The Kelly Criterion

A closer look at how estimation errors affect portfolio performance

by Adrian Halhjem Sælen

Advisor: Professor Steinar Ekern

Master Thesis in Financial Economics

NORGES HANDELSHØYSKOLE

Summary

The Kelly criterion are being used by many succesful investors and are often linked to their success in amassing great fortunes for both themselves and the funds they are running. This simple formula has made such a great impact on their results, that it is often cited as the main reason for these fund managers outperforming the indexes and other investors. However, there are several problems with using the Kelly criterion, and the main one involves the big risk that is being taken by sizing ones portfolios using the formula.

This thesis will show that the problem with large risk becomes even greater when the errors in our estimation of the inputs start to get larger and larger. When our investor's estimation errors increase, the volatility of the portfolio will not only increase as a result of this, but the growth of the wealth will fall. The conclusion that can be drawn from this is that risky assets, that is, assets that are very hard to determine the actual returns and volatility of with good accuracy, should use a fractional Kelly size ($c < 1$) when investing, to avoid overbetting on the asset, or even refrain from using the Kelly criterion at all.

Contents

1. Introduction

1.1 Background

When investors are to allocate their money into favorable investments, usually their most important goal is to make as much money as possible. For those investors who only are concerned with maximizing their wealth, using the Kelly criterion when sizing their portfolio positions makes perfect sense.

The Kelly criterion, which was first introduced by John R. Kelly in his famous article on the information rate, has proven to optimize the growth of an investors' wealth if one follows the formula of Kelly diligently. More precisely, Kelly tells you how much to invest in a given asset where you know the expected return and the volatility of the asset if you want to maximize the expected growth of your wealth. This presentation, however, makes it seem that Kelly is the ultimate portfolio strategy for every investor out there, but that is not the case.

When planning a portfolio strategy, the two things usually one has to decide is what is more important to the investor, high return or low risk. The consensus in today's knowledge of asset management is that we for the most part cannot have both higher than average return and lower than average risk, and the investor needs to decide which one is the most important to them in their specific situation.

Institutional investor's, like governments and pension funds which have significant liabilities to take care of, would be more interested in keeping their risk profile low in order to not sustain a loss beyond their liabilities. For investors like high-risk hedge funds or private investors which are more interested in maximizing their wealth, having higher return often comes at the cost of having lower risk, and investors such as Edward O. Thorp in the Princeton-Newport Fund, Warren Buffet in his company Berkshire Hathaway, John Maynard Keynes' King's College Chest Fund and different gambling syndicates all have in common that they have applied the Kelly criterion when creating their impressive results.

The formula that has led to such fortunes is quite easy to recite and use:

$$f^* = \frac{Bp - q}{B}$$

f^* = The Kelly fraction (the percentage of your wealth you want to wager)

B = The net return amount you are given if you win

p = Probability of winning

q = Probability of losing ($q = 1 - p$)

The only calculations that is needed, is the probability of winning (which in turn gives you the probability of a loss) and the net return you expect to get when you win. So why is it that so many are not using the Kelly, which has proven to optimize the growth of an investors' wealth over the long run?

One of the biggest problems with using the Kelly criterion, and any portfolio strategy in general, is the problem of getting accurate estimates for the inputs needed to calculate the return, and in turn, the fraction of our wealth to invest in the investment objective at hand. This might seem trivial enough for basic betting opportunities like playing roulette in a casino. In such a scenario, you would know the exact probability of winning and the exact return you get if you win. However, once we start to invest in areas where we encounter more uncertainty, like sports betting where we don't know our exact win probability, or the stock market where we neither know our exact win probability, nor the exact return we would expect to receive, we have to start estimate our inputs. This adds uncertainty to what our actual return and risk will be, and may lead us to bet either too much or too little, and as we will see, may lead to disastrous results for our wealth.

My motivation for writing this thesis about the Kelly criterion is that I would like to know to what extent this problem with estimation errors hurts an investor's return and growth of his wealth. The values that we calculate later in our examples in chapter 2 will show the optimal amount to bet on a proposition. If we deviate from this, be it betting more or less than the optimal fraction f^*, we must expect to see lesser gains (and possibly higher variance if we bet a fraction $f > f^*$). This will happen frequently when all we can do is to make assumptions about investment objectives and how their values and returns will change as we

move into the future. I want to look at how these estimation errors affect the growth and return we get on our portfolio in different scenarios and hopefully gain more insight as to what we need to watch out for when we are dealing with a lot of uncertainty.

There are several examples throughout history where investors have underestimated the probability of extreme events happening, which make investors drastically overbet compared to if they had known the "true" probabilities for every event that could have occurred. The best example of this might probably be the fall of Long Term Capital Management which went bankrupt doing arbitrage trading in Russian government bonds.

1.2 Thesis question

In this thesis, I would like to direct my focus onto the performance of portfolios that make use of the Kelly criterion and specifically how the performance shifts as the uncertainty in our estimations of the investment objectives expected return and/or variance increases.

First I will present the Kelly criterion in chapter 2, as well as give examples of how it is used, discuss the strengths and drawbacks, and expand upon the problem with estimation errors. Chapter 3 will contain the main part of this thesis with the simulation of a basic betting proposition, where one bets on the outcome of a coin flip. Here I add some uncertainty in the form of imperfect information provided to the bettor and note the impact this has on the wealth. We then move on to a more complex scenario where I will take a closer look at how the Kelly criterion will perform in a stock market that we will simulate using excel and Brownian motion. Chapter 4 concludes and sums up what we find in our simulations.

2. The Kelly Criterion

2.1 The Kelly fraction

For investors, the main objective is to find investments where one can expect positive returns. This, however, is only part of the equation in becoming a profitable investor. When you know how to find investments that will provide higher expected returns and similar risk as what the stock market will offer, you need to figure out how much to invest in these ventures to make the most of every opportunity presented to you. One example of such a portfolio strategy is using the Kelly criterion to calculate how much you are willing to invest in an asset. Kelly is a very aggressive investment strategy, which seeks to maximize the growth of the users' wealth. It's important to note that we are maximizing the expected growth of our wealth, and not the absolute return amount of our portfolio. If we were to maximize our expected return, we would simply bet everything we had on every opportunity that will be presented to us that promised above average returns.

However, there are several shortcomings and problems with using the Kelly criterion. As we will show in this thesis, Kelly is most certainly not for everyone, not even professional money managers. So firstly I will introduce the Kelly criterion and explain it using a simple coin tossing example. Next I will go over some of the critique that it has received by other famous economists throughout the years, and then focus on the practical problems that it faces when you are calculating the expectation of different assets and the uncertainty surrounding these calculations.

The problem at hand is to find how much we want to invest in an opportunity presented to us. If we use the Kelly criterion when we are deciding how much to invest, we are looking to find a fraction of our current wealth to wager. To calculate this fraction, we need to know the net return we are expected to receive from our investment and the probability of our investment to payoff and indirectly we find how often we will lose our investment as well. The formula for the optimal Kelly fraction, as I presented it in the introduction:

$$f^* = \frac{Bp - q}{B}$$

To better illustrate the formula, let us use a simple game of coin tossing as an example:

Let's say that you offered a bet on a single coin toss from a stranger. The stipulations are such that you can wager whatever amount you would like and if you guess correctly on which side the coin will land, you get a return of twice the amount you wagered and lose your whole wager if you guess incorrectly. Here P_n is your wealth after the wager is complete and P_0 is your wealth before the wager. Your arithmetic net return rate after a single coin toss if you win or lose is as follows:

$$r_W = \frac{P_n - P_0}{P_0} = \frac{3 - 1}{1} = 2 = B$$

$$r_L = \frac{P_n - P_0}{P_0} = \frac{0 - 1}{1} = -1$$

We assume that we are dealing with a fair coin in this game, which gives us the probabilities of heads and tails of $\frac{1}{2}$ each. Since it doesn't matter whether we choose heads or tails in this game, our chance of winning is equal in both instances, our probabilities for winning and losing becomes:

$$p = p_W = \frac{1}{2} \ , \qquad q = p_L = \frac{1}{2}$$

From this we can calculate the expected arithmetic return of this bet:

$$E[r] = p_W r_W + p_L r_L = \frac{1}{2} \cdot 2 + \frac{1}{2} \cdot (-1) = \frac{1}{2} \to 50\%$$

As we find that we have a positive expected return in this proposition, we know that we at least want to invest some of our wealth, at least if the condition of $E[r] > R_F$ holds. This can be formally proved in an expected utility context, as illustrated below.

