

Imperial College London
Department of Computing

Extending Kelly Staking Strategies to Peer-to-Peer Betting Exchanges

Edmund Noon

January 2014

Supervised by Dr. William J. Knottenbelt

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Computing of Imperial College London
and the Diploma of Imperial College London

Abstract

In his seminal paper J. L. Kelly Jr. linked information theory with a staking system for calculating the stake of a favourable bet. Kelly's system is elegant, and under certain conditions can be said to be optimal. In the fifty years since, his ideas have been extended to financial markets, with much success. There have been fewer such extensions to betting. In many countries gambling has been liberalised and the range of bets available has increased markedly. One recent innovation has made a particular impact, i.e. the introduction of peer-to-peer betting exchanges.

Betting exchanges allow their members to place bets with each other, behaving as bookmakers, something Kelly did not consider. This and the Internet have increased the number of bets available on a given event. These are often highly correlated. We extend Kelly's method to incorporate lay bets; and we do so analytically. We extend this numerically to incorporate highly correlated bets and commission, another feature unique to exchanges. We further extend this solution to include bets previously placed, making possible a profit before the match has started.

Whilst extending the optimal solution to cover new situations we make use of the increase in data available to examine the applicability of the original assumptions. We run Monte-Carlo simulations to show when these assumptions are likely to fail.

One of the features of betting exchanges is the ability to increase or reduce the bet size as the odds change over time. We examine the optimal bets to add to an existing portfolio of bets if the odds have changed. We briefly consider changes in the modelled probabilities.

Finally, we suggest a novel way of reversing Kelly's strategy to provide a market making strategy, even when there is some uncertainty in where the market prices should be. However we also show that if that uncertainty is too great such a strategy is unlikely to be profitable.

Acknowledgements

I would like to thank:-

- Dr. William Knottenbelt, my supervisor, for his enthusiastic support.
- Dr. Daniel Kuhn, for nudging me to find the analytical solution in Chapter 3.
- The Engineering and Physical Sciences Research Council (EPSRC), for funding the Doctoral Training Award (DTA), which in turn funded this research.
- My family, in particular Kevin Noon, for their support and proof reading.

Contents

1. Introduction	15
1.1. Motivation	16
1.2. Objectives	17
1.3. Contribution	18
1.4. Thesis Outline	18
1.5. Publications and Declaration of Originality	20
2. Background	21
2.1. Measuring Risk	21
2.1.1. Daniel Bernoulli	21
2.1.2. John L. Kelly Jr.	22
2.1.3. Harry M. Markowitz	25
2.1.4. Edward O. Thorp	26
2.1.5. Paul A. Samuelson	27
2.1.6. Thomas M. Cover	29
2.2. Football Modelling	33
2.2.1. Michael J. Moroney	33
2.2.2. Charles Reep	34
2.2.3. Michael J. Maher	34
2.3. Interior Point Methods	36
2.4. Gambling: Liberalisation, the Internet, and Betting Exchanges	37
2.4.1. United Kingdom Regulation	38
2.4.2. Costs and Benefits	41
2.4.3. Internet and Betting Exchanges	47
3. Extending Kelly’s Fractional Staking to Lay Bets	50
3.1. Lay Bets	50
3.2. Simple Example	54
3.2.1. Naive Kelly	54
3.2.2. Standard Kelly	55
3.2.3. Back and Lay Kelly	55
3.3. Football Example – Real Data	55
3.3.1. Data	55
3.3.2. Model	57

3.3.3. Results	58
4. Further Extending Kelly for Betting Exchanges	60
4.1. The general case	60
4.2. Simple Example (Part II)	63
4.3. Football Example	63
4.4. Ignoring Commission	64
4.5. Commission	65
4.6. Simple Example (Part III)	66
4.7. Constraint Choices	67
5. Analysing Kelly's Assumptions	72
5.1. Maximum Effective Bet Size	72
5.2. Probability of Large Losses	75
5.2.1. Kelly Cannot Go 'Bust', but	75
5.2.2. Probability of Substantial loss: Illustrative Examples	76
5.2.3. Probability of Substantial loss: Convergence	79
5.3. Minimum Bet Size	81
5.3.1. Minimum Bet Size: Primary	82
5.3.2. Minimum Bet Size: Secondary	84
5.3.3. Minimum Bet Size: Large Loss	85
5.4. Fractional Kelly	87
6. Dynamically Hedging a Portfolio of Bets	91
6.1. Increasing or Unwinding Previously Placed Bets	91
6.1.1. Proebsting's Paradox	91
6.1.2. Previous Bet: Increasing Risk	92
6.1.3. Previous Bet: Reducing Risk	93
6.2. Including Previous Bets	94
6.3. Moving Probabilities	95
6.4. The Uncertain Biased Coin	96
7. Market Making with Inverse Kelly	99
7.1. Market Making	99
7.2. Example	101
7.3. Small Scale Live Test	104
7.4. Market Making: Final Comments	106
8. Conclusion	110
8.1. Achievements	110
8.2. Should We Follow Kelly?	112
8.3. Future Work	113

A. Glossary	114
A.1. Guide to Betting Terminology	114
B. Common Football Bets	119
B.1. Bets Determined by Final Score	119
B.2. Bets not Determinable by Final Score	120
Bibliography	121

List of Tables

2.1.	The distribution of regulated gambling in the UK and USA for the year 2009. The UK casino and booking numbers exclude remote gambling. Definitions as in Figures 2.5 and 2.6.	42
2.2.	The size of regulated gambling in the UK and USA for the year 2009. Definitions as in Figures 2.5 and 2.6. The exchange rate used is the current exchange rate (£1=\$ 1.54) in case 1, and the middle of the 2009 range (£1=\$ 1.75) in case 2.	42
3.1.	The odds available to an imaginary tribe on the roll of a fair die.	54
3.2.	The staking required by the Kelly criterion and the capital multiplier for each outcome of the fair die. The wealth relative is the factor by which wealth changes for that outcome.	55
3.3.	The staking required by our updated Kelly with lay bets and the capital multiplier for each outcome of the fair die. The wealth relative is the factor by which wealth changes for that outcome.	56
3.4.	A comparison of the per bet growth rate achieved by each of the betting strategies for the fair die. The fraction staked is the total fraction of wealth required to be staked at each stage.	56
3.5.	The betting data on English Premiership matches from the 2009/10 season used in this chapter is taken from Betfair's API at approximately 10:00 on the day of the match. Included here are only those markets with a back and a lay price for each of the four markets listed.	57
3.6.	Comparison of growth rates when i) using back bets only, and ii) using back and lay bets. The data consists of 330 English Premiership matches from the 2009/10 season (described in more detail in Table 3.5). Test assumes an initial wealth of £500 and is allowed to grow. G is the (natural) logarithmic growth rate. In this example the bets are unrestricted in size (this will be considered later in Section 5.1). The probabilities are calculated by fitting the entire season to Maher's model.	58
4.1.	The odds available on the parity of the roll of a fair die to an imaginary tribe.	63

4.2.	The staking required by our generalised Kelly for two markets, and the capital multiplier for each outcome of the fair die. The wealth relative is the factor by which wealth changes for that outcome and includes the Odd/Even bets.	63
4.3.	Three results repeated from Table 3.6 with the new result of combining all three. As before, the data consists of 330 English Premiership matches from the 2009/10 season (see Table 3.5). Test assumes an initial wealth of £500 and is allowed to grow. In this example the bets are unrestricted in size; this will be considered later, in Section 5.1. The probabilities are calculated by fitting the entire season to Maher’s model.	64
4.4.	Poor bet caused by lack of commission adjustment	69
4.5.	Bets adjusted for commission	69
4.6.	A comparison of the results from Table 4.3 with the results using the commission adjusted objective function. In both cases, when calculating the wealth, commission is accounted for. Back and lay bets are placed. Combined 4 indicates all four markets combined together, combined 3 excludes correct score. Data as before.	70
4.7.	The exchange odds available on the parity of the roll of a fair die to an imaginary tribe.	70
4.8.	The staking required by the generalised Kelly for two markets and the capital multiplier for each outcome of the fair die, when commission is ignored.	70
4.9.	A comparison of the wealth relatives of bets calculated without taking account of commission, and the wealth relatives when commission is taken into account.	70
4.10.	The staking required by the generalised Kelly for two markets bets and the capital multiplier for each outcome of the fair die, when commission is included in the optimisation. The wealth relatives have had commission deducted.	71
5.1.	Results from using constrained sizes: otherwise as Table 4.6. The fraction staked at each stage is restricted to that available at the time. A substantial fall in terminal wealth and growth rates is seen.	72
5.2.	The probability, p_ψ , of wealth falling below a fraction, ψ , of initial wealth appears to be close to ψ , at least for the moderate values of p and o chosen here ($p = 0.5, o \times p = 1.2$).	79
5.3.	Table showing the decline of the incremental value of each iteration of Equation 5.2, where $p = \frac{1}{3}$, and $o = 3.5$. The iteration number is the last time that scale is exceeded so, for example, all iterations with $k > 1217$ are smaller than 10^{-7} . This suggests that for large enough k the series becomes roughly geometric, dividing by 10 every 500 iterations.	81

5.4.	The probability that a Kelly strategy is impacted by the minimum bet of £2, where $p = \frac{1}{3}$, and in the first case $o = 3.5$ and in the second $o = 3.15$. The second strategy cannot place a bet below £86.	88
5.5.	The logarithmic growth rates using fractional Kelly staking, where $p = \frac{1}{3}$ and $o = 3.5$	90
5.6.	The probability of being unable to continue with a fractional Kelly strategy for various fractions, where $p = \frac{1}{3}$ and $o = 3.5$	90
7.1.	Payout value for each of the possible match results for the four bets in the example. We have £100 lay bets of 3.9 for the draw, 4.0 for the away win and two for the home win of 2.04 and 1.99.	101
7.2.	Details of football matches which occurred during the market making test. Fallow matches were left untouched by the market making software.	105
7.3.	Details of the bets matched during the market making test. Only one match, only, had placed Correct Score bets.	105

List of Figures

2.1.	A graph comparing a Shannon constantly rebalanced portfolio (50% cash, 50% stock) and the underlying equity. The equity randomly doubles or halves in price each day.	29
2.2.	A graph comparing a Shannon constantly rebalanced portfolio (50% cash, 50% stock) and the underlying equity. The equity randomly doubles or halves in price each day. On this occasion the stock price fell significantly over the 500 days of trading, and yet the portfolio still did well.	30
2.3.	A graph showing what would have happened had the stock crashed suddenly in price from day 200 (perhaps Boo.com or Marconi). The stock falls quickly, with little volatility. So when compared to Figure 2.2 the portfolio does very badly (from 1 000 000 to 10 000).	30
2.4.	A graph comparing a Shannon constantly rebalanced portfolio with commission and the underlying equity. Commission charges of 0.5% are deducted from both purchases and sales. The dotted lines show a similar pair of portfolios but this time the underlying equity has a much lower volatility.	31
2.5.	A graph showing the size of the three main US market sectors. Casinos includes Commercial Casinos, Native American Casinos and Card Rooms, Lotteries includes Charitable Games and Bingo, and Bookmaking includes Pari-Mutuel Wagering. Data from Statistical Abstract [1].	42
2.6.	A graph showing the size of the four main market sectors. Casinos includes Casinos and Arcades, Lotteries includes Bingo, and Remote includes all forms of remote betting regulated from the UK, but not remote lottery transactions which are included in lotteries. Data from Gambling Commission [2], except lottery data which is from the National Lottery Commission [3] and [4].	43
2.7.	A graph showing the over-round for Premiership match odds for various bookmakers. Ladbrokes and William Hill are the main UK high street bookmakers. Interwetten is Austrian and not particularly active in the UK, whilst the rest are internet or telephone bookmakers.	47
3.1.	Four different cases for choosing bets to back and lay: no favourable bets; back only bets; lay only bets; and a combination of both back and lay bets.	53

4.1. Matrix: four horse race.	62
4.2. Excerpt of matrix for football bets. The first column is for cash, the next six are for Match Odds bets (3 back and 3 lay), the following four are for Under/Over 2.5 goals, followed by the first few back bets of Correct Score.	62
5.1. Match odds market. When wealth goes above £10 000 market size appears to be an impediment to growth.	73
5.2. Correct score market. When wealth goes above £50 000 market size appears to be an impediment to growth.	74
5.3. Over under 2.5 goals market has lower volumes than the previous two. When wealth goes above £2 000 market size appears to be an impediment to growth.	74
5.4. All 4 markets combined. When wealth goes above £30 000 market size appears to be an impediment to growth.	75
5.5. A graph of probability of substantial loss versus expected value of the bet being wagered for different bets with different probabilities of winning. In each case, as the expected value increases the probability of large loss decreases. When more favourable bets are won a greater number of losing bets will be required to bring about a substantial loss.	77
5.6. A graph of probability of substantial loss versus the underlying bet probability for different expected values. In each case, as the bet probability increases the probability of large loss decreases. When more favourable bets are won a greater number of losing bets will be required to bring about a substantial loss.	77
5.7. A graph of probability of substantial loss, p_ψ , versus both expected value and probability of winning the bet being wagered.	78
5.8. A graph of the number of iterations used to calculate the probability of substantial loss versus expected value and the probability of the bet being wagered. In each case, as the expected value increases the probability of large loss decreases. When more favourable bets are won a greater number of losing bets will be required to bring about a substantial loss.	78
5.9. A graph of the minimum wealth (£) needed to be able to place a bet of £2 at the given odds and probability. The minimum probability plotted is 0.001, and the minimum expected value is 1.001.	82
5.10. A graph of the probability that wealth will increase sufficiently so that the Kelly bet is no longer below the minimum bet before the strategy goes bust. Starting wealth is £100, and $p = \frac{1}{2}$. There are two possible definitions of bust here. One is that it is no longer possible to place any bet, i.e. wealth has fallen below the minimum bet, and the other is that the required bet is less than half the minimum bet.	83

5.11. A graph of the probability that wealth will increase sufficiently so that the Kelly bet is no longer below the minimum bet before the strategy goes bust. When starting wealth is larger than £200, and for most values of p this value is reasonably stable.	84
5.12. A graph of the ratio of two growth rates in one case the smaller bet is ignored, in the other it is included. The ratio is calculated for different bets of differing probabilities, in each case the wealth is £10 000, the primary bet size is £10, and the smaller bet size is £1. A few of the points on the graph are not possible (when the sum of the probabilities is greater than one and have been filled using the nearest value.	85
5.13. A graph of the ratio of two growth rates and standard deviations for different sizes of primary bet. The growth ratio is the factor by which growth is lower when the secondary (£1) bet is excluded. The risk impact is the factor by which the standard deviation of return increases. In each case the wealth is £10 000, the primary probability is 0.1, the secondary bet probability is 0.01 and the smaller bet size is £1.	86
5.14. A graph of probability that wealth falls so much that the required bet is below the minimum bet size by starting wealth. The minimum bet is £2, $p = \frac{1}{3}$ and odds $o = 3.5$	87
5.15. A graph of probability that wealth falls so much that the required bet is below the minimum bet size by starting wealth. The minimum bet is £2, $p = \frac{1}{3}$. Note that the EV scale is reversed. Two bets, with expected value 1.1, have too little starting wealth and have been plotted with probability 1.	88
5.16. A graph of probability that wealth falls so much that the required bet is below the minimum bet size by starting wealth. The minimum bet is £2, the expected value is $\frac{3.5}{3}$, which is the same as used in Figure 5.14. Note that the p scale is reversed. If the initial wealth is too low to place any bets, given the minimum bet size, then a value of 1 is used for the probability.	89
6.1. A graph showing the changing match odds for the Premiership match between Manchester United and Chelsea on 3 rd April 2010. Manchester United, the home team, is the clear favourite, with odds shown in red and blue, but then United's odds increase suddenly on Tuesday evening. At this moment Wayne Rooney, one of Manchester United's star players, limps off the pitch during a match against Bayern Munich and there is concern that he might not be able to play against Chelsea. Later in the week it is announced that his injury is serious enough to cause him to miss this match. Data taken from Betfair API.	96

7.1.	Initial prices in football market making example. Prices on the left (in blue) are those that the market can back (i.e. we would lay) and those on the right (in pink) are those that the market can lay. The yellow dotted line indicates fair value given by $p = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$, and $f = 0.02$. All prices are rounded to Betfair valid prices.	102
7.2.	Prices in football market making example after first bet is matched. The red square indicates that the customer has backed the draw (so we have laid it). Prices have now adjusted. The yellow dotted line indicates fair value given by $p = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$, and $f = 0.02$. All prices are rounded to Betfair valid prices.	102
7.3.	Prices in football market making example after second bet is matched. The second red square indicates that the customer has backed the away win (so we have laid it). Prices have now adjusted. The yellow dotted line indicates fair value given by $p = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$, and $f = 0.02$. All prices are rounded to Betfair valid prices.	102
7.4.	Prices in football market making example after third bet is matched. The third red square indicates that the customer has backed the home win (so we have laid it). Prices have now adjusted. The yellow dotted line indicates fair value given by $p = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$, and $f = 0.02$. All prices are rounded to Betfair valid prices.	103
7.5.	Prices in football market making example after fourth bet is matched. The fourth red square indicates that the customer has backed the home win for a second time (so we have again laid it). Prices have now adjusted. The yellow dotted line indicates fair value given by $p = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$, and $f = 0.02$. All prices are rounded to Betfair valid prices.	103
7.6.	Time taken before the matched bet was matched, compared to expected value of trade, measured using trade weighted average price.	107
7.7.	Length of time between matching and start of match, compared to expected value of trade measured using traded weighted average price. . .	108

1. Introduction

Gambling appears to have been regulated throughout its long history [5]. There are laws written before 2000 BC in hieroglyphics restricting gambling [6], and detailed accounts from Ancient Rome of their changing relationship with gaming and bookmaking [5]. Across many countries and cultures there appears to have been some practice similar to gambling [7]. In its most general form gambling can be thought of as the agreement to transfer assets, usually money, based upon the outcome of some as yet unobserved event. Notwithstanding that, different traditions have produced different terminology. The terminology used in this thesis is that prevalent in the United Kingdom.

