniel Bernoulli used the utility function $\log x$ to "solve" the St. Petersburg Paradox. (But his solution doesn't eliminate the paradox because every utility function which is unbounded above, including log, has a modified version of the St. Petersburg Paradox.) The utility function $\log x$ was revisited by J.L. Kelly (1956) where he showed that it had some remarkable properties. These were elaborated and generalized in an important paper by Brieman (1961). Markowitz (1959) illustrates the application to securities. For a discussion of the Kelly criterion (the "geometric mean criterion") from a finance point of view, see McEnally (1986). He also includes additional history and references.

I was introduced to the Kelly paper by Claude Shannon at M.I.T. in 1960, shortly after I had created the mathematical theory of card counting at casino blackjack. Kelly's criterion was a bet on each trial so as to maximize $E \log X$, the expected value of the logarithm of the (random variable) capital X. I used it in actual play and introduced it to the gambling community in the first edition of Beat the Dealer, Thorp, (1962). If all blackjack bets paid even money, had positive expectation and were independent, the resulting Kelly betting recipe when playing one hand at a time would be extremely simple:

bet a fraction of your current capital equal to your expectation. This is modified somewhat in practice (generally down) to allow for having to make some negative expectation "waiting bets", for the higher variance due to the occurrence of payoffs greater than one to one, and when more than one hand is played at a time.

Here are the properties that made the Kelly criterion so appealing. For ease of understanding, we illustrate using the simplest case, coin tossing, but the concepts and conclusions generalize greatly.

2 Coin Tossing

Imagine that we are faced with an infinitely wealthy oppponent who will wager even money bets made on repeated independent trials of a biased coin. Further, suppose that on each trial our win probability is $p > 1/2$ and the probability of losing is $q = 1 - p$. Our initial capital is X_0. Suppose we choose the goal of maximizing the expected value $E(X_n)$ after n trials. How much should we bet, B_k, on the kth trial? Letting $T_k = 1$ if the kth trial is a win and $T_k = -1$ if it is a loss, then $X_k = X_{k-1} + T_k B_k$ for $k = 1, 2, 3, \ldots$, and $X_n = X_0 + \sum_{k=1}^{n} T_k B_k$. Then

$$E(X_n) = X_0 + \sum_{k=1}^{n} E(B_k T_k) = X_0 + \sum_{k=1}^{n} (p - q) E(B_k).$$

Since the game has a positive expectation, i.e., $p - q > 0$ in this even payoff situation, then in order to maximize $E(X_n)$ we would want to maximize $E(B_k)$ at each trial. Thus, to maximize expected gain we should bet *all of our resources* at each trial. Thus $B_1 = X_0$ and if we win the first bet, $B_2 = 2X_0$, etc. However, the probability of ruin is given by $1 - p^n$ and with $p < 1$, $\lim_{n\to\infty}[1 - p^n] = 1$ so ruin is almost sure. Thus the "bold" criterion of betting to maximize expected gain is usually undesirable.

Likewise, if we play to minimize the probability of eventual ruin (i.e., "ruin" occurs if $X_k = 0$ on the kth outcome) the well-known gambler's ruin formula in Feller (1966) shows that we minimize ruin by making a *minimum* bet on each trial, but this unfortunately also minimizes the expected gain. Thus "timid" betting is also unattractive.

This suggests an intermediate strategy which is somewhere between maximizing $E(X_n)$ (and assuring ruin) and minimizing the probability of ruin (and