The definition of final wealth:

$$\tilde{Y} = W + fW\tilde{R} + (W - fW)R_F \to \tilde{Y} = W[(1 + R_F) + f(\tilde{R} - R_F)]$$

$\tilde{Y}$ $\quad$ = Final wealth

W $\quad$ = Starting wealth

f $\quad$ = Fraction bet on the proposition

$\tilde{R}$ $\quad$ = Stochastic net return on the risky asset

R_F $\quad$ = Return on the risk free asset

$u(Y)$ $\quad$ = Utility function (with positive marginal utility $u'(Y)$)

Now we want to maximize the utility of our final wealth:

$$E[u(\tilde{Y})] = E[u[W[(1 + R_F) + f(\tilde{R} - R_F)]]]$$

To maximize this, we differentiate the expression with respect to f and get the first order optimality condition:

$$E[u'(\tilde{Y})(\tilde{R} - R_F)W] = 0$$

Trying out $f = 0$:

$$E[u'(\tilde{Y})(\tilde{R} - R_F)W]_{|f=0} = Wu'(W(1 + R_F))E[\tilde{R} - R_F]$$

With a positive marginal utility, this expression will have the same sign as $E[\tilde{R} - R_F]$. This means we will always invest some of our wealth in the risky asset when $E(\tilde{R}) > R_F$.

But the question still remains as just how much one wants to invest on this coin flip. Some might say that this question is trivial and of lesser importance than actually finding these investments that provide such great returns. I shall illustrate why this statement may lead to

severe trouble for an investor. Let's look at another example of a favorable investment opportunity where we this time will generalize the inputs:

The same stranger offers you yet another proposition, this time it is a lottery that is presented to you. From a pool of N tickets (where $N > 1$), the lottery master will sell you one (and only one) ticket for a price you choose yourself and if your ticket is drawn, will reward you with $1{,}5N$ times the amount you bought your ticket for. If you lose, you get zero and lose the amount you wagered. Again we calculate the arithmetic returns:

$$r_W = \frac{P_n - P_0}{P_0} = \frac{1{,}5N - 1}{1} = 1{,}5N - 1$$

$$r_L = \frac{P_n - P_0}{P_0} = \frac{0 - 1}{1} = -1$$

The probabilities that we deal with here are not the same as in the coin tossing example, where $N = 2$. Now the probability of winning is 1-to-N, such that:

$$p = p_W = \frac{1}{N} \quad and \quad q = p_L = \frac{N - 1}{N}$$

From this we can calculate the expected return from this lottery:

$$E[r] = p_W r_W + p_L r_L = \frac{1}{N} \cdot (1{,}5N - 1) + \frac{N - 1}{N}(-1) = \frac{1}{2} \rightarrow 50\%$$

We see from this that both of these propositions have the exact same expected arithmetic return of 50% of our wagered amount, but would we wager the same amount on both of these bets?

If we first look at the two extremes of betting a fraction, f, of our wealth, W, of 0 and 1 respectively, we will first recall that betting nothing ($f = 0$) would probably never be a smart idea given that we have shown how this bet will produce above average return that severely trumps the risk free rates that we could hope to find in any market. Here we assume that the net return is 50% and finding an investment that will provide greater return and be essentially risk free in the real world is impossible.

Investing the full amount ($f = 1$) of one's wealth can also never be the right answer. This would indeed maximize the expected return of his capital, but this would also include a

probability of $\left(\frac{N-1}{N}\right)$ of losing everything. Putting yourself in a position where there is a non-zero chance of losing your entire wealth is bad practice for any investor, and the bigger N is, the riskier this bet is. So the answer for f must be somewhere between 0 and 1. This depends on the investor's aversion to risk, and thus the utility this individual places on money.

Utility theory states that item A has more utility than B to that individual, if presented with having to choose between the two, the individual would choose A. For most things, these are subjective and hard to put and objective figure on why one would choose one item over the other, for example choosing between an apple or an orange. However, it is easier when we are comparing amounts of money and risk. These are perfectly quantifiable and are easy to compare. If an individual had to choose between n amounts of money and $n + 1$ amounts of money, everyone would prefer the latter. Everyone prefer more money, rather than less.

However, for most people, there are diminishing utility as for more money you acquire. If you have $0, getting $10 would yield a lot of utility to you. If you have $100.000, having another $10 extra would not mean that much to you. Another way of stating this is that people have concave utility functions.
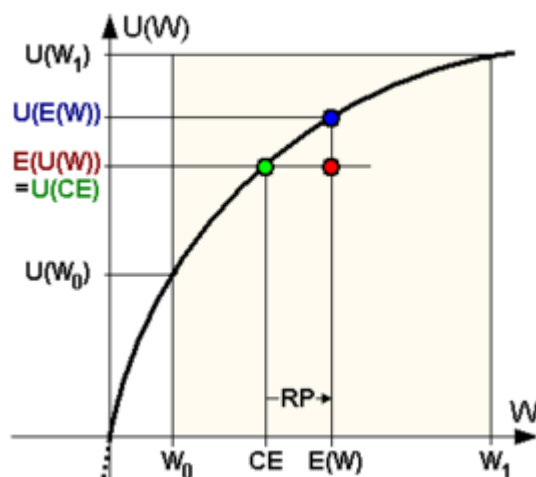

Figure 2.1

This shows that people are willing to accept a certain sum of money (the certainty equivalent, CE) that is smaller than the expected value $E(W)$ of a gamble. How much smaller depends on how risk averse an individual is. If the person is not risk averse at all, or for some reason is risk seeking, then they would want the $CE \geq E(W)$. For most rational people, this is not the case, and everybody has to some degree some aversion to risk, but just how much that is, is not always easy to quantify.

Let's assume that the investor invests a fraction f of his capital on each bet. Then we have the final wealth:

$$\tilde{Y} = (1 + fB)^M (1 - f)^L W$$

where N is the times we repeat the bet. M and $L(= N - M)$ are the times we won the bet and lost the bet respectively. If we also use the quantity G, called the exponential rate of growth of the investor's capital as defined by Kelly:

$$G = \lim_{N \to \infty} \frac{1}{N} \log \frac{\tilde{Y}}{W}$$

Taking logarithms, and moving $\frac{1}{N}$ inside the equation brackets:

$$G = \lim_{N \to \infty} \left[\frac{M}{N} \log(1 + fB) + \frac{L}{N} \log(1 - f) \right]$$

Taking limits, using the strong law of large numbers:

$$G = p \log(1 + fB) + q \log(1 - f)$$

$$p = \frac{M}{N}, q = \frac{L}{N}$$

If we want to maximize our growth, we maximize the expression with regards to f:

$$f = \frac{Bp - q}{B}$$

which we recognize as the Kelly criterion we presented above in chapter 1.

So if we go back to the question as to how much we would wager on the two different bets we were presented with, we can plug the parameters into the Kelly formula:

$$f_{Coin\,flip} = \frac{2 \cdot \frac{1}{2} - \frac{1}{2}}{2} = \frac{1}{4} = 25\%$$

$$f_{Lottery} = \frac{(1{,}5N - 1)\left(\frac{1}{N}\right) - (1 - \frac{1}{N})}{1{,}5N - 1} = \frac{1}{3N - 2}$$

The results here tell us which fraction of our wealth we should wager on each individual bet that we are presented with in the coin flipping and lottery examples respectively. As we can see, the amounts can differ vastly, depending on how large N will be. But as was pointed out earlier, the expected arithmetic returns for both of these bets are the same, no matter the value of N, so how come they end up with differing fractions to bet?

The answer could lie in the fact that there is a difference between the expected value from a bet and the expected growth. We use the formula we derived above to find these values:

$$G_{Coin\,flip} = (1 + fB)^p (1 - f)^q - 1 = \left(1 + \frac{1}{4} \cdot 2\right)^{\frac{1}{2}} \left(1 - \frac{1}{4}\right)^{\frac{1}{2}} - 1 = \sqrt{\frac{3}{2} \cdot \frac{3}{4}} - 1 \approx 6,066\%$$

$$G_{Lottery} = (1 + fB)^p (1 - f)^q - 1$$

$$= \left(1 + \left(\frac{1}{3N - 2}\right) \cdot (1,5N - 1)\right)^{\frac{1}{N}} \left(1 - \left(\frac{1}{3N - 2}\right)\right)^{\frac{N-1}{N}} - 1$$

$$= \left(1 + \frac{1}{2}\right)^{\frac{1}{N}} \left(1 - \left(\frac{1}{3N - 2}\right)\right)^{\frac{N-1}{N}} - 1$$

If we set $N = 2$ in the formula for $G_{Lottery}$, we will find that we get the exact same expression as for $G_{Coin\,flip}$:

$$G_{Lottery} = \left(1 + \frac{1}{2}\right)^{\frac{1}{2}} \left(1 - \left(\frac{1}{3 \cdot 2 - 2}\right)\right)^{\frac{2-1}{2}} - 1 = \sqrt{\frac{3}{2} \cdot \frac{3}{4}} - 1 \approx 6,066\%$$