Betting exchanges are a peer-to-peer betting platform. For every bet that something will happen, there is a bet that it will not. In the UK we call the bet that it will happen a back bet, and the bet that it will not a lay bet. Betting exchanges lend themselves to descriptions drawn from financial markets. Having placed a back bet, a person placing a lay bet later might be said to be unwinding his position, and back and lay become synonymous with bid and ask or buy and sell. Those unfamiliar with gambling terminology or familiar with that of a different country may wish to examine the glossary in Appendix A.1.

The general form of gambling defined above covers a much wider range of activities than those commonly covered by the term. In his letter Daniel Bernoulli [8] moves seamlessly between lottery tickets, dice and shipping insurance. Most forms of insurance are economically identical to betting; it is the regulatory environment which, quite properly, is different. Of course the motivation of an individual buying car insurance is different from that of someone betting on the winner of the Grand National. Also different in motivation and regulation, but otherwise similar, is betting on snow on Christmas day, and weather insurance/derivatives. Many of the techniques in this thesis apply to a wide range of insurance and derivative contracts and other investments. In many cases the purchaser of insurance will not need to calculate the optimal size of the bet – it is the same as the value of her car, but in the case of the underwriter some extension of this thesis might be of interest.

The most common form of betting in the UK is fixed odds betting. In this form of betting the ratio of money transferred on a win or loss is fixed in advance. This is the type of betting considered by Kelly [9] and it will be our main consideration.

1.1. Motivation

Kelly developed a staking algorithm to calculate the amount of wealth which should be staked on a favourable bet. Subsequently Breiman [10] showed that this system is optimal in more than one sense. It will eventually outperform any essentially different system and will minimise the number of bets required to reach some predetermined amount of wealth.

Since 2000 the introduction of betting exchanges to some countries has allowed consumers to place lay bets (see Section 3.1) in addition to the back bets Kelly considered. In its current form Kelly's methodology does not address these bets.

In recent years the range of bets available on any given event has increased markedly. Regulatory changes might account for some of the increase: in some countries, including the UK, there has been a decrease in regulation; state sponsored lotteries have allowed gambling to become more acceptable; betting exchanges have increased competition. Technological changes have also been significant. The Internet has made it easier to compare odds offered by competing bookmakers; it has also allowed companies to offer their services internationally (sometimes flouting local regulations), offering products common in other countries. Computers make it easier for a bookmaker to quote prices in an increased number of markets.

Thorp [11] and others have considered multiple favourable bets when those bets are uncorrelated, but that cannot be applied here. Breiman [10] did consider a general form of the problem which would have included this case but claimed a general solution was not possible. Since then, interior point optimisation methods have increased the range of problems for which numerical solutions have become available.

Bookmakers (in the UK) quote odds. A *better* places a stake on an outcome. If that outcome loses then the better loses the stake; otherwise, she wins her stake multiplied by the odds. The bookmakers hope to produce a profit by offering lower odds than those they believe to be fair. Betting exchanges work differently.

Betting exchanges allow individuals to bet with each other. In most cases the exchange takes no risk on the events, it is the exchange's customers who are taking the risk with each other. To pay their costs and make a profit the exchange charges commission. Betfair, currently the largest exchange in Europe, charges the winner of the bet the commission and the loser nothing. This is common to many of the other leading exchanges available in the UK. It might seem that a reduction in the odds by the appropriate commission would solve this problem. Unfortunately in many cases this does not work. One of the features of Kelly's system is that in addition to backing outcomes which seem favourable, others are also backed. The commission is on the net winning amount, and therefore needs to be included in the calculation.

Kelly's claim of asymptotic optimality required two assumptions; wealth is infinitesimally divisible, and there is no maximum bet size. It is possible (likely – as it is an asymptotic result) that there will be a significant run of bad luck which reduces the available capital. Because of the infinitesimal divisibility of capital Kelly never goes bankrupt. He is always able to place the next bet even if it is a fraction of a cent. And because this strategy has exponential growth then eventually wealth becomes large again. When modifying this strategy for financial products it is possible to assume a large starting capital, in which case the probability of bankruptcy is small. Similarly the financial markets are deep so the placing of large bets/investments is possible. The depth available on betting exchanges is not as large, and the minimum bet is typically £2 (which is larger than the smallest transaction possible in some financial markets). It appears that the two assumptions might be mutually exclusive.

By allowing both back and lay bets, betting exchanges offer the enticing possibility of profit without event risk. It is possible to take bets off before settlement, which might be before the event being considered has even started. Having placed a favourable back bet, it is likely that the market prices will change. If others also consider this bet to be favourable it is possible that the market price will move in a direction favourable to the better, i.e. lower. If the price moved too far the better could place a lay bet which would also be favourable and would *unwind* the original bet. That is profit now with no risk on the event. This is obviously an ideal scenario. One imagines an unwind should be considered at fair value, but perhaps it should be considered even earlier?

1.2. Objectives

The key objectives are:-

- To extend Kelly's fractional staking to include lay bets.
- To solve Breiman's generalisation of Kelly numerically, and to provide examples of this method in practice.
- To examine the problems caused by commission charged on winning bets by exchanges and to improve on the two trivial approximations of either ignoring commission or adjusting the odds by the commission.
- To use the large amount of data made available by betting exchanges to examine the practical implications of Kelly's assumptions, in particular what is the maximum effective wealth, and how does the minimum stake interfere with optimal growth?
- To adjust the new staking methods to take account of bets already placed.

1.3. Contribution

This thesis shows that Breiman's generalisation of Kelly's fractional staking should be considered solved, at least numerically. Breiman [10] states that finding a general solution is difficult. We show that by using interior point optimisation it is computationally straightforward, taking milliseconds on a standard desktop computer when combining four overlapping markets. For one specific case we solve it explicitly. This specific case is likely to be the most common extension that users of Kelly staking are likely to use as it is back and lay bets on a single (otherwise) disjoint market, which is the case for horse races, football matches, winner of a tennis tournament and many other common bets.

We further extend our numerical solution to include an approximation of the effects of commission. This approximation may be made as close as needed to the actual effect, with calculations taking milliseconds on a standard (circa-2009) desktop PC when calculations are within £0.05 on bets of the order of hundreds of pounds.

We developed a formula for the probability of a substantial loss. Using this with the minimum bet size we measured the probability of a strategy failing (because the next bet is below the minimum size). The Kelly Criterion promises exponential growth. In all markets eventually wealth will be much greater than the market can bear. We examine this size for bets on English Premiership matches. When considered alongside the minimum bet research we show that there is a relatively narrow range of wealth where a Kelly strategy will be unencumbered. We show that for all but the smallest wealth a fractional Kelly strategy will have a lower probability of falling below the minimum bet size.

We present a novel method of using the above methods to make markets. Rather than take odds from a website and then decide how much to bet, as above, this method places bets into a (perhaps empty) market and adjusts the prices as trades happen. It might be possible to make money even if the initial model is not accurate.

1.4. Thesis Outline

The remaining chapters are organised as follows.

Chapter 2 describes the background to the Kelly Criterion, as gamblers call it, also known as Maximum Expected Log (MEL) or Growth Optimal Portfolios (GOP) by economists and investment analysts respectively. This covers the controversy of how generally applicable this idea is. This thesis relies on Interior Point optimisation to solve the general problem and so this chapter also contains a brief summary of this work. As many of the examples rely on football models these are

similarly covered. Whilst these ideas apply to investments, insurance and other decisions most of this thesis uses betting examples and betting exchanges in particular. A discussion of the UK regulation which led to the current situation is included. This contains a few criticisms of the current regulations, and possible problems with betting exchanges.

Chapter 3 extends Kelly's original algorithm to include lay bets. This is demonstrated first with a simple example and then using football prices taken from the Betfair API.

Chapter 4 considers situations where there are numerous bets available all dependent on one game, such as the different bets typically available on a football match. This is the generalisation of Kelly described, but not solved by Breiman. These are solved numerically and demonstrated with a simple example and an example using real football betting prices. Betting exchanges do not charge an over-round as do bookmakers. They charge commission. Next this commission is included in the numerical solution.

Chapter 5 analyses some of the assumptions made in proving that MEL is optimal. In particular this strategy is asymptotically optimal, and yet as wealth increases it is not possible to place bets large enough to continue that growth. The second key assumption is that no matter how small wealth falls it is always possible to place a bet. Betfair has a minimum bet size of £2. How high is the probability that wealth will fall so low that a bet may not be placed?

Chapter 6 considers dynamic trading of a betting portfolio. If, having backed a runner in a race the odds increase should the size of the bet be increased, and by how much? If, on the other hand, the odds fall, should the bet be unwound? What are the implications for growth. What about if the probabilities also move? These are all answered.

Chapter 7 suggests an unusual application of Kelly to market making, that is placing bets in a market with no pre-existing prices, and then changing these prices as things trade. It is shown that it might be possible to make money using this method even if the model is not perfect, but it has to be reasonably close. One of the advantages of investigating betting exchanges rather than stock exchanges is that it is possible to test some ideas with a few pounds. The market making application was tested.

Chapter 8 concludes the thesis with a summary of the key results, and a discussion of further applications and research.

Appendix A has a guide to betting terminology and the glossary.

Appendix B describes some of the common football bets available on betting exchanges

aimed at UK customers.

1.5. Publications and Declaration of Originality

I declare that this thesis is my sole work. Throughout this thesis, other peoples research and results have been used. This use has been fully cited and referenced in the bibliography.

The following publications arose from the work carried out during my PhD.

3rd International Conference on Mathematics in Sport [12] extends Kelly's staking algorithm to incorporate lay bets, which are bets that something will not happen and are available on betting exchanges, in addition to back bets. This continues providing a numerical solution to Breiman's generalisation, which is a convex optimisation problem. This work has been cited in [13]. The material featured in Chapters 3 and 4 is based upon this paper.

IMA Journal of Management Mathematics, Volume 24, Number 3 [14] extends the ideas above and provides an example using football data from Betfair. It then incorporates commission, providing an arbitrary close approximation. This paper also examines the size of some football markets, in particular considering how large wealth would need to be before market size is a constraint to growth. The material featured in Sections 4.5 and 5.1 is based up this paper.

4th International Conference on Mathematics in Sport [15] suggests a novel use of Kelly to calculate prices used in market making, and performs a test in an active market. The material featured in Chapter 7 is based upon this paper.

2. Background

2.1. Measuring Risk

Whilst gaming has existed for at least four thousand years and gambling for at least two thousand, published studies appear to be relatively recent. The first formal discussion of what we now call probability theory appears to have arisen when Chevalier de Méré wrote to his friend Blaise Pascal asking for help on the *problem of points* which Pascal and Pierre de Fermat solved in 1654 in their correspondence with each other as described by Ore [16]. The problem of points asks for the fair distribution of stakes for an incomplete game (somewhat similar to Duckworth-Lewis [17], but two hundred years earlier, and for dice). Their ultimate solution included a number of breakthroughs. In particular they considered what might have happened had the game continued, and importantly introduced expected value. Poundstone has written a very entertaining account [18] of most of the following people and it is to be recommended to anyone interested in the following sections.

2.1.1. Daniel Bernoulli

In 1738 Daniel Bernoulli [8] solved a problem which had been posed by his cousin, Nicolas Bernoulli 25 years earlier. The problem became known as the St. Petersburg Paradox after the solution's published location. Nicolas imagined a game of chance with a payout of 1 unit if a fair coin toss reveals a tail, and doubling each time a head is shown before the first tail. If you assume that the casino offering the game has unlimited resources, the question is what should one pay to play such a game. In finite games the solution is normally the expected value (EV), but not in this case.

$$EV = 1 \cdot \frac{1}{2} + 2 \cdot \frac{1}{4} + 4 \cdot \frac{1}{8} + 8 \cdot \frac{1}{16} + \dots = \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \dots = \infty$$

Bernoulli solves this by suggesting that people value a fixed monetary gain differently dependent upon their wealth. In particular a poor man values a fixed gain more than a rich man. He considers a very poor man finding a lottery ticket with a 50% chance of paying out twenty thousand ducats and a 50% chance of paying nothing. He claims

that it would be sensible for him to sell this ticket for 9000 if he were able, and yet it would be sensible for a very rich man to buy such a ticket at that price. Bernoulli suggests a utility function, which has to be increasing and concave, and he chooses the log function.

When the log utility function is applied to the St. Petersburg Paradox what you are prepared to pay depends upon your wealth. Starting with a wealth of 13 units you would be prepared to pay 3 of them, to be prepared to pay 6 you would need a wealth of about 1100, and to pay 10 your wealth would have to be more than 250 000.

Because of the concavity of the log function a fixed loss hurts more than the same gain benefits. “This is Nature’s admonition to avoid the dice altogether” claims Bernoulli as both players in a fair game expect to make a loss.

Bernoulli continues with a discussion of shipping insurance. With the log utility function it is possible to work out if insurance should be taken out, or underwritten. Again Bernoulli explains why it might make sense both for someone to purchase insurance and someone to underwrite it, given different amounts of wealth, despite this being a zero-sum game. Bernoulli points out that for the underwriter there must be a positive expected value in the policy. The expected log clearly shows the value of diversification, both to shipping, and to financial investments. Bernoulli claims that as these mathematical results “accord so well with experiences of common practice we can be comfortable with the assumption” (that we operate using a log utility function of our total wealth).

2.1.2. John L. Kelly Jr.

John Larry Kelly, Jr, also acknowledged a problem with expected value. Should a gambler be lucky enough to find a bet with positive expected value how much should he stake? Following an expected value maximising approach would have the better stake his entire wealth on this bet. Repeating this, would soon lead to certain destitution. Kelly’s solution [9] is to bet a fixed fraction of wealth, and to chose that fraction such that it maximises the rate of growth of capital.

Kelly worked at Bell Laboratories, where a few years earlier Claude Shannon had published his work [19] on Information Theory. Answering the question of how much to stake on a favourable bet, Kelly, writing in the same the journal as Shannon, uses the the language of information theory for much of his paper. In the particular case of a bookmaker offering *fair odds*, the maximum growth rate predicted by Kelly bears a strong resemblance to Shannon’s rate of transmission. Poundstone reports [18] that Kelly was inspired by news of a betting scam. One of the most popular television programmes of the time was a quiz show called ‘The \$64,000 Question’. In this show a contestant was asked a series of questions with his prize money doubling on every correct answer.

The show was popular and some viewers would gamble on the outcome. However, the show was shown at prime time on both the East coast and West coast. Some gamblers had friends in New York telephone the results through to them several hours before the show started on the west coast. This supposedly was the inspiration for Kelly's noisy transmission line.

Starting with capital V_0 , Kelly maximises the rate of growth (G) of total capital (V_N) at stage N where

$$G = \lim_{N \rightarrow \infty} \frac{1}{N} \log \frac{V_N}{V_0}$$

At each bet the gambler maximises G , the expected value of the log of his capital. It must be noted (as it was by Kelly) that this is not imposing any utility function on the gambler; the log arises naturally as we maximise compound growth. There is no measure of his value of increased capital. This is interesting, as the result is mathematically very similar to that of Bernoulli, which was first translated into English the year before Kelly submitted his paper. Bernoulli's work had previously been translated into German and the results were well known amongst many mathematicians. However, Bernoulli does assume that the gambler has a log utility function and that a fixed gain for a very wealthy man is less valuable than the same gain for a poorer one.

Whilst the simple case with which Kelly starts has a parallel with Shannon's information theory as Kelly includes a bookmaker's take this becomes less clear. And it is these results that are of most interest. The paper noted one particularly interesting feature which many gamblers do not like. If a particular game is such that there is no bet which is favourable then the Kelly strategy states that nothing should be bet on any event. However, if there is at least one event which is favourable then the Kelly strategy may require a bet on not only that event, and any other favourable events, but also on some unfavourable events. This made it unpopular with some gamblers who objected to betting on events which on their own have a negative expectation, or betting where the gambler's estimate of the winning probability is lower than that implied by the odds. It is easy to show that they should not be concerned. Imagine being offered odds of 2.20 for heads and 1.99 for tails on the toss of a fair coin. Anyone betting on heads only stands a 50% chance of loss and so needs to limit the size of the bet. Conversely, anyone betting £11 on the unfavourable bet for every £10 staked on the favourable bet is sure to win about £1, and should bet his entire wealth and borrow and bet more.

Kelly shows that eventually his strategy will get ahead and stay ahead of any other essentially different strategy. Most observers, for example Markowitz [20], take Kelly to be the maximum which should be staked as John Haigh explains [21], a smaller bet than that suggested by Kelly will reduce the growth rate, and the risk (the chance of substantial loss) whilst a larger bet than this will increase the chance of loss yet still have a lower growth rate.

There are assumptions made by Kelly which need to be noted. This is an asymptotically optimal strategy, which has no restriction on bet size, or on the division of capital. Whilst it will eventually prove to be the best strategy the path taken might be very uncomfortable for many. In the real world we do have restrictions on the division of capital; there is a minimum bet size. If our capital falls such that the bet we need to make is consistently smaller than the minimum bet, we should consider ourselves to be broke. We might also find ourselves in the much more pleasant position that our capital has grown so much that we need to place bets that are too large for the market. People tend not to worry about that as much. Kelly has been applied to the financial markets. In many cases the starting portfolio is so large compared to the minimum transaction size (which can be less than £1) the infinitesimal divisibility of capital assumption is reasonable. And as many financial markets are liquid and deep, very large transactions are usually possible.

We have noted that betting the Kelly fraction is optimal, but we have not been precise. Leo Breiman showed that such a strategy is optimal in more than one sense. Breiman had already shown [22] that it does not matter if the game changes on each iteration. As long as the optimal fractions are chosen on each occasion then the whole process will be optimal. A year later he published a substantial paper [10] which proved a number of results. Perhaps the most significant is that the Kelly criterion not only maximises the asymptotic wealth, but also minimises the expected number of bets required to reach a predetermined fortune, which might be a key factor for an investor/gambler.

Rather than focus on a single horse race, as Kelly did, Breiman was interested in a more general version of the problem, allowing overlapping sets of outcomes called events, each with odds. He wanted to find the fraction of wealth, z_j , staked on each bet j , where "... X is a random variable taking values in the set $I = \{1, \dots, s\}$ such that $P\{X=i\}=p_i$ and let there be a class \mathcal{C} of subsets A_i of I , where $\mathcal{C} = \{A_1, \dots, A_r\}$, with $\bigcup_j A_j = I$, together with positive numbers (o_1, \dots, o_r) ."

$$G = \sum_i p_i \log\left(\sum_{i \in A_j} z_j o_j\right)$$

Unfortunately, Breiman then stated that the solution to the above equation is "difficult in the general case".

All of the definitions so far, and indeed the thesis under consideration assume the availability of a favourable bet; a bet in which there is some expectation of a gain. Kelly's paper has a simple definition of a favourable bet, that there must be at least one event i , with probability of success p_i and decimal odds o_i available such that $p_i o_i > 1$.

The above definition is too simplistic. For example, Thorp [23] shows in his Blackjack card counting system that it is often necessary to make bets which are not immediately favourable, but which allow access to favourable bets later. Breiman takes a very general

view in his definition of a favourable game. If S_n is the capital after n plays then a game is favourable if almost surely

$$S_n \rightarrow \infty.$$

Breiman actually offered credit for this definition to Dubins and Savage [24]. Their book concerns unfavourable games, but offers insights into optimal betting on a favourable game in the finite case.

2.1.3. Harry M. Markowitz

Harry Markowitz's work has been acknowledged by diverse organisations. In 1989 he received INFORMS John von Neumann Theory Prize, awarded for a fundamental and sustained contribution to theory in operations research, for his work developing sparse matrix tools [25]. A year later he received (jointly) the Nobel Prize [26] for his "pioneering work in the the theory of financial economics". Much of the theory of finance of the last thirty years has been based upon the Capital Asset Pricing Model, which in turn was based upon a paper Markowitz published in 1952 [27], and more detailed extension into a book [28].

Markowitz argued that a portfolio could be characterised by its expected return and its variance. For any given return there was a portfolio with a minimum variance, or for any given risk (or variance) there was a portfolio with a maximum return. This set of portfolios is called the efficient frontier; no portfolio on this frontier could be said to dominate another. It is the attitude of the owner of the portfolio to risk which should determine the position on the efficient frontier.

This idea, and its extensions, by amongst others Sharpe [29] and Miller [30], who shared the Nobel prize with Markowitz, became CAPM – the prevailing received wisdom for many years, especially when combined with the work of Fama on the randomness of stock market prices [31]. Even today, whilst a few appear to be able to make abnormal returns (see Thorp below) most accept that the market is mostly efficient, and sales of index funds have expanded with new products such as Exchange Traded Funds (ETF). The financial crisis of the last five years has led some [32] to question market efficiency. Others argue that whilst it has flaws it is still better than anything else available [33].

Latané [34], Thorp [11] and Markowitz [35] each produced work trying to reconcile Maximum Expected Log (MEL) or Growth Optimal Portfolios (GOP) as the generalisation of Kelly's work to stock markets was often termed, with portfolio theory. Markowitz had himself discussed it in Chapter 6 of his book [28]. Christensen [36] refers to some of the research comparing the GOP to mean-variance portfolios. He concludes that the GOP appears to be more risky than the efficient frontier, but "performs rather well" even on shorter than expected time horizons. However, he finds that these results are

not statistically significant.

2.1.4. Edward O. Thorp

Edward O. Thorp has probably done more to popularise Kelly staking than any other writer. He has written extensively on gambling and investing, both academically and popularly. Significantly he has followed his own teaching with great financial success. His book [23] was one of the first to suggest a practical use of the Kelly criterion (as he so named it). He had realised that as the casino game Blackjack was played without replacing the seen cards back into the deck, then towards the end of a set of deals a player who could remember all of the cards might have an advantage over the house. In fact he was able to create a system of counting the number of high cards that had been played. High cards (tens and aces) are better for the player than for the house. If fewer are played then there must be more left in the remaining deck, which would be better for the player. He used the Kelly fraction to calculate how much to stake, although with the modification referred to earlier. He still had to bet a minimal amount on unfavourable games in order to remain at the table to take advantage of the next time the deck became favourable. Several hundred thousand copies of this book were sold.

Unsurprisingly, most research on growth optimal staking has concentrated on financial markets rather than betting markets. Thorp was an early and prolific advocate of Kelly's method in managing financial risk. Financial markets typically have continuous returns rather than the discrete returns afforded by betting markets; otherwise, they are remarkably similar. Whilst Latané [34] was one of the earliest, and Thorp one of the most prolific, for example, [11],[37] and [38] (with Rotando), many of the leading financial authors have published in this area, see Merton [39] and Markowitz [35]. Indeed Thorp founded Princeton-Newport Partners in 1969. Between then and 1988 the fund had a compound average annual return of 15.1%, after fees [18]. The S&P500 averaged 8.8% over the same period. The standard deviation of that return was much lower than comparably successful funds of the time (such as those of George Soros and Warren Buffett).

In much of the research there is little difference between that concentrating upon financial products and that on betting. The typical research (see most of those in the previous paragraph) takes a discrete time period approach. Each period may be considered as a bet on a portfolio of stocks, each with a random return. Given that in financial markets it is quite usual to buy something and then sell it later, and that Thorp is very successful in financial market, it is surprising that he did not consider doing the same in betting markets until 2008 [40]. Thorp was reporting an email conversation with Proebsting who had noticed something surprising. He calculated the Kelly fraction of a bet, and then considered what to do if the odds improve. If one solves the exponential

log equation as Thorp did, (and the present writer [41] although with too little analysis) the result is such that the new fraction plus the first fraction is larger than the fraction for the new odds on their own. Thorp named this Proebsting's paradox. It leads to the disturbing conclusion, that if successive improvements in odds occur and are backed, then as the odds increase the fraction bet will tend towards one.

Aaron Brown has provided a solution to this in a private communication with Thorp, which is reported by Wikipedia [42]. He considers that the improvement in odds should be considered as a loss. The existing bet should be notionally sold back at the new odds, and then a new bet considered, but on the new, lower capital.

Much of this thesis refers to betting markets. Apart from their scale and social usefulness there is another difference between many financial markets and betting markets. It would be possible to divide each betting market into small periods of time and consider each period independently as is done in much of the above referenced work. However, in each case there will be a final settlement price which determines the payout and thus it is also possible to think of all of these periods as correlated. In classical derivatives pricing, such as Black and Scholes [43], the derivative is based upon the price of an underlying investment, the price of which is visible. It is possible to consider a no arbitrage assumption to price the derivative. In some other derivatives such as those involving weather [44] or [45], this is not possible. There is no underlying security to observe. The pricing of sporting bets might be considered comparable to these derivatives.

2.1.5. Paul A. Samuelson

Adoption of Maximising Expected Log has not been universal. Most notably, another Nobel winning economist, Paul A. Samuelson has been particularly critical [46] and [47] (with Merton). One early paper by Hakansson [48], which compared mean-variance efficient portfolios with the MEL portfolio did contain an error, which was pointed out in Merton and Samuelson's paper. Hakansson later corrected this [49] and argued that this error did not change his fundamental argument. One of Samuelson's criticisms is that whilst maximising the geometric mean at every stage leads to a higher terminal wealth, almost certainly, after sufficient time, it does not maximise the expected utility in general. Thorp [50] agrees, and claims to have proved as much. He shows [51] that in general one utility function will not dominate another. He also states [40] that this point is not relevant to any of his and Breiman's or Latané's claims. Markowitz [35] does appear to claim that in some circumstances the geometric mean may approximate that of a bounded utility function, which he later corrects.

MEL maximises long term gain and avoids ruin. These are the two features which recommend it to many; they are also the features of which Samuelson is critical. He

argues that for many the long term goal is not relevant, in which case, argue Thorp and Markowitz and others, do not use it. Avoiding ruin is more difficult to dismiss. The financier Jarrod Wilcox has a good analogy. Every day millions of Americans drive their cars, and yet every day, on average, 89 people will die (2011 data). If we used the Kelly Criterion to select our method of transport we would not drive. Most people do not appear to think that way. Most are happy to take substantial risks if the probability is low.

It appears that Samuelson, as an economist, was unable to consider anything which was incompatible with utility theory. Whilst also an economist, Latané had worked on Wall Street and perhaps was able to see that utility theory was a theoretical step too far for many. Bell and Cover [52] "... are not interested in utility theory in this paper... [but are] ... interested in objective aspects of portfolio selection". Remember that Bernoulli did want to impose a log utility function on everyone; Kelly, Markowitz, Thorp and the rest did not. The use of MEL is a result of maximising compound growth. Markowitz, as an economist, is happy with utility theory. He does not argue that those who use utility functions such as those provided by Goldman [53] as a counter example to Thorp should use MEL; he argues that they should not use any utility function other than log. Luenberger [54], years later, extends this. He shows that under certain conditions, including assuming long term goals, log is the only sensible choice of utility function.

Samuelson had provided another criticism [55] concerning the strategy's asymptotic nature. If an investor is not willing to accept one bet, why would he accept a sequence of them? Bell and Cover answered this in 1988 [56] by showing that in some senses the portfolio which maximises the geometric mean is also game theoretically optimal in both a single play and a finite number of plays.

Poundstone concludes his chapter [18] (page 221) on this debate with an amusing anecdote. Merton and Samuelson [47] respond with a snide aside to an imaginary question from someone who wants to maximise his compound return "It is a free country. Anybody can set up whatever criterion he wishes." They continue though that anyone who understands this matter has a "duty to help people clarify the goals they will, on reflection, really want". Thorp's response was that when he explained the Kelly criterion to investors most people say "Yeah, sounds great to me, I want that".

In an attempt to have the last word Samuelson published a paper written entirely in words of one syllable [57]. He agrees that an MEL investor with probability approaching one as time increases has more wealth than another investor following a different strategy. He argues that this second investor should not change his strategy because of this result, and that this is not an argument to promote such a strategy in general.

The decision to follow this strategy is left to you, dear reader; if you do not like a strategy which maximises the probability of increasing your wealth more than any other strategy please feel free to choose an alternative. We have a brief discussion of why you might




Figure 2.1.: A graph comparing a Shannon constantly rebalanced portfolio (50% cash, 50% stock) and the underlying equity. The equity randomly doubles or halves in price each day.

not want to in Section 8.2.

2.1.6. Thomas M. Cover

Thomas M. Cover co-authored one of the leading books on Information Theory [58], Google Scholar lists it as having been cited more than thirty thousand times. As quoted above, Cover (with Bell) [52] argued that MEL was a utility-independent optimality. This paper shows that choosing the Kelly portfolio also maximises the probability that an investor will win a game with another investor when they both have to choose portfolios for a single period. In a further paper [56], they extend this further. The MEL portfolio is “game theoretically optimal in a single play” or in any number of plays of the stock market. The log optimal portfolio is robust independently of the knowledge of the length of the game.

This answered Samuelson. We should choose the log optimal portfolio for a single bet and for a sequence, even if the bets change at each stage; we merely adjust our bets at each stage as appropriate.

Shannon gave two lectures on an investment idea he had [18], but never published any details. He suggested something that we would now call a constantly rebalanced portfolio (CRB). He imagined a two instrument portfolio half in cash and half in a particular equity. This equity is very volatile, but without any fixed drift, its price

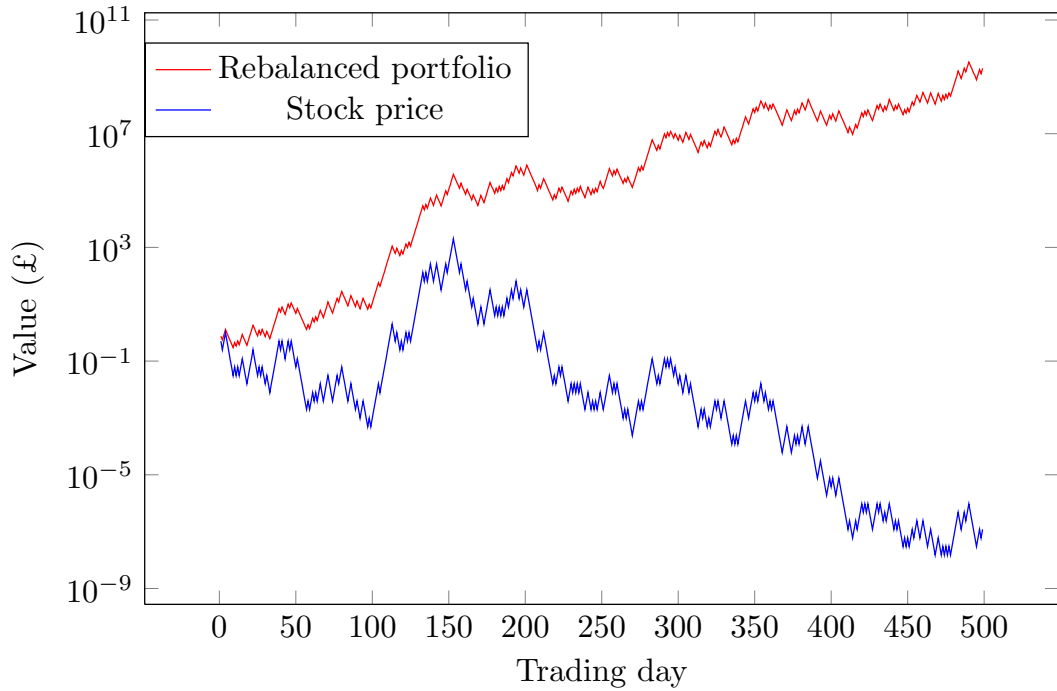


Figure 2.2.: A graph comparing a Shannon constantly rebalanced portfolio (50% cash, 50% stock) and the underlying equity. The equity randomly doubles or halves in price each day. On this occasion the stock price fell significantly over the 500 days of trading, and yet the portfolio still did well.