If we increase the value and set $N = 3$:

$$G_{Lottery} = \left(1 + \frac{1}{2}\right)^{\frac{1}{3}} \left(1 - \left(\frac{1}{3 \cdot 3 - 2}\right)\right)^{\frac{3-1}{3}} - 1 = \sqrt[3]{\frac{3}{2} \cdot \left(\frac{6}{7}\right)^2} - 1 \approx 3,292\%$$

We can see that as N increases, G shrinks. Thus for bets where we have larger values of N, the expected growth shrinks. What this means is that we would rather invest in assets that provide a higher win probability than others, if the expected return is equal in both of the bets. Intuitively, this makes sense as investing in high win probability assets will give a more steady income stream than low win probability assets and thus we can expect to compound our wealth for a longer time than if we invest in the low win probability asset which could

experience a lot longer stretches of losses. We can also see that the coin flip gamble has a second-order stochastic dominance over the lottery (as long as $N > 2$), which is preferred by most all of the rational investors that exist in the world that has an increasing and concave utility function. A second-order stochastic dominance implies that an individual prefers the coin flip over the lottery, based not on first-order stochastic dominance (which means that the $E_{Coin\,flip}[r] > E_{Lottery}[r]$, which is not the case), but on the fact that $\sigma_{Coin\,flip} > \sigma_{Lottery}$, which is true. The risk in the lottery will be higher than the coin flip when $N > 2$.

What this example illustrates is that the expected return an asset promises to provide is not the be all and end all goal to find. We also need to consider the probability of getting these returns. If we set the *N* high enough in the general formula, we can even get negative growth values.

2.2 Sensitivity to the mean

As seen above in our coin flipping example, the calculation of Kelly itself is quite easy when all of the input values are known to us. However, when we explore areas where there are unknown inputs that we need to estimate ourselves, problems start to arise, not only when it comes to actually being able to estimate these somewhat correctly, but also doing this on a consistent basis. Let's introduce another scenario to better explain this.

A simple example that may look like the coin flipping experiment, could be betting on the outcome of a sporting event, like head to head betting on Tour de France (cycling). Here we simply bet on who we think would come first over the finish line between two cyclists for a given race. Let's use Cadel Evans and Thor Hushovd. If we also assume that we make a wager with a friend and thus there is no vigorish (this is how the bookmaker makes money, by shading the lines a little and so they pay out a little less money than they get in). Now let us say that your friend offers you 1-to-1 payoff (or a 100% arithmetic return) on a bet on Thor Hushovd ($r_T = 100\%$). This implies that he thinks Cadel Evans will win more than $\frac{1}{2}$ of the time $\left(p_C = \frac{1}{2}\right)$, and thus the same probability of Hushovd winning $\left(p_T = \frac{1}{2}\right)$. If this where the case, than the expected return on this wager would be:

$$E[r] = p_T r_T + p_C r_C = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot (-1) = 0$$

For this bet to have a positive expected value for you, Hushovd needs to win more than $\frac{1}{2}$ of the time. But how can you come up with a good estimate for Hushovd's win probability?

There are several methods that people use when they are trying to calculate these scenarios, some more robust than others. Most people will simply ask their brain how he has done in the races you have seen for yourself and make a subjective decision as to whether or not they think Hushovd has more than $\frac{1}{2}$ chance of winning, which for most people is almost impossible to do accurately. And we need to know not only if Hushovd will win more than $\frac{1}{2}$ of the time, we need the exact probability that he will do so for us to get a correct Kelly fraction to bet. Doing this subjectively is, to put it lightly, not an easy feat.

Most professional sports bettors and those who actually do find profitable bets on a consistent basis do so by applying some form of objective prediction here. The most common one is to try to create a predictive model for the bet at hand. By collecting data for the individuals involved and use this to determine both the contestants chances of ending up ahead of each other, we can place an estimate on both the athletes win probabilities.

However, these are just that, estimates. We might be able to guess correctly on the win probabilities on average and thus end up ahead even though we most of the time are wrong (both on the good and bad side). When I say ahead here, I mean that if we flat bet an amount equal to $1 on every bet, we expect to have a larger wealth after we have made our bets on average. But a Kelly bettor does not only care about expectation. If we were just looking to maximize our expected win amount, we would risk everything we own on this bet. For most people though, one can clearly see that this strategy has severe risk of going broke, as sooner or later, one will face a losing bet and thus go bankrupt. What the Kelly bettor is concerned with though is maximizing the growth of their wealth, and seeing as risking ones entire wealth on a single bet that has a probability $p_L > 0$ of losing would bankrupt the investor, this cannot be optimal.

This is where the term expected growth comes into play. As earlier noted, most people in different facets of life and especially in finance and economics that deal with different types of investments and bets that are uncertain as to how much it will pay off and how often it will do so, will be familiar with the term expected value. This is what we would expect to

happen on average when we invest in an asset, or if we invested in it several times over at the exact same conditions, we would expect to end up with the expected value. As also noted earlier, with the extreme comparison of the unlikely lottery and the coin flip proposition which each had the same amount of expected value, we obviously note that we will invest different sums of money on those two, and expected growth can show why this is.

Expected growth simply states how much we expect our wealth to grow for a given investment. This is also just as arbitrary as thinking of expected value, where the expected value of a bet could be 0,50, but we either win 2 or lose 1, so we will never see the expected value come up after only one bet, or we might never end up with the value we expect. The process of compounding causes the distribution of future investment results to become skewed in such a way that there are fewer values in excess of the average value, but they exceed it by a greater amount, on balance, than the amount by which the more plentiful below-average values fall short of the average value. The same goes for expected growth. In the coin flipping example, we would expect a growth of 6,066% per bet (I arrive at this value by using the formula: $E[G] = (1 + Bf)^p (1 - f)^q - 1$), but we will never see an exact increase of such an amount, but if we repeat the bet several times, we can expect our wealth to grow by this amount per wager we make.

As I previously noted, we only got estimates for our win probabilities in our case, and while we on average will be right on the money with our expectation, we will not always be there on the individual wagers. One would initially think that if we on average get the amount we should bet right, that should even out and result in the same outcome as getting it right everytime, meaning we end up with what we would expect. However, this is not the case when it comes to growth of wealth. We maximize the growth of our wealth when we get the completely accurate inputs and bet the Kelly fraction we calculate from this. But if we overestimate the win probability of our bet, we will overbet and if we underestimate it, we will bet too small or maybe not even at all. Let's show an example of this:

Let us go back to our coin flipping example for a moment and this time we offer the bet to a friend with the odds being $\frac{11}{10}$ ($B = 1,1$). We know that we are dealing with a fair coin and that there is a 50% chance of it landing on heads and another 50% chance of it landing on tails. Now we offer this bet to a friend of ours that is allowed to examine the coin and after he has done this, he is convinced that the chance of the coin landing on heads is: $\hat{p} = \frac{11}{21}$. We,

off course, know better, and are fully aware that this is a perfectly fair coin. This friend of ours is also familiar with the Kelly criterion, and he calculates his optimal betting amount and expected growth from these values:

$$\hat{f} = \frac{B\hat{p} - \hat{q}}{B} = \frac{\left[\frac{11}{10} \cdot \frac{11}{21} - (1 - \frac{11}{21})\right]}{\frac{11}{10}} = \frac{1}{11} = 9{,}090\%$$

$$\hat{G} = (1 + \hat{f}B)^{\hat{p}}(1 - \hat{f})^{\hat{q}} - 1 = \left(1 + \frac{1}{11} \cdot \frac{11}{10}\right)^{\frac{11}{21}}\left(1 - \frac{1}{11}\right)^{\left(1 - \frac{11}{21}\right)} - 1 = 0{,}4549\%$$

Our friend agrees to the wager and will bet an amount equal to $\frac{1}{11}$ of his wealth on each coin flip, in the belief that he will be able to grow his wealth with 0.4549% on each coin flip.

Since we know that his estimates of this coins bias are wrong, we can calculate his actual expected growth:

$$G = (1 + \hat{f}B)^{p}(1 - \hat{f})^{q} - 1 = \left(1 + \frac{1}{11} \cdot \frac{11}{10}\right)^{\frac{1}{2}}\left(1 - \frac{1}{11}\right)^{\frac{1}{2}} - 1 = 0{,}00\%$$

This gives us a growth rate of exactly zero. This means that our friend cannot expect to grow his wealth betting on this coin flip. His overestimation of this coin's probability of landing on heads has lead him to overbet. If he had known the true probabilities, he would have calculated the Kelly fraction and the growth to be:

$$f^* = \frac{Bp - q}{B} = \frac{1.1 \cdot 0.5 - 0{,}5}{1.1} = \frac{1}{22} = 4{,}545\%$$

$$G_{Coin\,flip} = (1 + f^*B)^{p}(1 - f^*)^{q} - 1 = \left(1 + \frac{1}{22} \cdot \frac{11}{10}\right)^{\frac{1}{2}}\left(1 - \frac{1}{22}\right)^{\frac{1}{2}} - 1 = 0{,}1136\%$$

We can see from this example that he is betting exactly double the optimal bet amount. This is the same as betting double Kelly. If we overbet, in the sense that we bet more than the Kelly fraction we would get if we had calculated the true win probability, we not only will decrease the growth rate of our wealth, but also increase the risk that our investments will have. We will see the same effect when we underbet, where the growth will be lower than usual, however, this is a more forgiving scenario for the investor as the risk will at the same time be lowered.

Figure 2.2

This graph shows the different growth rates and standard deviations we get when betting a multiple of the optimal Kelly fraction. A multiple of 1.00 is the same as betting full Kelly and will produce the optimal growth. The formulas used to calculate the values can be shown, for example with the multiple of 1,5:

$$E[r] = p_T r_T + p_C r_C = \frac{1}{2} \cdot \frac{11}{10} + \frac{1}{2} \cdot (-1) = 0,05$$

This value for the expected return is generalized value. For every 1 unit of money that we wager, we expect to return 1,05 from that, for a net return of 0,05. In our example, we have calculated an optimal Kelly fraction of $\frac{1}{22}$ and will also use a fraction $c = 1,5$ for this example.

$$G_{Coin\ flip} = (1 + fB)^p (1 - f)^q - 1 = \left(1 + 1.5 \cdot \frac{1}{22} \cdot \frac{11}{10}\right)^{\frac{1}{2}} \left(1 - 1.5 \cdot \frac{1}{22}\right)^{\frac{1}{2}} - 1$$
$$= 0,08519\%$$

$$\sigma^2{}_{Coin\ flip} = \sqrt{p\big((cfB) - cfE[r]\big)^2 + q\left(\big(cf \cdot (-1)\big) - cfE[r]\right)^2}$$

$$= \sqrt{\frac{1}{2}\left(\left(1,5 \cdot \frac{1}{22} \cdot \frac{11}{10}\right) - 1,5 \cdot \frac{1}{22} \cdot 0,05\right)^2 + \frac{1}{2}\left(\left(1,5 \cdot \frac{1}{22} \cdot (-1)\right) - 1,5 \cdot \frac{1}{22} \cdot 0,05\right)^2}$$

$$\approx 0{,}51253\%$$

2.3 Critique of Kelly

I have now shown what Kelly is, and have gone into detail on how this formula will grow your wealth optimally if used correctly. However, there are several reasons for not using the Kelly criterion, even though it promises such vast fortunes if taken into account. Even the famous Edward O. Thorp, who has made this formula famous by using it to beat blackjack and create an astonishing wealth in his investment fund, cites several shortcomings of the Kelly criterion.

The first and most prominent problem with applying the Kelly criterion to a portfolio strategy is the fact that it comes with great risk. Even though people prefer to make the most money possible out of their investments, they still have to weigh in on how much risk they are willing to take to possibly be able to make all that money. This off course depends on the individual who are investing, and that persons tolerance to risk. The obvious answer to such an individual who might not be comfortable with betting such large sums of money that Kelly demands, might be more inclined to apply a fractional Kelly, f=cf*, where 0<c<1. By using a smaller fraction of the optimal Kelly fraction, the individual will experience less growth, but also drastically lower the risk on their portfolio.

Another problem with the Kelly criterion is the practical problem of how to handle simultaneous investments, how to account for several investment ideas that you have come up with at the same time. Let's say you have found two companies to invest in, that have promising prospects for the future and you calculate that both warrants a f*=1/2. Now, how do you account for this problem where 2f*=1. Since f*<1, it is a non-zero chance of losing a lot of money at both these investments, and you are leaving yourself vulnerable to bankruptcy. The alternative is to position yourself in one of the companies first, and then in

the other afterwards, so that only 3/4 of your wealth is invested. However, this will not be optimal either, as both investments where deemed equally good, and should be given the same amount of money. In addition to all of this, is also the problem of correlation between investments. It may be hard to calculate these numbers accurately, and they can pose a threat where one overestimates the Kelly fraction if we do not take correlation between assets into consideration.

The Kelly criterion is not the be all, end all for money management, there is more to it than that. However, if you are in the business of maximizing your wealth, you should look into Kelly. The following chapter will take a look at how estimation errors plays a role in using it in practice.

3.   Simulations

The purpose of this thesis is to evaluate how the Kelly criterion behaves when we introduce uncertainty to our estimates of the inputs used in the formula. It has already been shown that investing Kelly will grow your wealth optimally, but we want to see how the performance of our portfolio changes as the errors in our estimations increases.

3.1 Analytical calculations

Before we delve into the actual simulations of this scenario, I would like to try and calculate what results we could expect to find when we apply the Kelly criterion to our coin flipping example. We will also look at how the derivative of the Kelly criterion performs, half Kelly (where we calculate the Kelly fraction we bet and divide it by two), and betting a flat amount on every coin flip, regardless of our wealth growth.

First we look at the possible outcomes that can come of this scenario. We either guess right and win b times our wager amount, for a net return of B, or we guess wrong and lose our wager, which gives us a net return of -1. If we assume that we will be betting a fraction f of our wealth, we either get the net return of fB or $-f$. We can describe the possible outcomes from the coin flipping example by using binomial trees:


Figure 3.1 and 3.2

The stipulations for the bets never change and we repeat it with the same relative values in each period. Given this, we can calculate both an up and down factor (u and d), where

$u > 1$ and $d < 1$. By calculating these values, we can easily expand our calculations to n subsequent trials:

$$1 + fB \equiv u$$
$$1 - f \equiv d$$

We start out with our starting wealth W at time 0, and the ending wealth after one period will be either W_u or W_d, with probabilities p and q respectively. The factors in our example are:

$$u = \frac{21}{20} \ and \ d = \frac{21}{22}$$

Figure 3.3

We find these values for this one period game:

Mean gross return: $\varepsilon(r_A) = 1 + \frac{1}{2}\left[\frac{1}{20} + \left(-\frac{1}{22}\right)\right] = 1 + \frac{1}{440} \approx 1{,}002273$

Standard deviation: $\sigma(r_A) = \frac{1}{2}\left[\frac{1}{20} - \left(-\frac{1}{22}\right)\right] = \frac{21}{440} \approx 0{,}047727$

Geometric return: $G = \left(\frac{21}{20}\right)^{\frac{1}{2}} \left(\frac{21}{22}\right)^{\frac{1}{2}} - 1 \approx 0{,}001136$

Now we take this analysis one step further and look at how a two period game would look like. This time we flip the coin two times and we still decide to use the Kelly criterion when deciding how much to bet. Notice that this time, we will sometimes bet more or less in our second trial since we have either increased or decreased our wealth after the first trial

depending on the result of the coin flip. Since the changes in our wealth are all relative to our absolute amount, we can generalize it, like we have already done, and simply need to multiply the factors in the outcome tree to find the different possible ending wealth's we can end up with:


Figure 3.4

The results here came about by simply multiplying the relevant factors with each other:

$$uu = \frac{21}{20} \cdot \frac{21}{20} = \frac{441}{400} = 1,1025$$

$$ud, du = \frac{21}{22} \cdot \frac{21}{20} = \frac{441}{440} \approx 1,0023$$

$$dd = \frac{21}{22} \cdot \frac{21}{22} = \frac{441}{484} \approx 0,9112$$

Calculating the basic values for this two period game:

Mean gross return: $\varepsilon(r_A) = p^2 u^2 + 2pqud + q^2 d^2 - 1 = \frac{1}{4} \cdot \frac{441}{400} + \frac{1}{2} \cdot \frac{441}{440} + \frac{1}{4} \cdot \frac{441}{484} \approx$ 1,004551

Standard deviation:

$$\sigma(r_A) = \sqrt{p^2\big(u^2 - \varepsilon(r_A)\big)^2 + 2pq\big(ud - \varepsilon(r_A)\big)^2 + q^2\big(d^2 - \varepsilon(r_A)\big)^2} \approx 0{,}067688$$

Geometric return: $\qquad\qquad G = \left(\dfrac{21}{20}\right)^{\frac{1}{2}\cdot 2}\left(\dfrac{21}{22}\right)^{\frac{1}{2}\cdot 2} - 1 \approx 0{,}002273$

Again we expand our game. This time, we will look at a three period game ($N = 3$). We flip the coin three times, use Kelly in deciding the fraction of our wealth to bet and update our wealth for each of our trials.