Figure 2.3.: A graph showing what would have happened had the stock crashed suddenly in price from day 200 (perhaps Boo.com or Marconi). The stock falls quickly, with little volatility. So when compared to Figure 2.2 the portfolio does very badly (from 1 000 000 to 10 000).




Figure 2.4.: A graph comparing a Shannon constantly rebalanced portfolio with commission and the underlying equity. Commission charges of 0.5% are deducted from both purchases and sales. The dotted lines show a similar pair of portfolios but this time the underlying equity has a much lower volatility.

doubles or halves every period. At the end of every period the proportions are re-adjusted back to half-half. Figure 2.1 shows one particular random simulation of this. The equity is relatively unchanged but the portfolio had grown considerably. In addition the portfolio has had a much lower volatility. Even if the stock fell significantly as in Figure 2.2 the portfolio will still do well (though not as well as Figure 2.1). This is assuming that both money and stock is infinitesimally divisible. In reality when the price fell the cost of transactions would increase beyond any reasonable assumption. Shannon was asked by an audience member if he had tried it himself and he replied that he thought that commission would kill it as an idea. He was probably trying to demonstrate that clever stock market strategies were not incompatible with the work of Fama and others suggesting that markets moved like a random walk. However, his own investing appears to have been the exact opposite, following a buy and hold forever policy, perhaps because of tax considerations.

In fact it appears that commission need not be such a large problem. (Commissions have fallen substantially since 1966 when Shannon gave the first of these lectures.) Figure 2.4 is produced using the same simulator as Figure 2.1. This time a commission of 0.5% is charged on each purchase and sale. Commissions vary across products, countries and customers, but many investors will be able to do much better than that. The portfolio still does very well. The dotted lines on the same graph show a greater concern. The volatility in this simulation is very high. The dotted lines multiply or divide the price by

1.1 in each period instead of 2. This is still a high volatility, and yet when commission is included the portfolio does not perform particularly well, but at least it does not lose much.

The portfolio in Figure 2.3 shows what happens when a share price collapses, never to recover. All the way down in price more and more shares are being purchased, soon to be valued at zero. The equity in this graph could be any of the failed dot com stocks of the year 2000 crash, or Enron, Marconi, or Northern Rock. Almost as bad would be a company which fell significantly in value only to be taken over or de-listed. In this figure the crash happens on day 200, but of course it could happen on day two. The CRB loses 99% of its value in this case.

Cover extends Shannon's idea, creating Universal Portfolios [59]. He shows that his Universal Portfolio will asymptotically outperform the best individual stock. Assume $\mathbf{x}_i = (x_{i1}, \dots, x_{im})^T$ denotes the relative change in stock prices on day i , and $\mathbf{b}_i = (b_{i1}, \dots, b_{im})^T$, where $b_{ik} \geq 0$, and $\sum_j b_{ij} = 1$, the fraction of wealth invested in the j^{th} stock on day i is b_{ij} , then $S_n = \prod_{i=1}^n \mathbf{b}_i^T \mathbf{x}_i$ is the factor by which wealth increases in n days. The constraints on \mathbf{b} ensure that the investor is fully invested and neither borrows nor lends. If we choose \mathbf{b}^* with full knowledge of the future market so as to maximise the respective S_n^* . Cover shows that S_n^* exceeds the best stock and most stock market indices such as the Dow Jones average. Cover goes on to create a a portfolio without looking ahead \hat{S}_n such that

$$\frac{1}{n} \log \frac{S_n^*}{\hat{S}_n} \rightarrow 0$$

for every bounded sequence $\mathbf{x}_1, \mathbf{x}_2, \dots$. Cover's algorithm is exponential in the number of stocks. Helmbold et al. [60] have since developed a variation with slightly better properties with linear complexity in the number of stocks. They examined 22 years of market data for pairs of stocks. Had they a time machine to go back in time 22 years to purchase the best performing stock they would have multiplied their wealth by a factor of 52. By following their on-line portfolio selection they would have increased their wealth by a factor of over 110 – without any knowledge of the future. The best constantly rebalanced portfolio would have increased their wealth by a factor of 144. There has been much research in the last twenty years on Universal Portfolios, many having tried to model commission costs, for example [61].

Universal Portfolio can be said to be following a policy of diversification, as recommended by Bernoulli, Markowitz, and many others. But diversification is another rift in the prevailing view, and just MEL itself it appears as though people are arguing about the question rather than the answer. In the UK the top 350 stocks are the most liquid; in the US this number is at least 500. Purchasing such large numbers of stocks would prove problematic for all but the largest investors. The standard solution such as suggested by Malkiel [62] is that about 20 stocks provides almost as much diversification as 100. The argument assumes that stock market prices are independent and considers

the standard deviation of a portfolio. This falls with each additional stock added to the portfolio, but each additional stock added after the 15th reduces it only slightly. Acknowledging that stock prices are not independent, the suggestion is then to make sure that those 20 stocks are drawn from as wide a range as possible of the allowable stocks. In the same book Malkiel promotes such diversification.

On the other side of diversification are Warren Buffett [63], Edward Thorp, Maynard Keynes [64] and the followers of Security Analysis [65]. This book starts off by criticising Wall Street for its focus on earnings per share, which sounds familiar, but this book was published in 1934. Presented with 20 diverse shares which are good investments I suspect that all of the above would have happily diversified. The problem for the value investor, as these investors tend to be known, is that finding good investments, and monitoring them is hard work and, therefore finding 20 is an unlikely achievement. Graham and Dodd go further. They argue that if a stock is undervalued then its risk (in the sense that it will fall in price, rather than its volatility) is likely to be lower. Creating a portfolio by investing in your ten best opportunities might be actually safer than investing in 20, as the second ten are intrinsically more risky.

Buffett provided more clarity on his view in a meeting with Business students in 2008.

“I have 2 views on diversification. If you are a professional and have confidence, then I would advocate lots of concentration. For everyone else, if it’s not your game, participate in total diversification.”

“If it’s your game, diversification doesn’t make sense. It’s crazy to put money into your 20th choice rather than your 1st choice.”

Warren Buffett

He is advocating total diversification for almost everyone else other than a handful of experts. He thinks that the market is more efficient than the average person who has limited time to invest. Therefore he is unlikely to outperform the market. The efficient market followers would not argue with his advice for the small investor, and we can leave the experts to decide for themselves.

2.2. Football Modelling

2.2.1. Michael J. Moroney

In Facts From Figures [66] Moroney is one of the first to suggest using the Poisson distribution to model goals scored in football. He first discusses data from Bortkiewicz’, [67], ‘Deaths by Horse Kicks’; these deaths are a remarkably good fit to the Poisson distribution. Moroney’s football data, whilst “not too bad” a fit to the Poisson distribution it is not as good as it might be. Moroney considers that this is caused by factors

such as weather, and the opposing team. The over-dispersion seen is a common feature of non-homogeneous data, for example [68]. Moroney continued by suggesting instead the Negative Binomial distribution, which can be thought of as a Poisson distribution with a random λ drawn from a Γ distribution.

2.2.2. Charles Reep

Wing Commander Charles Reep is a somewhat controversial but undoubtedly influential [69] figure in English football, (and also to Norwegian football [70]). He has been referred to as the father of modern data capture and analysis, having collected passing data on thousands of matches during his lifetime. He, with others, [71] examined the number of passes each side makes in each period of possession and favours the negative binomial for this distribution. He confirms Moroney's findings regarding goals and extends this research to other sports [72]. He reaches two conclusions of particular note. One is that as goals appear to follow on from short sequences of passes he develops the *long game* – involving long passes. His other conclusion is that football is dominated by chance.

Hennessy writing in The Times [73] perhaps exaggerated this randomness, suggesting that using a pin was as good a way as any of choosing winning teams. Hill [74] responding to this rather than Reep's original work, collected the views of several pundits ahead of the season and compared their forecasts with the actual results, and unsurprisingly found a correlation. He concluded that whilst there is much randomness, other factors appear to appear relevant in the long run.

2.2.3. Michael J. Maher

Moroney and Reep (and others) had used a single parameter for the goal rate, and all had acknowledged that this parameter would be expected to vary with the strength of the opposition. However, the real breakthrough in football forecasting arrived when Maher produced a model which adjusted the rate dependent upon the opposition [75]. This built on work relating to American Football, some of which was cited by Maher, [76], [77], [78].

When team i plays against team j the score is (X, Y) , where X is team i 's score and Y is j 's. In the first section Maher assumes that X and Y are independent. Each team has two parameters α , a goal scoring rate, and β , a defensive factor. There is one further parameter κ which is the home advantage. In football the home team, on average, performs much better than the away team, well known for many years, and analysed by Clarke several years later [79]. κ is constant across the whole division. X

and Y are Poisson distributed with parameter λ and μ respectively such that

$$\begin{aligned}\lambda &= \kappa\alpha_i\beta_j, \\ \mu &= \alpha_j\beta_i.\end{aligned}$$

Some were still wedded to the negative binomial. Pollard, one of Reep's co-authors was still using it [80] several years later.

Maher detected a slight dependence between X and Y and provided a dependent variant using the bi-variate Poisson of Griffiths and Milne [81].

Maher was interested in measuring how well his model fitted the data. Dixon and Coles [82] made a few minor changes to Maher's model to predict the odds of the next (not yet known) match. They fitted past match data numerically to obtain their parameters, and in particular they weighted more recent matches more highly than older matches. Instead of using the standard maximum likelihood method they created a 'pseudo' likelihood. The contribution of each match is exponentially weighted by its age. They went on to collect odds data from bookmakers for 1995 and 1996. At the time bookmakers charged about 11% to place a bet. This is often called the over-round and is the bookmakers' expected profit per match. Despite this large cost of placing a bet they were able to overcome it, showing that their model was significantly better than those employed by the bookmakers. Unsurprisingly this is no longer the case [41], at least when using this model.

Since then there has been considerable effort employed building models for football (and all sports). Much of this work is done for bookmakers and companies which sell prediction data, and therefore do not publish their models. Nonetheless there has been an increase in interest in this field. In 1995, Clarke and Norman [79] examined the home advantage in English football, as had Barnett and Hilditch earlier, regarding artificial pitches [83]. Since then, home advantage has been in decline, which is true in many sports and in most countries [84].

The football models which use Maher's outline all predict the score. Obviously knowing the final score we can say who won - the result. Goddard, [85] compared such models to those which predict the result directly, in particular those such as [86], [87]. He concluded that the difference in the success of each approach is small, and a slightly better performance is obtained using a combination of the two. Teams tend to care about winning or losing rather than the goal difference, and it is not unheard of for the final rank of a table (by points) to have significant outliers of goal difference. In 1937-38 Manchester City scored the highest number of goals in the division (of 22 teams), had a positive goal difference, and still finished second from the bottom [88].

Moneyball [89], a book about Baseball, has had a surprising impact on football modelling. Michael Lewis, writing in 2003 about Billy Beane's attempt to improve the

Oakland A's performance by using statistics to choose players. Since then many other Baseball teams have started examining such data, and the same is true of a few football teams. In an attempt to encourage the next generation of football modellers, Manchester City give away their proprietary data [90] for a season.

A game popular in some countries is fantasy football. There are numerous fantasy games with differing rules, but in essence an entrant to the competition makes up a team from all players in the division, and scores according to the performance of those players in the real matches. Unsurprisingly, there is research into modelling fantasy football [91] or [92].

Of course there are times when the teams do not play their best game. Sometimes a certain result is a good result for both teams, and unsurprisingly, yet with no suggestion of deliberate match fixing such a result can be found [93]. Sometimes there are suggestions of unsporting agreements [94]. At the end of the season one team may be safe from relegation, but with no chance of winning a further trophy, whilst the opposing team may still have something to fight for. One is not surprised if such matches diverge from predicted results [95]. Such a problem (possibly corruptly so) is well known in other sports, perhaps most famously, because of a popular account in Freakonomics [96], Sumo [97].

This leads into extensive accounts of match fixing more generally. One substantial starting place would be Wikipedia [98], but unfortunately there are many scandals, some so enduring or sophisticated that they have acquired the status of names, for example [99] [100] [101]. Recent work includes [102], [103], [104].

2.3. Interior Point Methods

Whilst this thesis does not develop Interior Point methods, it does make extensive use of them. Indeed it is because of their development over the last 20 years that we claim Breiman's general formulation, described above, should be considered solved. In that regard it does appear worthwhile to outline a few key developments; for a fuller treatment see Wright's history [105].

A general constrained optimisation problem might look something like this

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \text{ s.t.} \\ \mathbf{g}(\mathbf{x}) \leq \mathbf{b}. \end{aligned}$$

\mathbf{f} and \mathbf{g} may be any function. Other variants might impose a range on \mathbf{x} , or include equality constraints. Clearly these can be included in \mathbf{g} by adding pairs of inequality

constraints. Interior point methods create a new function which is a combination of the objective function \mathbf{f} and the constraint, often something like $\log[\mathbf{b} - \mathbf{g}(\mathbf{x})]$. These had been used in the 1960's but by the 1970's had been superseded due to problems with ill conditioning near solutions which lie on a constraint (which is common).

Somewhat surprisingly using non linear techniques to solve linear programming problems appears to have been “revolutionary” [105], especially when you consider the canonical form for linear programming:-

$$\begin{aligned} & \min_{\mathbf{x} \in \mathbb{R}^n} \mathbf{c}^T \mathbf{x} \\ & \text{subject to } \mathbf{A}\mathbf{x} \leq \mathbf{b} \\ & \text{and } \mathbf{x} \geq \mathbf{0}. \end{aligned}$$

In 1984 Narendra Karmarkar published [106] a new method for solving linear programmes (LP). Prior to this Linear Programming was considered to be different from the more general optimisation above, of which it is a subset. The standard way of solving LP, was the well known simplex algorithm. Developed by George Dantzig during his time with the US Air Force [107] it became one of the most loved algorithms of the twentieth century. Despite Klee and Minty showing its complexity was exponential [108], and Khachiyan finding a polynomial time ellipsoid algorithm [109], the simplex algorithm was nearly always much faster at solving any practical problem.

After 1984, Interior Point methods and the Primal-Dual method in particular received much attention. Many of the problems which had previously prevented general use had been overcome. Interior point methods are used to solve both, and there are many libraries for both.

We have made extensive use of the non-linear solver Ipopt [110]. This is an open source solver re-written in C++ and maintained by IBM Research. It is based upon the PhD thesis of Andreas Wächter [111] written in Fortran. There have been numerous publications detailing the methods and refinements, for example [112] and [113].

2.4. Gambling: Liberalisation, the Internet, and Betting Exchanges

Even before the emergence of Betting Exchanges, the Internet was making an impact on the betting industry. Forrest [114] suggests that internationalisation brought about by the Internet (amongst other things) was making the UK football betting market

more efficient, a process which has since accelerated [115]. Also having an impact was the regulatory environment which in the United Kingdom, at least, has been relaxing restrictions on gambling for the last 50 years.

2.4.1. United Kingdom Regulation

This is a brief introduction to set betting exchanges into some sort of historical context. A good starting point for more details would be ‘An Economic and Social History of Gambling in Britain and the USA’ [116]. There have been numerous attempts to regulate gambling over the years, once dictated by the religious or moral views of the monarch, or the Lord Protector. More recently, one common theme has been, until the mid Twentieth Century, that the upper and ruling classes wanted to be able to gamble themselves and yet they felt concerned “over the fate of the poor and guileless” [116]. The recent introduction of the (the UK state sponsored) National Lottery is often cited as one of the reasons gambling has become more socially acceptable, and that is another feature common to history. Lotteries had been taking place since Elizabethan times, but the first lottery authorised by Parliament took place in 1694 and between then and 1826 170 lotteries took place. At the same time all other lotteries were illegal (until 1934) [116]. Interestingly, one of the (many) causes of disquiet amongst the American colonists was the extension (from Britain) of this ban on lotteries in 1769 [117]. Many regions had previously organised lotteries to raise funds for schools and universities. The lotteries in the UK funded improvements to London’s water supply, Westminster Bridge, and the British Museum [116].

The UK had already prevented large wagers from being enforceable in court but the Gaming Act 1845 ensured that no wagers were legally enforceable [118]. This led to a surge in cash betting shops which were closed a few years later [119] when off-course cash betting premises were banned. The Betting and Lotteries Act 1934 [120] clarified and tidied up some details. It also repealed much older legislation including that regarding lotteries. This made way for the Totalisator, a device to calculate parimutuel betting. Now referred to as the Tote this was allowed at horse and dog tracks.

Betting shops were legalised by the Betting and Gaming Act 1960 [121]. By design, they were forced to be as unappealing as possible. Rule changes in 1986 give one an idea of how bad they must have been: chairs and televisions were allowed, and instant coffee machines [122], [123]. However unwelcoming they were, by 1963 there were 14,388 [116]. This act also accidentally legalised casinos.

The Royal Commission that led to the 1960 Act had wanted to allow private gaming, such as bridge, but not allow commercial gambling and recommended no bankers games. However, the Act was implemented with slightly different wording which did not explicitly prevent such games [124], and many casinos allowed the players to take turns

at being banker if they so wished (and had resources, and were not intimidated by the bouncer). The Act had also intended to allow small scale gambling by other organisations (raffles, bingo, sweepstakes) as a fund-raising activity not related to their main purpose. But again the wording was not clear. The betting shops had been intended and they were licensed and taxed, whereas the casinos had no restrictions. There were an estimated 1200 [124] by 1968.

James Callaghan, the then Home Secretary spoke of the 1960 Act “the Act precipitated the very evil it was meant to prevent”. He introduced The Gaming Act 1968 [125]. This act was successful, although to some extent the genie was out of the bottle. Casinos were limited geographically, they were licensed and registered, by local magistrates and the newly founded Gaming Board which had wide discretion to exclude people deemed undesirable. By 1972 there were 125 casinos, and organised crime had largely been eliminated [126].

On 14th November 1994 the National Lottery was launched. “The lottery created hysteria ... people who had never placed a bet in their life queuing up ... before the Saturday-evening deadline” [122]. In the financial year 2011/12 the National Lottery had sales of £6.5 billion [4].

The National Lottery was a double edged sword for betting shops. Initially, it led to a drop in profits as many small gamblers switched to it [122]. But at the same time it became increasingly difficult for the government to defend harsh restrictions on gambling shops whilst running and promoting a lottery from every supermarket and newsagent in the country. The National Lottery did all but kill off the Football Pools. This had been a de-facto lottery for the vast majority of the weekly players. Entrants had to predict which football matches would end in a draw. Some used their knowledge, but most played the same numbers every week. The latter class of people switched to the lottery which was easier to understand and had larger prizes, a higher (less negative) return, and most of the withheld money went to good causes.. The pools, as the name suggests, were pooled betting, and as many left the prize money fell into insignificance.

A number of deregulations followed the introduction of the lottery to allow betting shops to compete [123]. In common with other shops they were able to open on Sundays. The blank window rule was removed; window advertising of products and prices was allowed. Some advertising was permissible, although not on radio or television. Shops were allowed to have a limited number of slot machines. This last rule would prove very important. In the year 2012 about half of betting shops’ Gross Gambling Yield (GGY) comes from slot machines [2]. GGY is the amount of money staked less the amount paid out in prizes. If one considers the cost of maintaining those machines compared to the rest of a bookmaker’s operation it can be argued that almost all profit comes from slot machines.

About this time betting duty was cut. Prior to 1996, 10% of all winnings from off course

betting shops were retained, 7.75% was a tax, and for bets on horse racing the rest was paid as a levy to the racing industry. In 1997, the tax was reduced by 1%. This cut would prove to be too little, too late. The same year one Victor Chandler founded a telephone betting service in Gibraltar. Bets on this service did not attract the betting duty. Three years later he moved his entire telephone business there. Within four years turnover increased to over £1bn [127]. The government abolished the tax in 2002, taxing gross profit instead, with the horse racing levy following suit. This failed to stop the march overseas. With the increasing use of the internet, overseas betting became even easier. The Gambling Commission, the main gambling regulator, report that GGY regulated in the UK accounts for about 4% of the total remote (non betting shop, or track) gambling [2]. Another estimate puts the amount of offshore much higher, in fact more than double [128].

As we argue elsewhere in this thesis there are many activities indistinguishable economically from gambling. The Financial Services Act 2000 [129] contains a section ensuring that contracts for difference are recoverable in court. Previous legislation had also included similar provisions as the Government was concerned about several financial activities which people might try and argue were betting losses. There had been such a case in the Court of Appeal a few years earlier. In *City Index Ltd. v Leslie*, Leslie had refused to pay money owed to City Index from losses on a contract for difference, arguing that they were gambling losses and not enforceable under the 1845 Act. He lost.

The Gambling Act 2005 [130], further liberalised betting and gaming. One high profile feature of this Act had been to allow a new super casino, Las Vegas style. Manchester was chosen as the location, but then this policy was reversed. Of the other casinos allowed by this legislation only one has opened so far. This is the first substantial piece of legislation to deal with the Internet. It had been hoped that some the companies which had followed Victor Chandler might return; in fact more moved away. The Act did allow television advertising. Restrictions on casino games were relaxed; previously the regulator maintained a list of approved games, now it has a list of banned games, which is currently empty. Thomas Aquinas, the Thirteenth Century Theologian, is said to have permitted gambling, providing “it was not motivated by covetousness, characterised by unfairness, or used to exploit the young or psychologically immature” [131]. The Act is remarkably similar and maintains the overall policy position prevalent for the previous fifty years, to allow gambling whilst:-

- preventing crime or disorder, and making sure it does not support crime,
- ensuring that gambling is fair and transparent, and
- protecting children and the vulnerable.

Whilst many object to gambling on moral or religious grounds, and some of those think

that no-one should be able to gamble it is probably a fair opinion to claim that in most cases these objectives are being met.

The Government has recently announced a change in gambling tax policy. There will be an attempt to tax at the point of consumption, rather than the location of the betting equipment. This might raise as much as £240 million [128]. Perhaps this will encourage some betting companies to return to the UK, but perhaps not. Betfair, for example, takes about half of its bets from customers in the UK [132]. Half of its commission would be liable for the new tax. If it relocated to the UK all of its commission would be liable.

2.4.2. Costs and Benefits

We cannot discuss gaming and betting without acknowledging that it has costs to society. This section is short merely because it is beyond the scope of the author's experience. It is worth remembering that this is focused on the UK, and gambling varies greatly across the world. Munting reports [116] that whilst horse racing was the prevalent form of gambling in the UK in the nineteenth century, whereas in Germany there was hardly any horse racing and most gambling was on the Klassenloterie. This lottery was too expensive for the working class and so was a tax on the wealthy. Such differences in the most popular types of gambling remain today.

Figure 2.5 shows the Gross Gaming Yield (GGY) of the US market[133], broken down into three sectors: casinos, including commercial, Native American, and Card Rooms; bookmakers, including pari-mutuel; and lotteries, including bingo and charity lotteries [1]. Figure 2.6 is a similar chart for the UK, but broken into four: casinos, including arcades; bookmaking; lotteries, also including bingo; and remote betting which includes bingo, gaming and betting, but not lotteries [2], [3], and [3]. Comparing the overall size in 2009 with that of the UK, shown in Table 2.2, we can argue that they are broadly similar, perhaps slightly larger in the USA. Using the exchange rate of £1 = \$ 1.54, prevailing on 28th July 2013 GGY per head is \$291 for the US, and \$238 for the UK. The US is a much wealthier country; if we examine the fraction of GDP arising from GGY we see 0.68% for the US and 0.57% for the UK. Indeed had we used £1 = \$ 1.81 for the exchange rate these would have been the same and in 2009 the GBP exchange rate varied between USD 1.40 and 1.90.

It is the breakdown within the markets in Figures 2.5 and 2.6 which reveal the difference. Table 2.1 shows the numbers for 2009, which is the only year common to both figures. In the UK lotteries (almost all of which is the National Lottery) take a greater share than in the USA, but both are significant. In the UK there is very little casino gaming whereas in the USA it is by far the largest sector.

The different types of gambling in each country might account for the different rate

Sector	UK	USA
Lotteries	46.3%	30.5%
Casinos	16.1%	66.2%
Bookmaking	37.6%	3.3%

Table 2.1.: The distribution of regulated gambling in the UK and USA for the year 2009. The UK casino and booking numbers exclude remote gambling. Definitions as in Figures 2.5 and 2.6.

	UK £	USA \$	USA £ ¹	USA £ ²
Total	£8.12 Bn	\$ 89.26 Bn	£58.0 Bn	£51.0 Bn
Per capita	£131	\$ 291	£189	£166

Table 2.2.: The size of regulated gambling in the UK and USA for the year 2009. Definitions as in Figures 2.5 and 2.6. The exchange rate used is the current exchange rate (£1=\$ 1.54) in case 1, and the middle of the 2009 range (£1=\$ 1.75) in case 2.

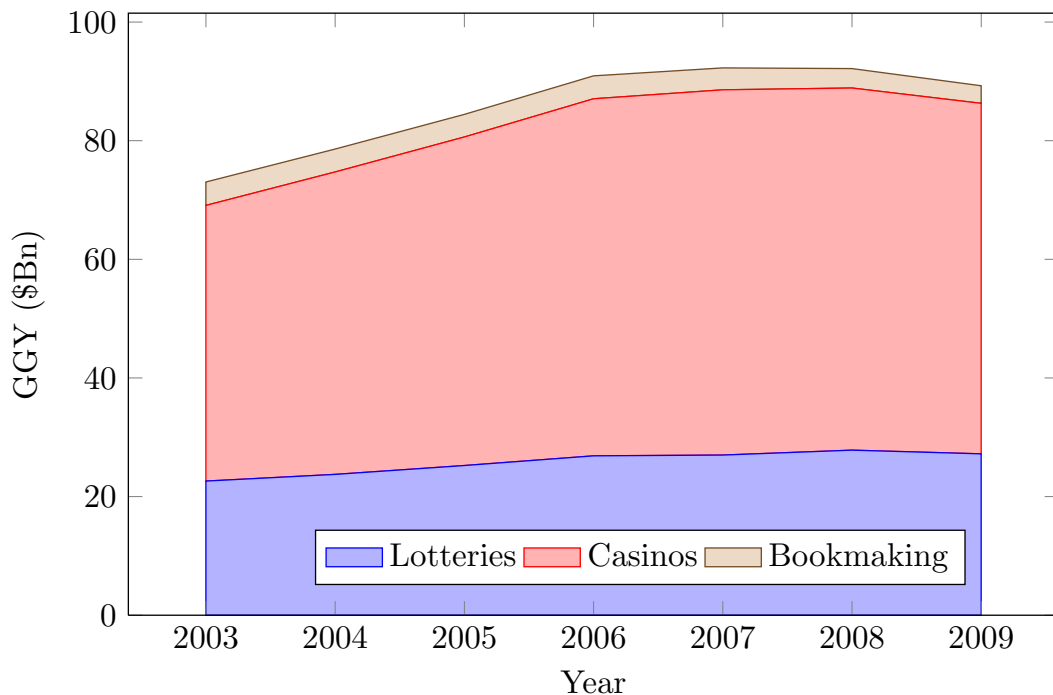


Figure 2.5.: A graph showing the size of the three main US market sectors. Casinos includes Commercial Casinos, Native American Casinos and Card Rooms, Lotteries includes Charitable Games and Bingo, and Bookmaking includes Pari-Mutuel Wagering. Data from Statistical Abstract [1].




Figure 2.6.: A graph showing the size of the four main market sectors. Casinos includes Casinos and Arcades, Lotteries includes Bingo, and Remote includes all forms of remote betting regulated from the UK, but not remote lottery transactions which are included in lotteries. Data from Gambling Commission [2], except lottery data which is from the National Lottery Commission [3] and [4].

of problem gambling reported in each country. For the UK the Gambling Commission conducts a survey every five years and the most recent [134] reports problem gambling below 1% of the adult population, whereas this is about 3.5% for the USA. This survey also reports that problem gambling is particularly low for lottery players, and particularly high for casino gamblers. Therefore it is possible that different types of gambling available in different countries account for much of that difference. The rest of this section will refer mainly to the UK.

In the UK in 2010, 73% of the population aged over 16 took part in some form of gambling [134]. The National Lottery was the single biggest activity with 59% of the population taking part at least once. Economically, it could be argued, that the National Lottery is one of the least fair forms of gambling available in the UK. For every £1 staked, only £0.50 is returned. However, for many, this is the reason it is popular as £0.28 goes to good causes. Most people will have passed by a sporting facility, a concert hall or a museum which has been built, or significantly improved with lottery money, totalling over £30 billion since 1994.

Obviously before gambling was legalised, by definition, there was an association with crime and gambling. Betting shops and casinos have large amounts of cash, which is very attractive to money launderers and other criminals. There was clearly a problem before 1968, but just as Nevada significantly reduced crime by a mixture of encouraging corporate ownership of casinos, and a strong regulator [135] the UK has significantly reduced criminal activity. There appears to be consensus that the Gaming Act 1968 largely succeeded in its aim of removing criminality from the gambling industry [126]. Haringey recently held a meeting to discuss betting shops in the Borough and reported “no evidence received to link betting shops with crime”, and the police reported [136] that betting shops were a “focal point for crime where this was already known to be a problem”, but did not think that they increased crime in themselves. There is no evidence that the crime rate in Atlantic City is higher than that of Orlando, when weighted by the number of visitors [135]. It does say that it is possible that criminals commit crime at home and then travel to a casino area, but that is (reasonable) speculation.

If we argue that those who want to stop gambling are in the minority, and that criminality has been largely removed from the industry, we still need to consider how well the regulator is doing with its other aims: ensuring fair play and protecting the vulnerable. It appears that it is not doing quite as well with these goals.

“Gambling is a bigger threat to sport than doping” read one recent headline [137]. Other headlines indicated that 50 jockeys might be involved in race fixing, 5 were later arrested, but not one was brought to court [127]. This is not a new phenomenon. Nearly 100 years ago eight Chicago Baseball players were banned for life, having been accused of intentionally losing. Four footballers were jailed for matches fixed in 1964 in the UK.

It has been claimed that this was responsible for the high rate of betting duty which was imposed shortly afterwards [116]. The Gambling Commission has responded, creating the Sports Betting Intelligence Unit, which is working with the sporting regulatory bodies and with other European regulators [138]. Between 1st April 2012 and 31st March 2013, 93 cases of suspicious betting activity have been reported to the Commission [2].

The other accusation against the Commission concerns understanding the payout of Fixed Odds Betting Terminals (FOBT). The return, or payout is often hard to spot and varies between each game. But more importantly the player recycles his winnings multiple times and so the return is usually much lower than the payout a player expects. Haringey have called for this to be displayed [136]. This is probably compounded by problem gambling being much more prevalent amongst the poorer, younger less educated.

Most objections regarding FOBTs are to their very existence, and failure of the Commission to protect the vulnerable. Under the 2005 Act each betting shop is allowed four machines. The recent development of high stake gaming machines has caused considerable concern. Some say that they are the “Crack Cocaine” [139] of gambling. A more measured quote is “The presence of relatively high-stake category B2 (FOBT) machines in high-street betting shops was a source of considerable concern to groups which aim to combat problem gambling. They argued that B2 machines posed a greater risk of causing problem gambling than other forms of gambling” [126]. One author argues that the fall in over-round seen for popular sports in high street betting shops is a loss-leader to get people in the door [140] to play the machines. They have changed the nature of betting shops.

The number of betting shops was in decline, as some betting moved overseas and online, but the Fixed Odds Betting Terminals have reversed that trend. The Gambling Commission report [128] an increase of 300 shops in the last two years. As four are allowed per shop it is not surprising that branches of the same chain can be found near each other. Haringey report [136] that betting shops have moved into areas with much higher footfall and become more prominent. However, other high street retailers and banks have withdrawn from the high street vacating some prime retail space. It may be that there is merely increased availability. Harriet Harman has called this a “blight” on London’s streets. FOBTs are very profitable. They account for at least half of betting shop profits [136], and perhaps much more.

The concern over FOBTs is part of a more general concern for problem gamblers. Whilst surveys show that the number of problem gamblers in the UK is below 1%, lower than in many countries such as USA, Australia and South Africa and broadly similar to European peers such as Norway, Switzerland and Germany, this still represents between 250,000 and 450,000 people. Whilst there are some who use this as part of an argument to ban all gambling [131], they would want to ban gambling in any case. As a very

small proportion of the community abuse kitchen knives [141], should we ban them? Or should we close public libraries because terrorists have been able to research bomb making by reading chemistry books [142]? No, of course not, but as 34% of serious knife wounds are caused by kitchen knives some are considering if it would be possible to have knives with sharp ends short and longer knives rounded [143]. Similarly it could be argued that more could be done to protect people with gambling problems. The profits from a FOBT are large enough that it should be possible to do more.

One type of help for those who realise that they have a problem is self exclusion. They attend their local betting shops and complete a form. There are just over 20,000 new self exclusions each year, which will include fewer individuals making more than one application. But it is poorly enforced. There were more than 12,000 breaches reported last year. Many argue that self exclusion is not working [144]. It should be quite easy to introduce technological solutions. One simple change would have credit and debit cards which have a gambling opt-out. Or introduce facial recognition. Casinos in Ontario have just such a system for self exclusion [145]. These could be built into the FOBTs, with the player being scanned by the machine and only allowed to play if he is not on a list of self excluded players. As these machines are connected to the Internet this should be straight forward to introduce.