This time we end up with these ending wealth'

Figure 3.5

The results came by multiplying the relevant factors with each other:

$$uuu = \frac{21}{20}\cdot\frac{21}{20}\cdot\frac{21}{20} = \frac{9261}{8000} \approx 1{,}1576$$

$$uud, udu, duu = \frac{21}{22} \cdot \frac{21}{20} \cdot \frac{21}{20} = \frac{9261}{8800} \approx 1{,}0524$$

$$udd, dud, ddu = \frac{21}{22} \cdot \frac{21}{22} \cdot \frac{21}{20} = \frac{9261}{9680} \approx 0{,}9567$$

$$ddd = \frac{21}{22} \cdot \frac{21}{22} \cdot \frac{21}{22} = \frac{9261}{10648} \approx 0{,}8697$$

Calculating the basic values for this three period game:

Mean gross return: $\varepsilon(r_A) = p^3 u^3 + 3ppquud + 3pqqudd + q^3 d^3 - 1 = \frac{1}{8} \cdot \frac{9261}{8000} + \frac{3}{8} \cdot$

$\frac{9261}{8800} + \frac{3}{8} \cdot \frac{9261}{9680} + \frac{1}{8} \cdot \frac{9261}{10648} \approx 1{,}006834$

Standard deviation:

$\sigma(r_A) =$

$$\sqrt{p^3\big(u^3 - \varepsilon(r_A)\big)^2 + 3ppq\big(uud - \varepsilon(r_A)\big)^2 + 3pqq\big(udd - \varepsilon(r_A)\big)^2 + q^3\big(d^3 - \varepsilon(r_A)\big)^2} \approx$$

$0{,}083136$

Geometric return: $G = \left(\frac{21}{20}\right)^{\frac{1}{2} \cdot 3} \left(\frac{21}{22}\right)^{\frac{1}{2} \cdot 3} - 1 \approx 0{,}003411$

By now we can see that we can generalize the formulas for mean return, standard deviation and the geometric return and calculate these for a N period game.

Let's now assume that we repeat the bet N times, this can be generalized:

Mean gross return: $\varepsilon(r_A) = p^N u^N + p^{N-1} q u^{N-1} d + \cdots + pq^{N-1} ud^{N-1} +$

$q^N d^N$

Standard deviation:

$\sigma(r_A) =$

$$\sqrt{p^N\big(u^N - \varepsilon(r_A)\big)^2 + p^{N-1}q\big(u^{N-1}d - \varepsilon(r_A)\big)^2 + \cdots + pq^{N-1}\big(ud^{N-1} - \varepsilon(r_A)\big)^2 q^N\big(d^N - \varepsilon(r_A)\big)^2}$$

Geometric return: $G = \left(\frac{21}{20}\right)^{\frac{1}{2} \cdot N} \left(\frac{21}{22}\right)^{\frac{1}{2} \cdot N} - 1$

3.2 Basic coin flipping proposition

First we want to take a look at the basic coin flipping example that we introduced earlier. We modify it slightly by changing the win amount to: $B = \frac{11}{10}$. The expected value of this bet can also be calculated:

$$E[r] = p_w r_w + p_l r_l = \frac{1}{2} \cdot \frac{11}{10} + \frac{1}{2} \cdot (-1) = \frac{1}{20} = 0{,}05$$

The problem lies in how much we should care to wager. If we want to grow our wealth optimally, we should use Kelly, and we shall use this strategy as a starting point when doing our sensitivity analysis.

With the inputs presented to us, we calculate our Kelly fraction:

$$f^* = \frac{Bp - q}{B} = \frac{\frac{11}{10} \cdot \frac{1}{2} - \frac{1}{2}}{\frac{11}{10}} = \frac{1}{22}$$

When we are betting Kelly, the amount of wealth we have is irrelevant as the bet size is a fraction of our wealth, not an absolute amount. And since the stipulations of the bet never changes, we will wager the same fraction of our bankroll on all of the bets that we make.

In our simulation, we will start out with a wealth $W_0 = 100$, and face 3000 subsequent bets with the exact same stipulations in every instance and the wealth will be updated after each bet made. We will repeat this scenario 1000 times and record the ending wealth for every bettor.

The method we are using is called Monte Carlo simulation. Monte Carlo means using random numbers in scientific computing. More precisely, it means using random numbers as a tool to compute something that is not random. For example, let X be a random variable and write its expected value as $A = E[X]$. If we can generate $X_1, \dots, X_n$, n independent random variables with the same distribution, then we can make the approximation:

$$A \approx \hat{A}_n = \frac{1}{n} \sum_{k=1}^{n} X_k$$

The *strong law of large numbers* states that as $n \to \infty$, $\hat{A}_n \to A$.

I will show how I update the bankroll of the bettor after every bet by using an example:

We start out with wealth $W_0 = 100$, and we calculate how much we intend to bet, which depends on which investor we are looking at. For this example, we will use the investor that bets full Kelly. This means that he will invest $f_0 = W_0 \cdot f^* = 100 \cdot \frac{1}{22} = \frac{100}{22} \approx 4{,}545$ in the first period.

For the bettors that we simulate having an error in their estimations, we will tweak our formula for calculating the amount that will be wagered on a given bet. We assume that the "true" win probability in this event is exactly 50%, but what if it's not possible to know this beforehand and that we have to try to estimate this probability somehow. And when we are estimating something, to some extent there will be some uncertainty added to the values we find. To account for this, we will create a new formula for the win probability p_i, were we add an element of uncertainty:

$$p_i = p(1 + kz_i)$$

where z_i is a standard normal distribution $\left(z_i \sim N(0,1)\right)$ and k is factor for how much of an error the investors are making, where we will be using values of 0.01, 0.02, 0.03 and 0.04 for k. The way that we produce these random values in excel is simply by using the two functions together: NORMSINV(RAND())

RAND() produces a random number between 0 and 1. NORMSINV now takes this random number between 0 and 1, and tells you how many standard deviations you need to go above or below the mean for a cumulative Gaussian distribution to contain that fraction of the entire population. This will produce an approximation of a normal distribution of random numbers.

This new value p_i will be used to calculate the Kelly fraction that will be bet by this investor, and it will be recalculated after each subsequent bet, just as for the other Kelly bettors.

Then we use the function RAND() in excel to create a random number between 0 and 1, and by using an IF statement in excel, we say that if the random number is smaller than 0,5, it becomes a loss, and if the number is larger than 0,5, we have a win. To update the bankroll for the next period, we simply multiply the invested amount with the return we get. Let us assume that we win in this example:

$$W_1 = W_0 + (f_0 \cdot B) = 100 + \left(\frac{100}{22} \cdot \frac{11}{10}\right) = 100 + 5 = 105$$

Now we calculate the new bet size:

$$f_1 = W_1 \cdot f^* = 105 \cdot \frac{1}{22} = \frac{105}{22} \approx 4,773$$

And again we calculate our new wealth:

$$W_2 = W_1 + (f_1 \cdot B) = 105 + \left(\frac{105}{22} \cdot \frac{11}{10}\right) = 105 + 5,25 = 110,25$$

We continue this until $n = 3000$, where we end up with the ending wealth, W_E, for each of the bettors, and we record this over in another sheet and then update the first one with new random numbers. Doing this a 1000 times would be quite tedious if done by hand, and I have programmed a simple macro that will do this automatically. You can find the code for this in the appendix A.

As for the bettors, we will look at those that use full Kelly, half Kelly, fixed sized bets and 4 different bettors that want to use full Kelly, but have errors in their estimates. In this instance, full Kelly means betting the whole fraction of the optimal Kelly fraction, f^*.

$$f_{Full\ Kelly} = cf^*, \text{where } c = 1$$

$$f_{Full\ Kelly} = f^*$$

Half Kelly means we are betting half the fraction of the optimal Kelly fraction, f^*.

$$f_{Half\ Kelly} = cf^* \text{ and } c = \frac{1}{2}.$$

$$f_{Half\ Kelly} = \frac{1}{2}f^*$$

Fixed size bets here means a bettor that determines an absolute amount to bet and bets this amount on every bet, regardless of how his wealth changes. For this instance, we will make him use:

$$f_{Fixed\ size\ bets} = f^* \cdot W_0 = \frac{1}{22} \cdot 100 = \frac{100}{22} \approx 4,545.$$

So this bettor will continue to bet an absolute amount of 4,545, despite any change to his wealth.