There are many other problems with gambling, and we will conclude with reference to two more examples. Firstly, gambling taxation is regressive, poor people pay a higher share of duty than wealthy people. Unlike the Prussian Klassenloterie, the UK lottery is more highly represented by lower socio-economic groups. Casinos and online betting exchanges have a lower proportion of poorer customers and yet face lower rates of taxation. It has been thus for over a hundred years. The football pools were similarly regressive. Secondly the Act appears to have failed too with regard to overseas gambling. If gambling has costs to society, but also benefits with tax and employment, then allowing overseas gaming to expand is not desirable as it causes harm and yet neither pays tax nor provides jobs here.

“The Act permitted UK-based companies to offer online gambling services and continued to allow non-UK providers to target UK consumers, only the UK-based providers were made subject to regulation by the Gambling Commission. The decision to regulate online gambling at the point-of-supply rather than the point-of-consumption and to allow non-UK regulated providers to operate into the UK was widely criticised. This together with significant tax increases had led to migration of online providers offshore. There was broad consensus that the 2005 Act has, thus far, failed to produce a future-proofed regulatory structure for the remote and offshore online gambling industry”

House of Commons Culture, Media and Sport [126]

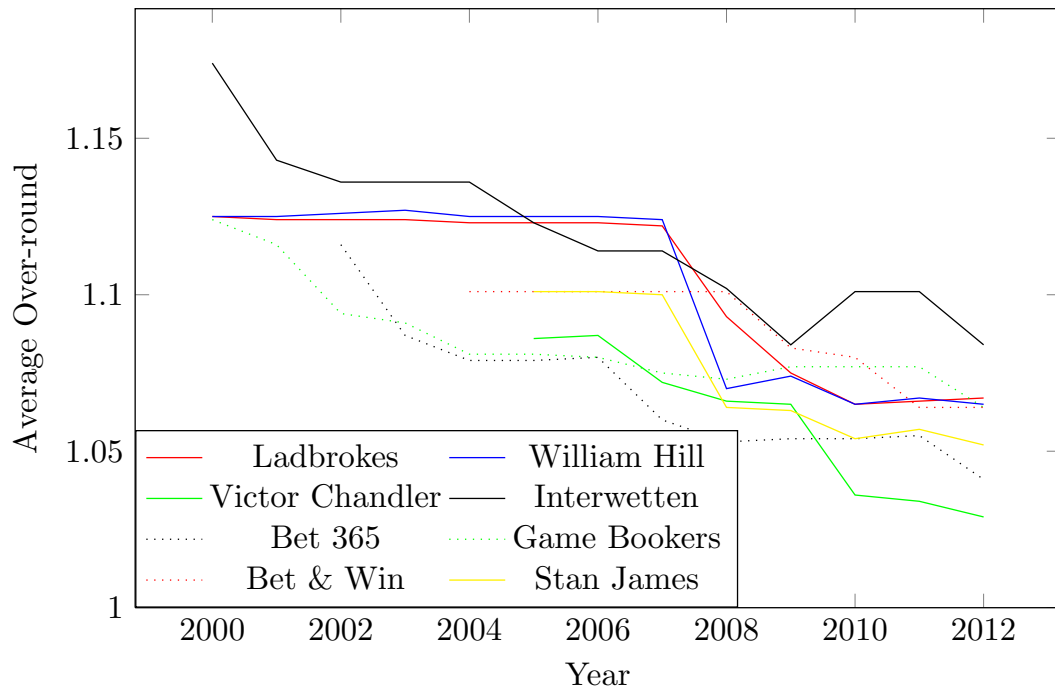


Figure 2.7.: A graph showing the over-round for Premiership match odds for various bookmakers. Ladbrokes and William Hill are the main UK high street bookmakers. Interwetten is Austrian and not particularly active in the UK, whilst the rest are Internet or telephone bookmakers. Data from football-data.co.uk.

Gambling is now widely accepted in the UK as a legitimate entertainment activity [126]. The National Lottery has raised over £30bn for good causes since its launch, in addition to generating £13bn in direct tax and £10bn in wages and indirect taxes. The bookmaking industry employs about 100,000 people and pays £1bn in tax each year, its impact on the economy being possibly as high as £3bn [146]. However, more can be done to help problem gamblers. Perhaps FOBT should be restricted. Haringey wants them to be licensed separately, and further independent research commissioned [136].

2.4.3. Internet and Betting Exchanges

There are numerous articles written about betting exchanges, which for all intents and purposes means Betfair. A good starting point is [147], but [148] and [149] have some interesting information. Prior to the launch of Betfair and Flutter, in 2000, the majority of people betting would use a fixed odds bookmaker, on the track, the high street or on the telephone. Betfair and Flutter merged in 2002 [150]. Betting exchanges claim to have brought several innovations.

Better Prices Bookmakers charged an over-round at the time of between 10 and 20%. Betfair charged a commission of between 4 and 5% for most customers, Flutter charged 2.5%. Better still, a customer could choose his own prices, perhaps making

no bet, or his chosen bet, but only at the right price.

Lay Prices Just like a stock exchange where shares are bought and sold, so the betting exchange has back and lay [151]. Lay prices are particularly important in markets with lots of entrants such as horse racing, although less so in football, and not at all in tennis.

Unwind Anytime A consequence of offering back and lay prices is that sometime after placing a back bet a lay bet can be placed, locking in a profit or loss.

In Running Betfair was an early adopter of in play betting, or betting after the start of the match or race.

Winners Welcome Because Betfair does not take risk, it argued that it welcomed winners, whereas bookmakers would often close the accounts of consistent winners.

The existing bookmakers have fought Betfair ever since. They argued that Betfair offers better prices because Betfair does not contribute to sport, or pay enough tax. They tried to say that offering Lay prices should be illegal as customers should have a bookmaking license. This failed and so they argued that offering Lay markets was enabling corruption; it is much easier to lose a match or race than to win it. Recently attention has turned to spot fixing. This is fixing some small element of the game, such as the number of no balls in one particular over, or the time of the first throw, but which should have little impact on the overall game. Some claim that the increase of in-running betting encourages spot fixing, although probably most famous football spot fixing in the UK is from 1995 [152], long before Betfair started. Co-incidentally 2000 was also the year that football's minimum treble rule was dropped. Fears over match fixing meant that unless a match was on live TV at least a three match accumulator had to be placed [153]. The bookmakers have launched their own in play betting, and a fast growing business it is. There is a controversy here, namely that most live footage is not actually live, but delayed by a few seconds. Some worry that customers do not realise this.

Figure 2.7 shows how the betting odds have improved for the customer over the last few years. The two high street bookmakers (there is a third, but we do not have any data) show a consistently higher over-round than the online competitors. The exception is Interwetten which is Austrian and probably not competing for UK business. In all cases the trend is been falling. Not included on this chart are Betfair, the leading betting exchange, and Pinnacle Sports, a large international bookmaker. We have data for Pinnacle Sports for last year only, which is the lowest of any at 1.020. Betfair has a very low headline over-round for leading markets often around 1.005, but commission needs to be added to this, and the more one bets with Betfair the lower the commission falls. If we consider a new gambler placing one bet per football match then Betfair over-round for the last few years would have been about 1.037 as he has no discount.

To reduce commission from 5% to 3% would need bets of £20 000 for many weeks. For such a serious gambler the over-round is 1.024. Betfair commission can fall as low as 2%, but this requires bets of £360 000 in one week, followed by bets of £60 000 per week to sustain it. If a gambler who is using Betfair bets on more than one runner in the same race, or takes his bet off before the end of the match then his over-round is likely to be lower than that quoted above.

How much of that fall was attributable to Betfair? The Internet was introducing competition from overseas Pinnacle, Victor Chandler (who had moved to Gibraltar by then) and others. The Internet made comparison easy. New bookmakers without the expensive high street infrastructure were competing for business. The Gambling Act 2005 had allowed more advertising. It is difficult to know the answer. The large incumbents have fought to restrict Betfair. Whilst Betfair has better prices than most it is not the best. Surprisingly Betfair recently became a bookmaker. Betfair does now contribute to the horse racing levy.

Many of the innovations listed above are associated with betting exchanges and as that may be fair for fixed odds betting, but these were common features of spread betting long before. Most spread betting companies offered buy and sell, even during many matches. Spread betting was common within financial districts as prices would be displayed on Reuters terminals [154].

Models used by bookmakers have had to cope with in play betting. One football model suitable for in play betting is Dixon and Robinson's [155], as it considers the current state of play.

It is a strange time to be considering betting exchanges. Betfair is currently by far the largest. It does not necessarily have the best prices. It has introduced Premium Fees which are like a tax (of up to 60%) for the most successful customers – winners are not that welcome. And, many would say [147], it has admitted that the exchange model does not work by becoming a bookmaker.

3. Extending Kelly's Fractional Staking to Lay Bets

As we have seen in Section 2.1.2 Breiman's research on the Kelly Criterion concentrated on the general case. Although he proved many properties which hold in general, he considered that finding a solution in the general case would be hard. Kelly, meanwhile, studied one special case only: that in which each event has exactly one winner. In particular, he discussed horse racing where it is possible to bet on a horse and it may or may not win. This was likely to be the only practical scenario available, and certainly the most common.

However, today there is another special case for which a solution would be very useful. For those able to bet in countries where betting exchanges are prevalent, such as the United Kingdom, it is possible to *back* an outcome, or *lay* it. In this chapter we extend Kelly's solution to cover both back bets and lay bets on a single market.

3.1. Lay Bets

Betting exchanges quote two odds for each event, i : the odds to back o_i and the odds to lay O_{L_i} . Recall that laying i is betting that it will not happen. The odds used throughout this paper are decimal odds, such that for each unit of money staked at odds of o_i the backer receives a total of o_i if his bet is considered a winning bet. As this includes the stake, the total net winnings per unit stake is $o_i - 1$. It is usual to quote the laying odds, O_{L_i} , as the equivalent odds the backer of the bet would receive. For every $O_{L_i} - 1$ units risked, should event i not occur, the layer will receive net winnings of one unit. Here it is convenient to consider instead the lay odds as $l_i = \frac{O_{L_i}}{O_{L_i} - 1}$, or the payout with a unit stake if successful for the layer (i.e. i does not win and the lay bet does pay).

We define a *game* as a future contest (or possibly a past contest with an as yet unknown result) upon which we may bet, be it a football match, a horse race, or an election. If Ω is a sample space of *outcomes* for a game, and an *event* is a set of outcomes taken from Ω , we will consider bets on these events, and an event will be considered winning if it contains the winning outcome.

In this language a horse race with only back bets is a mutually disjoint set of outcomes, each event containing one outcome. If we now introduce lay bets then in addition to an event for each outcome, containing that outcome, there is another set of events which is the complement of the first.

If p_i is the probability of event i occurring, x_i is the fraction backed at odds o_i , y_i is the fraction laid at odds l_i and b is the fraction not staked on any of the N possible events, then we must find the b , x_i , and y_i which maximise G , the long run growth rate, where

$$G = \sum_i p_i \log(b + x_i o_i + \sum_{j \neq i} l_j y_j) \quad (3.1)$$

We assume that if a market has odds such that it would be profitable to back and lay an event at the same time without risk of loss, a better would do this first in as large an amount as possible until this dislocation no longer exists. Therefore, we assume that $x_i y_i = 0$, i.e. no bet will be both backed and laid.

Kelly had no decision to make when he chose his constraint as $b + \sum x_i = 1$. When considering back bets on disjoint events the total amount of money bet is the sum of each bet placed. When considering lay bets, or otherwise non-disjoint events, things are more complicated. Consider a horse race with two favourites, both with back odds (O_B) of 9.8 and lay odds (O_L) of 10.0. Placing a £1 lay bet on the first of these will cost the layer £9. If the layer goes on to place a further bet on the second of these horses with a different betting exchange, this too will require £9, a total outlay of £18. If instead the second bet is placed with the same exchange as the first, and if that exchange allows offsetting, then the second bet actually reduces the total liability to £8. In this chapter we maintain the general case assuming that bets may be distributed amongst any number of counter-parties. In subsequent chapters we will be using data from just one, Betfair, and so in those chapters we assume all bets are placed with that one exchange and it allows offsetting.

To find the b , x_i , and y_i which maximise G in Equation 3.1, subject to $b \geq 0$, $x_i \geq 0$, $y_i \geq 0$ and $b + \sum x_i + \sum y_i = 1$ we use the standard KKT [156] method for constrained optimisation; we introduce λ for the equality constraint and μ_b , μ_{x_i} , and μ_{y_i} for the inequality constraints, all of which are greater than or equal to zero. For the Lagrange multiplier $\mathcal{L}(b, \mathbf{x}, \mathbf{y}, \lambda, \mu)$ we now consider each partial derivative in turn:

$$\left(\frac{\partial \mathcal{L}}{\partial b} = 0\right) \quad \sum_i \frac{p_i}{b + x_i o_i + \sum_{j \neq i} l_j y_j} = \lambda \quad (3.2)$$

$$\left(\frac{\partial \mathcal{L}}{\partial x_i} = 0, x_i > 0\right) \quad \frac{p_i o_i}{b + x_i o_i + \sum_{j \neq i} l_j y_j} = \lambda \quad (3.3)$$

$$\left(\frac{\partial \mathcal{L}}{\partial x_i} = 0, x_i = 0\right) \quad \frac{p_i o_i}{b + \sum_{j \neq i} l_j y_j} \leq \lambda \quad (3.4)$$

$$\left(\frac{\partial \mathcal{L}}{\partial y_i} = 0, y_i > 0\right) \quad \sum_j \frac{p_j l_i}{b + x_j o_j + \sum_{k \neq j} l_k y_k} - \frac{p_i l_i}{b + \sum_{k \neq i} l_k y_k} = \lambda \quad (3.5)$$

$$\left(\frac{\partial \mathcal{L}}{\partial y_i} = 0, y_i = 0\right) \quad \sum_j \frac{p_j l_i}{b + x_j o_j + \sum_{k \neq j} l_k y_k} - \frac{p_i l_i}{b + x_i o_i + \sum_{k \neq i} l_k y_k} \leq \lambda \quad (3.6)$$

Defining sets

$$S = \{i \in S : x_i > 0\} \text{ and } T = \{i \in T : y_i > 0\}$$

Making substitution $\beta = b + \sum_k l_k y_k$ and solving equalities 3.2, 3.3, and 3.5 yields

$$\begin{aligned} \lambda &= 1 \\ x_i &= p_i - \frac{\beta}{o_i} \\ y_i &= \frac{\beta}{l_i} - \frac{p_i}{l_i - 1} \text{ and} \\ \beta &= \frac{1 - p_B - p_L}{1 - \sigma_B - \sigma_L} \end{aligned} \quad (3.7)$$

where

$$\begin{aligned} p_B &= \sum_S p_i, \quad p_L = \sum_T p_i, \\ \sigma_B &= \sum_S o_i^{-1}, \text{ and } \sigma_L = \sum_T O_{L_i}^{-1} \end{aligned}$$

Solving the inequalities 3.4, and 3.6 yields

$$p_i o_i \leq \beta \leq p_i O_{L_i} \quad (3.8)$$

if $i \notin S \cup T$, giving a condition to determine those i which have zero fractions in the solution.

In his paper, Kelly is able to partition his bets into two sets, those which will have a positive stake (λ) and those with a zero stake. Kelly was able to find some conditions on these sets. Starting with λ empty he considers all possible bets in order of decreasing expected gain, adding each bet in turn until his conditions are met. The algorithm to find S and T is not as pleasing as Kelly's method as in our case β is not monotonic. In Figure 3.1 we show examples of the four possible cases that we need to consider: no bets

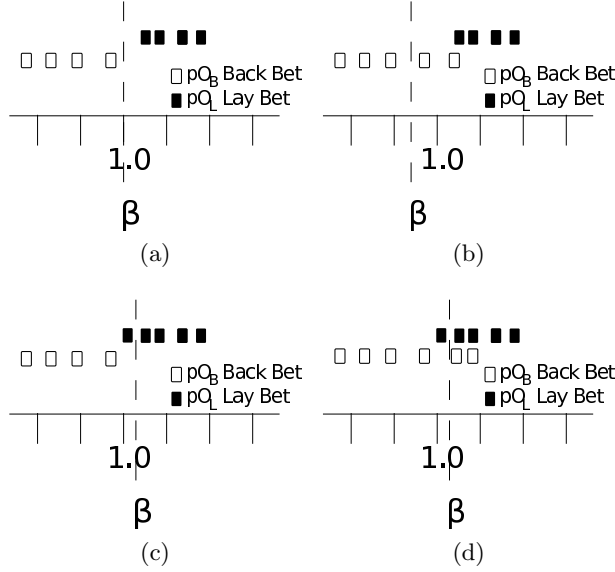


Figure 3.1.: When finding the sets S and T of bets to back and lay respectively we need to consider four cases: (a) both S and T are empty so there are no favourable bets and $\beta = 1$; (b) T is empty so there are only back bets; (c) S is empty so there are only lay bets; and (d) neither S nor T is empty so there are both back and lay bets.

are placed; only back bets are placed; only lay bets are placed and both back and lay bets are placed. We show in Section 4.7 that this is a convex problem and so a unique solution exists. The search to distribute the bets amongst S and T is $O(N \log N)$ as the $p_i o_i$ need to be sorted, and is thus easily tractable. Once sorted we need consider four cases; S and T both empty, T empty only, S empty only, and neither S nor T empty. One possible algorithm is:

1. Initialise $test_{threshold} = 1$; descending = true;
2. Include in S all back bets with $p_i o_i > test_{threshold}$
3. Include in T all lay bets with $p_i o_{L_i} < test_{threshold}$
4. Calculate β from equation 7 above.
5. If condition 8 holds solution is found. Terminate
6. If descending == true; set $test_{threshold}$ just smaller than the largest $p_i o_i \notin S \cup p_i o_{L_i} \in T$ and return to step 2.
7. If $p_i o_i \notin S \cup p_i o_{L_i} \in T$ was empty set descending = false;
8. If descending == false; set $test_{threshold}$ just larger than the smallest $p_i o_i \in S \cup p_i o_{L_i} \notin T$ and return to step 2.

This algorithm steps through all possible partitions until it finds a solution. A simple improvement would be to oscillate about 1 rather than head in one direction first and

Die Roll	Back Odds	Lay Odds
1	5.50	5.70
2	5.99	6.20
3	5.90	6.10
4	6.40	6.60
5	5.80	6.00
6	5.90	6.10

Table 3.1.: The odds available to an imaginary tribe on the roll of a fair die.

then return.

3.2. Simple Example

Consider an as yet undiscovered collective of six tribes, living in a tropical paradise with plenty of resources. These tribes play a six way sport and need to roll a die to select the team which will start. The tribes consider the number 1 auspicious – leading to better marriage prospects for members of the family, whereas the number 4 is ominous. This society has a highly developed gambling market and it is possible to gamble on the roll of the die. The (decimal) odds are shown in Table 3.1.

3.2.1. Naive Kelly

Suppose a passing explorer stumbles across this collective. She understands the Kelly Criterion, but can't bring herself to place bets on an unfavourable event. Because there is one favourable bet only, she calculates the fraction f of her possessions she is happy to wager. She is going to back outcome 4 at odds (o) of 6.4 with probability of winning $p = \frac{1}{6}$, where f is given by

$$f = p - \frac{1-p}{o-1}$$

In this case $f = 0.01234$ (rounding down, never up). When she loses, her wealth reduces by a factor of 0.98766, the wealth relative, when she wins it increases by a factor of 1.06663 ($0.98766 + 6.4 \times 0.01234$). After $6N$ bets we consider the value of her possessions V_{6N} . As $N \rightarrow \infty$ the proportion of winning bets $\rightarrow p$, so

$$V_{6N} \rightarrow V_0 \times (0.98766)^{5N} \times (1.06663)^N.$$

Die Roll	Back Stake	Back Odds	Lay Odds	Wealth Relative
1		5.50	5.70	0.98503
2	0.00222	5.99	6.20	0.99832
3		5.90	6.10	0.98503
4	0.01275	6.40	6.60	1.06663
5		5.80	6.00	0.98503
6		5.90	6.10	0.98503

Table 3.2.: The staking required by the Kelly criterion and the capital multiplier for each outcome of the fair die. The wealth relative is the factor by which wealth changes for that outcome.

As G the average per bet rate of growth of her possessions is given by

$$G = \lim_{N \rightarrow \infty} \frac{1}{N} \log_2 \frac{V_N}{V_0}.$$

$$G = 5.833 \times 10^{-4}$$

3.2.2. Standard Kelly

Imagine, now, that the visiting explorer was happy to follow Kelly's advice; it is optimal after all. In this case f is slightly increased for 4, and there is also a much smaller bet placed on outcome 2. Table 3.2 has the bets, and the wealth relative for each outcome. G has increased to $G = 6.006 \times 10^{-4}$.

3.2.3. Back and Lay Kelly

In the final example she makes use of the lay bets which are available to her. Her bet on 4 is reduced a little and that on 2 considerably; she lays 1. The details are in Table 3.3. When compared with Table 3.2 it is interesting to note that capital fractions for 2 and 4 have barely changed. Whilst 1 has clearly declined to a considerable loss each time the 1 is rolled there is a significantly lower loss for each of the other cases (3, 5 and 6). G has increased further to $G = 8.150 \times 10^{-4}$. Table 3.4 shows a comparison of the of the growth rates achieved in each of the above three cases.

3.3. Football Example – Real Data

3.3.1. Data

Whilst the extension discussed above applies to any set of disjoint back bets with corresponding lay bets, in this section we demonstrate the value of our techniques using

Die Roll	Back Stake	Back Odds	Lay Odds	Lay Stake	Wealth Relative
1		5.50	5.70	0.00829	0.95002
2	0.00017	5.99	6.20		0.99829
3		5.90	6.10		0.99728
4	0.01084	6.40	6.60		1.06665
5		5.80	6.00		0.99728
6		5.90	6.10		0.99728

Table 3.3.: The staking required by our updated Kelly with lay bets and the capital multiplier for each outcome of the fair die. The wealth relative is the factor by which wealth changes for that outcome.

Strategy	Fraction Staked	Growth Rate
Naive Kelly	0.01234	5.833×10^{-4}
Standard Kelly	0.01497	6.006×10^{-4}
Back and Lay Kelly	0.04997	8.150×10^{-4}

Table 3.4.: A comparison of the per bet growth rate achieved by each of the betting strategies for the fair die. The fraction staked is the total fraction of wealth required to be staked at each stage.

football data. These data were collected during the 2009/10 English football season from Betfair. The data were taken from the API close to 10:00 on the day of the match. In each case back and lay odds and amount available were collected for up to four markets: *match odds*; *correct score*; *over/under 2.5 goals*; and *over/under 3.5 goals*. The match odds market concerns betting upon the result (home win, away win or draw). The correct score market is indeed a bet on the final score (16 possible bets on all possible scores with no team scoring more than 3 goals, plus one bet called *other* which covers all scores with at least one team scoring 4 or more goals). Over/under 2.5 and over/under 3.5 are bets on the total score and in particular whether or not the total score will be over 2.5 goals or 3.5 goals respectively.

Out of 380 Premiership matches we collected some data on 350. Power cuts, system crashes, database outages and human error account for the lost 30 matches. Of the 350, ten are missing at least one of the four markets, and whilst a further ten have all four markets present, at least one of the markets has very poor prices. One possible explanation for the very poor prices may be due to Betfair outages. Prior to the closure of Betfair’s servers for maintenance, many prices are removed. We can not be sure why these matches have poor prices. Removing them from this test will likely increase the growth rate of both bets with and without lay bets; whereas leaving them will likely reduce rates in both cases. We have chosen to remove them.

A summary of the data used is seen in Table 3.5. The back size is the amount of money available for a back bet placed at the best price. The mean back size is the mean across all outcomes and all matches in the data. There is a further discussion of this

Market	Number Matches	Mean Back Size(£)	Mean Back Over-rnd	Mean Lay Under-rnd
Match Odds	330	6496	1.0054	0.9925
Correct Score	330	320	1.0206	0.9784
Ov/Und 2.5	330	406	1.0051	0.9936
Ov/Und 3.5	330	392	1.0083	0.9848

Table 3.5.: The betting data on English Premiership matches from the 2009/10 season used in this chapter is taken from Betfair’s API at approximately 10:00 on the day of the match. Included here are only those markets with a back and a lay price for each of the four markets listed.

in Section 5.1. The back over-round is the sum of the implied probabilities for a given market. It can be seen as a measure of how liquid a market is, or how fair it is, although this does not take account of commission. One may think of this (plus commission) as being the same as the over-round charged by a bookmaker. The higher this is, the worse for the better. Alternatively one may consider it to be the cost of unwinding a bet immediately after placing it. To get back 100% of your stake will cost the over-round - perhaps 102%. This over-round (plus commission) needs to be overcome by any model in order for the model to produce a profit. The mean back over-round is the mean across all matches. The over-round for Correct Score bets is higher than for the other three markets listed. Betfair does not allow bets to be placed at all prices. It offers only fixed permitted prices. As there are 17 outcomes in the Correct Score market compared to 2 or 3 in the other markets there will be more ‘rounding errors’ in the market with more outcomes.

3.3.2. Model

In this section we test the new method with the data outlined above, comparing Kelly’s back only with our new method. Our first consideration was to calculate the required probabilities for both methods. There have been many attempts to compare models to bookmakers’ prices, in particular [82] found some success. Indeed, they were able to use historical data to predict the probabilities of future football matches. Since then the bookmakers and users of Betfair have incorporated this research into their pricing models. In fact Dixon founded ATASS, and Coles works for Smartodds; both are companies modelling sporting outcomes. Most of the work of these and other companies has not been published. Using such models no longer produce bets with a positive expectation [41]. Producing a model which predicted football probabilities more accurately than the market would be a substantial piece of research on its own.

We could have used a more recent model such as that suggested by [157], but we feel that as this chapter concerned choosing bet size a simple model was preferred. We decided to use the first model from Maher [75]. This assumes that in any football

Market	Back Only		Back and Lay	
	Final Wealth (£)	Growth Rate	Final Wealth (£)	Growth Rate
Match Odds	200 493.70	0.0182	272 723.19	0.0191
Correct Score	48 431 624.57	0.0348	441 067 600.55	0.0415
Ov/Und 2.5	752 391.57	0.0222	848 074.35	0.0225
Ov/Und 3.5	14 366.95	0.0102	15 513.55	0.0104

Table 3.6.: Comparison of growth rates when i) using back bets only, and ii) using back and lay bets. The data consists of 330 English Premiership matches from the 2009/10 season (described in more detail in Table 3.5). Test assumes an initial wealth of £500 and is allowed to grow. G is the (natural) logarithmic growth rate. In this example the bets are unrestricted in size (this will be considered later in Section 5.1). The probabilities are calculated by fitting the entire season to Maher’s model.

match goals are Poisson distributed and those of the home team are independent from those of the away team. For each team there are two parameters: an attacking goal rate and a defensive factor. There is also an additional scaling factor for home advantage. The Poisson parameter for the home score is the product of the home advantage, the home team attacking rate and the away team defensive factor. Similarly, the Poisson parameter for the away score is the product of the home team defensive and the away team attacking rate. The factors are found by using a MLE fitted to match data. We fitted our parameters to data for the entire season.

3.3.3. Results

To be clear, we are using data not available to the modeller at the time the match began. It would not be possible to repeat this demonstration and make a profit. Table 3.6 has the results of this simple comparison. When calculating growth, Kelly used base 2 logarithms; this fitted well with Shannon’s entropy. The growth rates presented throughout the remainder of this thesis use natural logarithms as is the norm in economics and finance. In each case using lay bets in addition to back bets improved the rate of growth, though in two cases the increase was not significant. This is as expected. The two markets with only two possible outcomes (over/under 2.5 goals and over/under 3.5 goals) are the markets with little improvement. Obviously laying one outcome is the same as backing the other. Indeed, it is the market with the most outcomes which shows the greatest improvement. For some intuition of why this may be the case consider the following. On some occasions the modelled probability of one outcome will be higher than that implied by the betting odds. In this case the bet will be backed. On other occasions the modelled probability of one outcome will be lower. If no lay bets are available the optimisation will tend to back many of the other outcomes. When there are only one or two other outcomes this is not expensive, but as the number of

outcomes grows it does become more expensive. When lay bets are available this is not the case.

As these results use match results which would not have been available at the time the bets would have been placed it is not possible to obtain these results in practice. A model (or skill) superior to that implied by the market price would be required. Also this example assumes that matches take place consecutively. This is far from reality. On a Saturday afternoon nearly all of the lower division and some of the top division football matches take place at the same time. As stated in the introduction there are techniques to deal with this issue, but these would lower the average growth per match. We could have picked the best match of each group of matches with overlapping times, which might have improved growth, but at the cost of a much reduced set of data. As this problem applies equally to all the different methods we are comparing we decided that considering them serially is valid. However there is another reason that the results are unrealistic. In most cases towards the end of the season wealth has grown considerably. It would not be possible to place large enough bets, and so the growth rate would be reduced. A further discussion of this is in Section 5.1.

4. Further Extending Kelly for Betting Exchanges

In the previous chapter we demonstrated the advantages of including lay bets in each of four markets, i.e. an increased growth rate of wealth. In practice we would prefer to combine these markets and place bets on the best combination of all possible markets. There has been some work published in the case of independent games, such as [40] [158], but in our case these results all relate to the same game and are not independent. Breiman did consider such scenarios, indeed he described and proved several of their properties, but as previously noted offered no solution. Whilst it is possible that an analytic solution exists for some combinations of markets, it is likely to be different for each such combination, and indeed for each sport. Here we provide a fast numerical solution. Later in this chapter we demonstrate a further advantage of this numerical solution by modifying it to include commission.

4.1. The general case

Whilst an analytical solution to any problem is preferred, and will usually be more accurate and possibly faster than other approaches, there have been huge improvements in numerical solutions to optimisation problems in the last 20 years. Intentionally exaggerating (but only slightly) Boyd and Vandenberghe [159] state “. . . if you formulate a practical problem as a convex optimisation problem, then you have solved the original problem.” But G , in the previous chapter, (equation Equation 3.1) is a linear combination of logs of linear functions; maximising it is a convex problem. This is also the case with Breiman’s general equation

$$G = \sum_i p_i \log\left(\sum_{j \in A_j} z_j o_j\right)$$

Maximising G is convex and is defined on a closed bounded region and is itself bounded above – therefore there is a solution.

With this in mind, any specific instance of Breiman’s generalisation can be solved using standard optimisation methods such as interior point methods. This is exciting because it allows us to correctly select the size of bets when there are multiple markets on

one match: one betting exchange, for example, has over 40 markets on each English Premiership football match. We used the open source Ipopt [160] as our interior point solver throughout this thesis.

One method to facilitate this is to introduce a matrix \mathbf{M} , for each game, defined so that $m_{ij} = 1$ if event (bet) j is considered to be a win if outcome i occurs, or 0 otherwise. It is convenient to invest all money at each stage rather than having to treat the fraction not wagered, b , separately. As suggested by Breiman, in order to fully invest all money at each stage, we consider an event which wins with all outcomes and with have odds of 1.0 (i.e. it always wins and pays back the stake). So for \mathbf{M} we set the first column such that $o_1 = 1$ and each element of the first column is 1. If p_i is the probability of outcome i , and o_j the odds available on event j , we need to find z_j , the fraction of wealth to be wagered on event j which maximise G where,

$$G = \sum_i p_i \log\left(\sum_j m_{ij} z_j o_j\right).$$

Some of the events j may be lay bets which have had their odds converted to back bets.

To illustrate this we will consider two examples, the first of which is very simple, followed by a slightly more realistic example. Consider a race with four possible winners. In figure 4.1 these are labelled Horses A,B,C and D. Of course this need not be a horse race, but any race with four entrants and one possible winner. Each row of \mathbf{M} represents a possible outcome, the first is Horse A winning, and so on. Each column represents a different possible event, or bet. The first column is special as it represents the set of all possible outcomes and the optimisation will have odds of 1.0 assigned to it. Bets on this event always win, but win nothing. The next four columns are back bets on each of the horses, and the last four columns are lay bets on the respective horses. The column representing a lay bet is the complement of its respective back bet – betting that Horse A will not win is the same as betting that Horse B, Horse C or Horse D will win.