3.2.1 Results: Coin flipping

Ending wealth, W_E	Kelly Full	Kelly Half	Flat Betting	Kelly + ($k = 0.01$)	Kelly + ($k = 0.02$)	Kelly + ($k = 0.03$)	Kelly + ($k = 0.04$)
Average	126 539	3 425	718	185 450	275 669	354 742	372 420
Median	3 156	1 317	767	1 722	448	76	8
Std. Dev.	1 151 900	7 623	368	2 656 313	5 216 591	7 428 619	8 826 312

Table 3.1

The results from running this simulation are interesting. The thing that stands out the most is that the average ending wealth (W_E) for our investors gets larger as the error in our estimations increases. At first glance this might lead us to believe that it would be a good thing to have some uncertainty, which seems to increase our return. But our intuition knows better. It does not make sense that a strategy that is flawed in it's calculation of the inputs should outperform one that has complete information when they essentially are doing the same thing barring that they have different access to information. The reason for the bloated values for the average ending wealth might come from the aggressive nature of the investment strategy that is the Kelly criterion. Those select few samples we encounter that has an exceptionally good run of bets, and at the same time the errors these bettors make are on the positive side when we win and negative when we lose, that is, the correlation between winning and the errors are positive. This will make some of the ending wealth's seen by the bettors become exceptionally large. The largest W_E for the Kelly bettor with $k = 0.04$ is at 275 220 939, which exceeds the average by a large margin, and will in turn help bloat the results quite a lot. This problem arises because there are no ceilings as to how high the wealth can go, but it cannot shrink below 0, and thus the results will be skewed to the right.

If we on the other hand look at the median for W_E, the problems from having these estimation errors becomes more obvious. As the errors in the estimation of p_i increases, the median value decreases drastically and with $k = 0.04$, the median is all the way down to 8, which is far below our starting wealth $W_0 = 100$. This is exactly the problem with the estimation errors and the use of full Kelly; it is far too easy to start overbetting and thus increase our volatility and reduce our growth at the same time.

Figure 3.6

By summing up the times that W_E ended up above our W_0, we find the results seen in the figure above. We see that betting full Kelly leaves us profitable over the sample of 3000 bets more often than even flat betting a fixed amount every time. The reason for this is that we will always bet a fractional amount of our wealth with Kelly and in essence can never go broke using it (although we could reach a value of our wealth that is so small that we can't pony up enough cash for a minimum investment amount), but the fixed amount bettor will sometimes reach a low point where he has to bet his whole wealth and lose, which in turn would make it impossible for him to recoup his losses given that he is now broke and can't continue betting.

Again we see problems with the bettors that have estimation errors added to their inputs of the win probability. The likelihood of ending up with a profit after 3000 bets sinks drastically as the error margin increases. When the $k = 0.03 \: or \: 0.04$, it's more likely to end up with a loss after the all the bets have been placed, that is $W_E < W_0$.

Percentile (Coin flip)	Full-Kelly	Half-Kelly	Flat betting	Kelly (k=0.01)	Kelly (k=0.02)	Kelly (k=0.03)	Kelly (k=0.04)
1 %	6	58	0	3	0	0	0
5 %	43	154	0	18	2	0	0
25 %	568	558	548	269	44	5	0
50 %	3 156	1 317	767	1 722	448	76	8
75 %	21 229	3 417	968	13 545	4 527	1 108	172
95 %	306 141	12 985	1 245	210 841	163 225	117 684	39 367
99 %	1 554 415	29 240	1 407	2 511 633	3 719 347	3 610 566	2 203 991

Table 3.2

Looking at the percentile table for the coin flip simulation, we find the values for which the corresponding percentile amounts of the population fall beneath in our simulation. This gives us a clearer picture as to just how the ending wealth of the different investors are spread around. We can clearly see that the full-Kelly investor has a lot more risk involved than the half-Kelly investor and the flat bettor.

We also note the larger the spread becomes as k increases for the investors with errors in there estimations. For $k = 0.04$, the investor essentially goes broke more than 25% of the time. Another interesting result from this analysis is the fact that full-Kelly outperforms all of the other investment strategies up to the 99[th] percentile (barring the half-Kelly strategy). Both the half-Kelly and flat betting strategies are naturally outperformed by the full-Kelly strategy in the long run, but it is now shown that the error prone investors will suffer in all facets of the investment horizon because of their faults in estimation.

For further analysis of these findings, we can compare them with what we have found in part 3.1, our analytical calculations. Betting full Kelly ends up with a median wealth of 2869, that is after 3000 subsequent bets, so $n = 3000$. When we calculated the geometric return for Kelly, we found the generalized formula to be:

$$G = \left(\frac{21}{20}\right)^{\frac{1}{2} \cdot N} \left(\frac{21}{22}\right)^{\frac{1}{2} \cdot N}$$

We can also add in the starting wealth, $W_0 = 100$:

$$W_0 \cdot G = W_0 \cdot \left[\left(\frac{21}{20}\right)^{\frac{1}{2} \cdot N} \left(\frac{21}{22}\right)^{\frac{1}{2} \cdot N}\right] = 100 \cdot \left[\left(\frac{21}{20}\right)^{\frac{1}{2} \cdot 3000} \left(\frac{21}{22}\right)^{\frac{1}{2} \cdot 3000}\right] = 3012,1$$

This shows that we get an ending wealth when we calculate it analytically that is just a little bigger than what we got from our simulations.

To check the half Kelly bettor, we must find new values for the up and down factors:

$$u = 1 + \frac{1}{2}fB = \frac{41}{40} \ and \ d = 1 - \frac{1}{2}f = \frac{43}{44}$$

$$W_0 \cdot G = W_0 \cdot \left[\left(\frac{41}{40}\right)^{\frac{1}{2} \cdot N} \left(\frac{43}{44}\right)^{\frac{1}{2} \cdot N}\right] = 100 \cdot \left[\left(\frac{41}{40}\right)^{\frac{1}{2} \cdot 3000} \left(\frac{43}{44}\right)^{\frac{1}{2} \cdot 3000}\right] = 1286,7$$

Again, we find a value that is close to what we have found in our simulations, where the median W_E was found to be 1255.

Last we have the fixed size bettor. He will bet the same amount on every opportunity, regardless of his wealth. His bet size will be:

$$W_0 \cdot \frac{1}{22} = \frac{100}{22} = 4,545$$

Since the expected value from this bet will be the same on every bet, and the bet size also will never change, this is an additive function. We just multiply our expected gain after each bet with the times we are to bet:

$$E[r] = \ p_w r_w + \ p_l r_l \ = \frac{1}{2} \cdot \frac{11}{10} + \frac{1}{2} \cdot (-1) = \frac{1}{20}$$

Now we multiply our expected value with the bet amount and the times we will repeat the wager:

$$E[r] \cdot W_0 \cdot \frac{1}{22} \cdot N = \frac{1}{20} \cdot 100 \cdot \frac{1}{22} \cdot 3000 = \frac{300000}{440} = 681,8$$

The result with end up with here is the expected net return from investing in the market. We also need to add the starting wealth of W_0 to find the ending wealth W_E:

$$W_E = E[r] \cdot W_0 \cdot \frac{1}{22} \cdot N + W_0 = 681,8 + 100 = 781,8$$

We find that the expected ending wealth W_E to be larger than the average ending wealth found in the simulations. Given the low level of variance inherit in this strategy, combined with the large number of trials we have done, we would expect to arrive somewhat closer to the actual amount. The reason that we do not is because of the fact that this bettor will sometimes go broke (121 out of a 1000, to be exact), and thus not be able to continue betting anymore. When he does, he misses out on being able to continue taking advantage of this betting opportunity, and thus the cumulative expected value that we find here will be lower than the one we calculate analytically, which does not take into account the times he goes broke and can't continue to wager.

The conclusion we can draw from this is that our simulations seems to have been done correctly and have given us results that are in line with what we would expect to see. There is however no easy way to calculate the ending wealths for the bettors that use the Kelly criterion and have estimation errors in their calculations. These errors are supposed to be random, and if we were to use the average of the supposed estimations, we would arrive at the same answer as for full Kelly. This is clearly not the case, as the simulations show that we get values that differ vastly from the full Kelly bettor. The results however, are in line with what we would expect given our prediction in chapter 2, and specifically figure 2.2, where we show that if we increase the betting amount above the optimal Kelly fraction, we both increase the risk and reduce the growth of our wealth, which is not what an investor would want.

Since the error term is normalized to have a mean of 0 and a standard deviation of 1, we know that the error terms thus means that we have a standard deviation from our correct estimate of about $k \cdot 100\%$. For $k = 0,01$, this means that simply making a guesstimate of a coin flip, and being between 49% and 51% (when 50% is the correct rate) $\frac{2}{3}$ of the time, would result in worse returns and greater volatility than a bettor that gets it correct every time.

For $k = 0,04$, this window widens to 46% and 54%, which creates abysmal results. but think about this for a minute. If we go back to the cycling race example between Thor Hushovd and Cadel Evans again for a minute, think about how hard it must be to get an exact estimate on their actual win chances against each other. Being off by 4% in your estimates is not something that would be out of the ordinary, and you might even miss by more than this as well, making it even worse.

We shall continue to look at this performance of Kelly under the faults of estimation errors, and will take a look at the stock market, which has even more complex inputs to estimate.

3.3 The stock market

Since we generally will not be faced with such a great proposition as described in the coin tossing example, it seems reasonable to try to model a more familiar and realistic market where we can invest our money. Thus we want to try to make a model of stock prices.

The stock market is usually modelled in continuous time and the stock prices are considered to be continuous variables as well. These prices will move up and down in a random and unsystematic fashion, which makes us base our model on a stochastic process. A stochastic process that fits this profile pretty good is the Wiener process (also called standardized Brownian motion) which has a mean of 0 and a variance rate of 1.0 per year.

A variable k follows a Wiener process if it has the following two properties:

Property 1: The change Δk during a small period of time Δt is $\Delta k = \varepsilon\sqrt{\Delta t}$, where ε has a standard normal distribution.

Property 2: The values of Δk for any two different short intervals of time, Δt, are independent.

Wiener process

Figure 3.7

The figure above illustrates how a Wiener process behaves. The drift rate is here set to 0 and the volatility (yearly) is set to 40%. The process of whether the price will go up or down is randomized and if we were to recalculate this process, we would get a totally different path next time around. On average, the ending value will be 100 (as the drift is 0). However, if the stock market where to follow a Wiener process, no one would own stock considering that the expected return is zero.

What we want to use, is a generalized Wiener process (GWP). This process does not need the mean to be 0 and the variance be 1. By accounting for the fact that the stock market drifts upward in the long run, we can add a drift factor to the process. The GWP process can be written as such:

$$dx = adt + bdk$$

where *adt* is the drift term and *bdk* the noise term. The drift term decides how much we expect the process to increase in each time period dt and the noise term determines the fluctuations we will see around this expectation. This is illustrated in the figure below, where we have added a GWP with a drift of 10% (yearly) and variance of 40% (yearly). We also have added the drift without the noise added to it, to better illustrate that the GWP is moving around the drift term and not the starting value, which the standard Wiener process does.

Figure 3.8

When we use the Kelly criterion for investing in the stock market, we first notice that there are a couple of practical problems that will arise. The main difference between betting on a coin flip and investing in a stock market is the fact that the coin flip as an exact end time for the investment, while the investment in the stock market occurs in continuous time. This forces us to use continuous probability distributions when calculating our Kelly fraction for our investments.

Thus we have to use another formula for calculating the optimal betting fraction, since the formula that we have used so far only works when operating with finite time situations. From the article written by Jane Hung, we get the formula for the optimal betting fraction (Kelly) in continuous time:

$$f^* = \frac{\mu - r}{\sigma^2}$$

For this simulation, we will assume we have yearly rates of:

$$\mu = 10\%, \ r_F = 3\% \text{ and } \sigma = 40\%$$

which we convert to daily rates:

$$\mu_{Daily} = \frac{10\%}{252} \approx 0,0397\%, \ r_{F_{Daily}} = \frac{3\%}{252} \approx 0,012\% \text{ and } \sigma_{Daily} = \frac{40\%}{\sqrt{252}} \approx 2,5198\%$$

This leads to the optimal betting fraction:

$$f^* = \frac{\mu - r}{\sigma^2} = \frac{0,0397 - 0,012}{2,5198^2} = 43,75\%$$

Now that we know the optimal betting fraction, we need to simulate the actual stock market. We have already shown that we intend to simulate the daily drift of the stock market prices using the GWP, and the rates have been stated above here. The daily returns will be calculated:

$$r = N(10\%, 40\%) = \mu + \sigma z_i$$

where z_i is a standard normal distribution $\left(z_i \sim N(0,1)\right)$. These are also the logarithmic returns. When we want to calculate the ending wealth after a period of return, the formula we use is:

$$W_{t+1} = W_t \cdot e^{r_t}$$

We start out all of the investors with an initial wealth of 100. After each period, we update the wealth for the investor using this formula:

$$W_{t+1} = \left(W_t - (W_t \cdot f^*)\right) \cdot \left(1 + r_{F_{Daily}}\right) + (W_t \cdot f^*) \cdot e^{r_t}$$

What this formula does is first taking the amount that is not put into the market, which will stay in a bank account and earn the risk free rate $\left(r_{F_{Daily}}\right)$, and then we add the amount of the wealth we invest in the market and the return it gets from being invested. For the half-Kelly investor and flat rate investor, we do the same, but simply change the amount f^* that is invested in the stock market.

Now, we also introduce 4 other investors. All of them want to invest using the full-Kelly fraction. However, these investors have trouble predicting the exact daily drift that the stock market will assume. This is something that would be closer to the practical investors, as no one can accurately predict what the stock market will return in the future, but one can get pretty close. To add a degree of uncertainty to these investors, we have to create a new value for the daily drift:

$$\mu_i = \mu(1 + kz_i)$$

where z_i is a standard normal distribution $(z_i \sim N(0,1))$ and k is factor for how much of an error the investors are making, where we will be using values of 0.25, 0.50, 1.00 and 2.00 for k.

From this the optimal betting fraction can be calculated. Notice that this one has to be updated after every day. This is because of the error term that is random and will change every day as well. The formula for f^* is already known:

$$f^* = \frac{\mu_i - r}{\sigma^2}$$

This makes it so that we can update our formulas for the ending wealth after each period:

$$W_{t+1} = \left(W_t - \left(W_t \cdot \left(\frac{\mu_i - r}{\sigma^2}\right)\right)\right) \cdot \left(1 + r_{F_{Daily}}\right) + \left(W_t \cdot \left(\frac{\mu_i - r}{\sigma^2}\right)\right) \cdot e^{r_t}$$

Now that we have the formulas for updating the wealth for all of the investors we want to take a closer look at, we will use the same method as we have used for the coin flip simulation, the Monte Carlo method. We let 3000 days pass $(n = 3000)$, and record the ending wealth for every investor. We then update the random numbers, and record another set of ending wealths. This is done using a macro, see Appendix A.

3.3.1 Results: Stock market

Ending wealth, W_E	Kelly Full	Kelly Half	Flat Betting	Kelly $(k = 0.25)$	Kelly $(k = 0.50)$	Kelly $(k = 1.00)$	Kelly $(k = 2.00)$
Average	308	210	164	307	307	309	322
Median	262	203	165	251	233	169	53
Std. Dev.	205	65	17	216	257	441	2070

Table 3.3

The results we find from the stock market simulations shows us that having the errors in our estimations gives us roughly the same average ending wealth (W_E) for all the variations of Kelly betting, except for the fractional Kelly where we are betting half of the optimal Kelly fraction. We notice that the strength lies in the median, where betting full Kelly has the

highest median ending bankroll out of all the Kelly variations. Betting half Kelly gives us much lower W_E, but also significantly smaller standard deviation.

We see that the median for full Kelly ($k = 0.25$) is just a little smaller than the actual full Kelly ($k = 0.00$). The standard deviation is also somewhat bigger for the investor with the estimation errors, but not by a lot. These values are something that could be bearable, if we were able to continue estimating the daily returns within 0.25 standard deviations within the true returns. However, once the error term gets bigger, we start to see the problems arise again and the returns suffer and the standard deviation gets out of hand. This goes to show that figure 2.2 is correct. The errors lower growth and increase the volatility. We see that as k increases for the investors, W_E decreases and the standard deviation increases.


Figure 3.9

Full Kelly once again ends up above the starting wealth about 90% of the time, and half Kelly does so a little more, which should be obvious given the safer investment philosophy it adopts. Flat betting (or investing would probably be a more proper term) does better in this simulation. It is now the option that ends up with a positive ending wealth the most often out of all the investors. The reason for this might lie in the return nature of the stock market. In

the coin flip simulation, if you won your bet you get a return of B times your investment, or you lost your whole bet. When investing, you get an undetermined return on your investment. So even though the relative investment size for this fixed size bettor is larger than for the coin flip simulation, his returns are not polarized like they were in that example, here they are normally distributed and usually a lot smaller than the whole investment.

This frequency table shows us again how skewed the estimation error prone investors ending wealth's are. For the investors with $k = 0.25$ and $k = 0.50$, we see that the likelihood of ending up with less than the starting wealth is about the same as for full Kelly. Once we increase this to $k = 1.00$, this begins to drop sharply, and for $k = 2.00$, we will more often than not see ourselves ending up with a wealth lower than what we started with.

Percentile (Stock market)	Full-Kelly	Half-Kelly	Flat betting	Kelly (k=0.25)	Kelly (k=0.50)	Kelly (k=1.00)	Kelly (k=2.00)
1 %	66	102	73	62	49	16	1
5 %	94	122	113	88	72	34	3
25 %	168	162	185	161	138	85	16
50 %	245	196	229	241	230	173	56
75 %	384	245	284	375	374	333	189
95 %	678	326	353	694	791	999	1 122
99 %	1 178	429	415	1 270	1 540	2 739	4 446

Table 3.4

Unlike in the coin flip simulation, here the results are not as clear cut. We can see that full-Kelly outperforms the other strategies on the upside, and falls behind the safer strategies (half-Kelly and flat betting) on the downside.

However, we note that the difference between full-Kelly and Kelly ($k = 0.25$) is miniscule, and for all practical purposes essentially the same. The differences here can probably be attributed to randomness. It's worth noting that the investor with the error term added was almost always behind the full-Kelly investor, albeit with small values. As the error factor increases, we start to recognize the problems we found in the coin flip simulation. The median shrinks as the error increases, just as predicted from our theory.

This percentile table will show the risk that is being taken on by the investor when he decides to apply this investment strategy to his portfolio. It is a more intuitive presentation

rather than looking at the standard deviation of the ending wealth's, which is quite bloated for some of the investors given the widespread distribution values for the ending wealth.

So what does all this actually tell us? I doesn't take a genius to deduce that the Kelly criterion is a strategy that will increase the risk of your portfolio, and significantly so. However, what it does not tell us is that if you have a model for selecting stocks that will pick great on average, but has very volatile as far as accuracy goes, you will do very bad when utilizing the Kelly criterion. The volatility will increase drastically and the results will suffer from it. Let us look at an example of this:

If we say that we have a model that tries to predict the win probability of the coin flip proposition. It will half the time predict a win probability of 51% and the other half it will say the win probability is 49%. Now, we know that it for a fact is actually 50%, so it will half the time overestimate the odds, and the other half underestimate it. We can calculate the Kelly fraction we would bet should we observe these win probabilities:

$$f_O^* = \frac{Bp - q}{B} = \frac{\frac{11}{10} \cdot \frac{51}{100} - \frac{49}{100}}{\frac{11}{10}} = \frac{71}{1100} \approx 6{,}455\%$$

$$f_U^* = \frac{Bp - q}{B} = \frac{\frac{11}{10} \cdot \frac{49}{100} - \frac{51}{100}}{\frac{11}{10}} = \frac{29}{1100} \approx 2{,}636\%$$

Now, since the true probabilities is different than the one predict, we need to calculate the new growth rate as well:

$$G_O = (1 + f_O B)^p (1 - f_O)^q - 1 = \left(1 + \frac{71}{1100} \cdot \frac{11}{10}\right)^{\frac{1}{2}} \left(1 - \frac{71}{1100}\right)^{\frac{1}{2}} - 1 \approx 0{,}094\%$$

$$G_U = (1 + f_U B)^p (1 - f_U)^q - 1 = \left(1 + \frac{29}{1100} \cdot \frac{11}{10}\right)^{\frac{1}{2}} \left(1 - \frac{29}{1100}\right)^{\frac{1}{2}} - 1 \approx 0{,}094\%$$

Now, we can find the actual growth, and show how this model would have performed betting on this scenario:

$$G = G_O^{\frac{1}{2} \cdot 3000} G_U^{\frac{1}{2} \cdot 3000} \cdot W_0 = 1652{,}88$$

The results from this little digression should scare most people who would like to apply the Kelly criterion to their portfolio strategy. Simply by being off by 1% in your calculations of the win probabilities, you will cut your expected ending wealth in almost half. Not only that, but your volatility will also increase. If we increase the prediction discrepancy to 52% and 48%, we will find that:

$$G_O = G_U \approx 0,03345\%$$

$$G = G_O^{\frac{1}{2} \cdot 3000} G_U^{\frac{1}{2} \cdot 3000} \cdot W_0 = 272,73$$

We see that the expected ending wealth gets reduced dramatically once the estimation errors increase.

Now, I know that this calculation is quite simplified, and in the simulations, we have used a normal distribution to determine how much of an error we have made in our calculations, but the point still stands. If you are to use the Kelly criterion, you should be really confident in your estimations, and should devote much time and money to make sure that they are kept correct.

Also, this means that you should look for other ways to size your portfolio positions when you are dealing with investments that have a lot of uncertainty attached to it. If you are planning on investing in very volatile assets, like small cap stocks, using full Kelly should be avoided, or a fraction of it should be used instead.

Figure 3.10

As we can see from figure 4.1 above, the yellow lines indicate the growth and corresponding standard deviation that the bettor experience with the given errors. If we were to use half Kelly instead of full Kelly, we would end up on the purple lines instead. This means that the expected growth will be lowered for one of them, and increased for the other. The sum of this is a lower total expected growth:

$$G = G_{O\,Half\,Kelly}^{\frac{1}{2} \cdot 3000} G_{U\,Half\,Kelly}^{\frac{1}{2} \cdot 3000} \cdot W_0 = 1107,16$$

But we will also experience a much lower standard deviation when using a half Kelly instead of full Kelly.

However, if you have a way of quite accurately predict the win probability of an event, using Kelly makes more sense. An example of this might be betting on sports. Here you know the odds that you receive beforehand, and if you are to bet on let us say, the first half result of a football match. If we apply the efficient market hypothesis for the bigger markets that have high limits, we can use the full time results market, which often are large markets that will be pretty accurate about what the probability of an event occurring is (which is implicitly

stated from the odds you receive). Now, you can use this price to derive a probability for the half time result, which is a derivative of the full time result.

4. Conclusion

We have now shown how aggressive and volatile the application of the Kelly criterion as an investment strategy is. And not only are investors under the scrutiny of risk inherit in investing a large fraction of their wealth in an investment, but also from estimation error risk. Most of the investments made by investors include a lot of uncertainty that needs to be assessed correctly for one to make a sound investment. Using the Kelly criterion under such conditions can lead to disastrous results for your wealth, and Kelly should certainly not be used by novices, maybe not even professionals in markets that include a lot uncertainty.

The results found in this thesis should make people think twice about using the Kelly criterion, and if one decides that this is the portfolio strategy that is right for them, then using a fraction of the optimal Kelly fraction should be given a serious thought. A fractional Kelly strategy might result in a lower growth, but it also has a lot lower risk involved, and it will prevent you from making the mistakes that overbetting will produce where you receive the worst from both worlds, lower growth and larger risk. So the more uncertainty you are dealing with in your modeling, the lower your fraction of the Kelly criterion should be used.

References

Thorp, E. O. (2008) *Understanding the Kelly Criterion*, Wilmott, May and September, p. 3.

Chopra, V. K. and Ziemba, W. T. (1993) *The Effect of Errors in Means, Variances and Co-Variances on Optimal Portfolio Choice*, Journal of Portfolio Management, 19, p. 6-11.

J. R. Kelly, Jr. (1956) *A New Interpretation of Information Rate*, Bell System Technical Journal, 35, p. 917-926.

Hull, J. C. (2008) *Options, Futures and other Derivatives*, Pearson Prentice Hall, p. 261-268.

Hung, Jane. (2010) *Betting with the Kelly Criterion*, p. 3-4. Available at: <http://www.math.washington.edu/~morrow/336_10/papers/jane.pdf> [Downloaded 15. December 2012]

Kritzman, Mark P. (2000) *Puzzles of Finance: Six practical problems and their remarkable solutions*, p. 65-75

Ziemba, W. T. (2005) *The Symmetric Downside-Risk Sharpe Ratio and the Evaluation of Great Investors and Speculators*, Journal of Portfolio Management, 32(1), p. 108-122.

Goodman, Jonathan. (2006) *Principles of Scientific Computing*, p. 179-181. Available at: <http://www.cns.nyu.edu/~fan/docs/GoodmanBook.pdf> [Downloaded 15. December 2012]

Appendix A: Macro codes

I will include the simple macro codes that I have used in the simulations here:

```
Sub Coin flip()
'
' Coin flip Makro
'
   Let x = 0
   Do While x < 1000
   Sheets("Coin flip").Select
   Calculate
   Range("E4:L4").Select
   Selection.Copy
   Sheets("Monte carlo (coin)").Select
   Range("B1").Select
   Selection.End(xlDown).Select
   Selection.Offset(1, 0).Select
   Selection.PasteSpecial Paste:=xlPasteValues, Operation:=xlNone, SkipBlanks _
      :=False, Transpose:=False
   x = x + 1
   Loop
End Sub


Sub Stock market()
'
' Stock market Makro
'
   Let x = 0
   Do While x < 1000
   Sheets("Stock market").Select
   Calculate
   Range("F5:M5").Select
   Selection.Copy
   Sheets("Monte carlo (stock)").Select
   Range("B1").Select
   Selection.End(xlDown).Select
   Selection.Offset(1, 0).Select
   Selection.PasteSpecial Paste:=xlPasteValues, Operation:=xlNone, SkipBlanks _
      :=False, Transpose:=False
   x = x + 1
   Loop
End Sub
```