The previous example is straightforward, in fact we can solve it without resorting to numerical methods (although when we introduce commission in Section 4.5 we will use this method even for a single market). As previously stated there are many bets available on the biggest football matches, a few include: Match Odds (betting on the winner); Correct Score (betting on the final score); Over/Under 2.5 (betting if the total number of goals is 0, 1 or 2, or otherwise); To Score (betting if a named player will score a goal); Corners Match Bet (betting which team will have the most corners). Which markets will be combined will depend upon the sophistication of the available model. One class of football model discussed in Section 2.2.3 is that which models the final score. Match Odds bets, Correct Score and Over/Under bets are all determinable using the final score. These three are included in the next example. For a list of common football bets on betting exchanges see Appendix B.

$$\mathbf{M} = \begin{pmatrix}
1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\
1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 & 1 \\
1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 1 \\
1 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0
\end{pmatrix}
\begin{array}{l}
\text{Horse A} \\
\text{Horse B} \\
\text{Horse C} \\
\text{Horse D}
\end{array}$$

CASH
BACK
LAY

Figure 4.1.: Matrix for four horse race.

$$\begin{pmatrix}
1 & 0 & 1 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 0 \\
1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \\
1 & 0 & 0 & 1 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\
1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\
1 & 0 & 1 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 1 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
1 & 1 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0
\end{pmatrix}
\begin{array}{l}
0 - 0 \\
1 - 0 \\
0 - 1 \\
2 - 0 \\
1 - 1 \\
0 - 2 \\
3 - 0
\end{array}$$

CASH
BACK
LAY
BACK
LAY
BACK

MATCH ODDS
UNDER 2.5
CORRECT SCORE....

Figure 4.2.: Excerpt of matrix for football bets. The first column is for cash, the next six are for Match Odds bets (3 back and 3 lay), the following four are for Under/Over 2.5 goals, followed by the first few back bets of Correct Score.

In figure 4.2, as in the previous example, the rows represent outcomes and the columns events (or bets). As with Section 3.1 any lay odds included in the numerical problem would have to be converted into equivalent back odds. The first column again represents the fraction of wealth retained (i.e. not staked). The next three columns are back bets on the match odds markets, home win, draw, and away win respectively. The next three columns are the respective lay bets. Following those are four columns for the Over/Under 2.5 Goals. As with all markets with only two events it is easy to see that laying the second is the same as backing the first. Finally, we can see the first few columns of the Correct Score market, the rest being truncated. Also truncated are many of the rows. In theory we would have a very large number of rows to cover all possible scores. It is not immediately clear what is the maximum number of goals in a game, but it would be large. Fortunately scores of ten and above are very rare (no occurrence in the top four divisions of English league football in the last ten years); therefore, it is usual to create a few outcomes which include all those not otherwise specified.

Parity	Back Odds	Lay Odds
Odd	1.85	1.97
Even	2.07	2.2

Table 4.1.: The odds available on the parity of the roll of a fair die to an imaginary tribe.

Die Roll	Back Stake	Back Odds	Lay Odds	Lay Stake	Wealth Relative
1		5.50	5.70	0.00335	0.95002
2		5.99	6.20		1.02307
3	0.00241	5.90	6.10		0.98333
4	0.00681	6.40	6.60		1.06666
5		5.80	6.00		0.96911
6		5.90	6.10	0.00105	1.01667
Odd		1.85	1.97		
Even	0.02607	2.07	2.20		

Table 4.2.: The staking required by our generalised Kelly for two markets, and the capital multiplier for each outcome of the fair die. The wealth relative is the factor by which wealth changes for that outcome and includes the Odd/Even bets.

4.2. Simple Example (Part II)

Recall the explorer and the tribe from Section 3.2. She notices that another bookmaker offers odds on the parity (odd or even) of the roll of the die, see Table 4.1. As 1 is such an auspicious number tribe members tend to prefer to back odd rolls. To ensure a *balanced book* the bookmaker has to offer very favourable odds on the even roll. Fortunately she has her laptop and therefore can incorporate these odds into her calculation. On a 2009 desktop¹ this calculation takes about 130ms and, therefore, will not significantly impact the explorer's battery. For these combined markets $G = 1.313 \times 10^{-3}$. (The new parity market on its own would have required a fraction of 0.03271 to back Even and would have afforded a growth rate $G = 8.254 \times 10^{-4}$ – the combination is better than either on its own.)

4.3. Football Example

Table 4.3 shows the results of following this procedure with the examples and data from Section 3.3. Again, this staking is based on an initial wealth of £500; the available size in the market is once again not considered. The combination is a considerable improvement over any single market on its own. We should emphasise that the risk

¹Core 2 Duo CPU @ 3.00GHz, 8GB RAM, Ubuntu 12.04, running the Ipopt Java interface through Eclipse

Market	Back Only		Back and Lay	
	Final Wealth (£)	Growth Rate	Final Wealth (£)	Growth Rate
Match Odds	200 493.70	0.0182	272 723.19	0.0191
Over/Under 2.5	752 391.57	0.0222	848 074.35	0.0225
Over/Under 3.5	14 366.95	0.0102	15 513.55	0.0104
Combined	39 097 529.28	0.0341	47 422 608.99	0.0347

Table 4.3.: Three results repeated from Table 3.6 with the new result of combining all three. As before, the data consists of 330 English Premiership matches from the 2009/10 season (see Table 3.5). Test assumes an initial wealth of £500 and is allowed to grow. In this example the bets are unrestricted in size; this will be considered later, in Section 5.1. The probabilities are calculated by fitting the entire season to Maher’s model.

taken on the combination is the same as that taken on any one of the single markets. The next section will explain why at this stage we do not show the results of combining all four markets.

4.4. Ignoring Commission

Table 4.4 has the details of a seemingly good set of bets, except that when commission is correctly taken into account this set is actually very poor. The key problem is that commission on betting exchanges, such as Betfair, is accounted for on a per market basis, rather than a per match basis. When betting on multiple markets simultaneously this might leave the better exposed should some outcomes yield large profits in some markets and large losses in another. This is particularly true when there is a *backwardation* across markets. We use backwardation to mean that excluding commission it is possible to place bets in several markets such that a profit is guaranteed. This is the situation found here. In essence the bets are betting that there will be more than 2.5 goals in the Over/Under market, whilst betting that there will be fewer than 2.5 in the Correct Score and Match Odds markets. Backwardations small enough not to be tradable once commission is included are surprisingly common. There are over 20 in the 330 matches considered in this thesis.

The score in this particular match was 1 - 0. The gross winnings for this result of the bets in Table 4.4 total £52.27. However this small profit consists of a large loss (£2024.34) in one market and a slightly larger profit (£2076.61) in the other markets. Commission of 5% on any profitable market reduces the previous profit to a loss of £51.86. The Kelly method maximised the expected log, but if we examine the expected result of this one bet we will see that the result actually obtained is typical. The expected profit (without considering commission) is £20.91, once commission is included the expected profit reduces to a loss of £72.33.

In Section 4.5 we propose a solution to this problem. Following this method the payout (after deducting commission) would have been £30.70, and the expected return would have been £7.44. The bets produced by this method are shown in Table 4.5.

4.5. Commission

As previously stated betting exchanges are peer to peer, charging commission unlike bookmakers who adjust their prices to make a profit. In the previous examples the fractions to be wagered ignored commission, although the accounting of profits and losses in the tables did include commission at 5%. One obvious consequence of this is that the strategies above were all slightly over-betting. As the actual returns would have been lower than those implied by the odds used, the required fractions should have been slightly lower. Initially, we decided that this could be corrected by reducing the stake by a small amount and need not be of great concern. Only after discovering the bet in Section 4.4 did we realise that was not acceptable.

The principal issue is not the size of the commission, but that it is paid on each market. Generally, it is not possible to offset losses in one market against profits on other markets relating to the same match. The match underlying the bets in Table 4.4 ended 1 - 0. The total gross profit was £50, but this concealed a gain of over £2 000 in one market and a similar sized loss in another market. The commission on this profit is over £100, or double the anticipated profit. Fortunately with the numerical version of our new methods we are able to adjust the problem. This adjusted problem has a profit of over £30.

The intuition behind what is occurring in Section 4.4 is that the system thinks it has found a backwardation. It attempts to back all the low scores in the Correct Score market, whilst simultaneously backing Over 2.5 Goals in the Over Under 2.5 market. But as commission is charged on the profit of each market the backwardation does not really exist. We need to modify our objective function to take account of commission.

As stated in Section 2.3 we are using the Ipopt optimisation library. This is a fast, large scale interior point solver written in C++, with several interfaces including Matlab, Java, AMPL, and R. One condition is that the objective function (and any constraints) need to be twice differentiable. But commission is not a smooth function, being $C\%$ of any profit and 0% of any loss. Whilst there are other methods for such problems we decided to use a standard approximation to the Heaviside function.

As before we are trying to maximise G where

$$G = \sum_i p_i \log\left(\sum_j m_{ij} z_j o_j\right)$$

We now introduce our commission function $\mathcal{C}(x)$ with commission rate C

$$\mathcal{C}(x) = \begin{cases} x(1 - C) & \text{if } x \geq 0 \\ x & \text{if } x < 0 \end{cases}$$

so if $H(x)$ is the unit step Heaviside function

$$\mathcal{C}(x) = x - CH(x)$$

We choose to approximate H by H^* where

$$H^*(x) = \frac{1}{1 + e^{-2kx}}$$

and k is a constant which determines the closeness of the approximation when x is close to zero. Introducing a new matrix \mathbf{N} where $n_{lj} = 1$ if bet j is in market l and $n_{lj} = 0$ otherwise. We can now rewrite G as

$$G = \sum_i p_i \log\left(\sum_j m_{ij} z_j o_j - C \sum_l H^*\left(\sum_j n_{lj} m_{ij} z_j o_j\right)\right).$$

By introducing this approximation we have sacrificed convexity. The status quo was, however, not acceptable. When trying to maximise all four markets in the previous section, on 64 occasions (out of 330) a match with a positive expected return was changed to one with a negative expected return. Whilst we are not maximising the expected return, a positive expected return, or alternatively a favourable bet is a minimum requirement. More work needs to be done to see if the second order sufficient conditions are being met.

The results seen in Table 4.6 show a small improvement in growth rate for three of the individual markets, and a small reduction for the fourth. However, significantly we are able to combine all four markets together. For comparison we also re-calculate a combination of three markets and this, too, is improved.

4.6. Simple Example (Part III)

Consider again our gambling explorer. Imagine instead of finding two bookmakers she found a betting market place (presumably open outcry rather than electronic). This betting market charges each member a commission of 5% on the net winnings of each market, and nothing on a losing market. The odds for the rolled number are unchanged from Table 3.1, but the gross odds for the parity market have improved. They are visible in Table 4.7.

Not yet having incorporated the results from Section 4.5 into her solver, the explorer

decides to ignore commission, thinking that if she reduces her bets slightly she should be safe. Table 4.8 shows the results. The explorer's entire wealth is staked. However, not much is actually at risk as about half of her wealth is staked on the even numbers in the parity market, whilst about half is staked on the odd numbers in the die roll market. Growth is much higher than before at $G = 2.857 \times 10^{-3}$, ignoring commission. When we include commission, however, we can see the problem. Whichever number is chosen in the roll the betting above produces a large win in one of the two markets and a loss in the other. As commission has to be paid on the win in each case, but with no account taken of the loss this is a poor bet. Table 4.9 shows the new wealth relatives, the multiple by which wealth changes on each roll, adjusted for commission. Ignoring commission has turned an expected gain into an expected loss.

Fortunately, our explorer is aware of the problem commission causes when some markets have large wins and other markets have large losses; therefore, she decides to implement the techniques from Chapter 3. Table 4.10 shows the results. You will recall that we have added a commission on all winning bets but we only improved the odds on one of the markets; we are not surprised that the growth rate is not as high as in the bookmaker example, being $G = 3.862 \times 10^{-4}$.

4.7. Constraint Choices

Kelly, considering a disjoint set of back bets, had a simple constraint; the sum of all stakes plus cash retained should be one. This would be the correct constraint today in a similar situation. As explained in Section 3.1, once we are considering overlapping bets (such as lay bets) the position is not as simple. Two lay bets in one market cannot both lose. The cash required to place a collection of bets on a betting exchange is set to be equal to the worst possible loss on that market. Placing two lay bets at the same odds in one market requires less money than placing either one on its own, but only if both are placed with the same exchange.

Kelly could have placed his bets with any bookmaker. He could move around the track to find the best odds for each horse and use those odds for his calculation. If we do so with lay bets, not only will we not benefit from the cash offsetting described above, but we might pay a higher rate of commission. If one of the lay bets loses and the other wins we will pay commission on the gross winnings, not the net winnings which would be the case if both were with the same exchange.

Correcting the second of these problems is straightforward. We can modify \mathbf{M} and \mathbf{N} to include markets from multiple exchanges as though they are different markets. The relevant columns in \mathbf{M} will be the same, but there will be no commission netting. Similarly, bookmakers' prices with no commission could be included with the commission rate set to zero. At the outset of this thesis most of the volume of betting was concentrated in

one exchange, Betfair [161]. Now another is starting to increase volumes traded. More detailed work is required on the impact of splitting betting between multiple exchanges.

The remainder of this thesis assumes that all bets are placed on the same exchange.

Type of Bet	Outcome	Bet	Size	Odds
Correct score	(0 - 0)	BACK	£172.11	22.00
Correct score	(0 - 1)	BACK	£43.39	55.00
Correct score	(0 - 2)	BACK	£12.06	180.00
Correct score	(0 - 3)	LAY	£1.94	800.00
Correct score	(1 - 0)	BACK	£421.82	7.60
Correct score	(1 - 1)	BACK	£231.09	15.50
Correct score	(1 - 2)	LAY	£13.79	90.00
Correct score	(1 - 3)	LAY	£2.47	550.00
Correct score	(2 - 0)	BACK	£514.93	6.20
Correct score	(2 - 1)	LAY	£37.20	12.00
Correct score	(2 - 3)	LAY	£3.44	410.00
Correct score	(3 - 0)	LAY	£39.08	8.20
Correct score	(3 - 1)	LAY	£34.08	14.00
Correct score	(3 - 2)	LAY	£9.92	55.00
Correct score	(3 - 3)	LAY	£0.81	250.00
Correct score	(Any Unquoted)	LAY	£97.93	4.20
Match odds	(Birmingham)	BACK	£39.03	25.00
Match odds	(The Draw)	LAY	£64.55	7.20
Over/Under 2.5	(Over 2.5)	BACK	£2024.34	1.75

Table 4.4.: Output of naive optimiser, which ignored commission. The match is Manchester United vs Birmingham City on 16th August 2009.

Type of Bet	Outcome	Bet direction	Size	Odds
Correct score	(0 - 0)	BACK	£8.15	22.00
Correct score	(0 - 1)	BACK	£2.00	55.00
Correct score	(1 - 0)	BACK	£4.22	7.60
Correct score	(2 - 0)	BACK	£3.02	6.20
Correct score	(2 - 1)	LAY	£4.37	12.00
Correct score	(3 - 0)	BACK	£4.33	8.00
Correct score	(3 - 1)	LAY	£5.96	14.00
Correct score	(3 - 2)	LAY	£2.78	55.00
Correct score	(Any Unquoted)	LAY	£5.89	4.20
Match odds	(Birmingham City)	BACK	£3.79	25.00
Match odds	(The Draw)	LAY	£3.95	7.20

Table 4.5.: Output of commission adjusted optimiser. Bets are significantly smaller than in Table 4.4. The match is Manchester United vs Birmingham City on 16th August 2009.

Market	Back Only		Back and Lay	
	Final Wealth (£)	G'th Rate	Final Wealth (£)	G'th Rate
Match Odds	272 723.19	0.0191	322 369.27	0.0196
Correct Score	441 067 600.55	0.0415	753 151 809.35	0.0431
Ov/Und 2.5	848 074.35	0.0225	722 352.72	0.0220
Ov/Und 3.5	15 513.55	0.0104	22 145.13	0.0115
Combined 3	47 422 608.99	0.0347	65 313 782.08	0.0357
Combined 4	-	-	4 341 762 884.14	0.0484

Table 4.6.: A comparison of the results from Table 4.3 with the results using the commission adjusted objective function. In both cases, when calculating the wealth, commission is accounted for. Back and lay bets are placed. Combined 4 indicates all four markets combined together, combined 3 excludes correct score. Data as before.

Parity	Back Odds	Lay Odds
Odd	1.88	1.94
Even	2.1	2.17

Table 4.7.: The exchange odds available on the parity of the roll of a fair die to an imaginary tribe.

Die Roll	Back Stake	Back Odds	Lay Odds	Lay Stake	Wealth Relative
1	0.06110	5.50	5.70		0.91699
2		5.99	6.20	0.09147	1.03369
3	0.06826	5.90	6.10		0.98367
4		6.40	6.60	0.07582	1.10039
5	0.06657	5.80	6.00		0.96705
6		5.90	6.10	0.09523	1.01990
Odd		1.88	1.94		
Even	0.48565	2.10	2.17		

Table 4.8.: The staking required by the generalised Kelly for two markets and the capital multiplier for each outcome of the fair die, when commission is ignored.

Die Roll	Wealth Relative	
1	0.91699	0.89686
2	1.03369	1.00698
3	0.98367	0.96021
4	1.10039	1.07368
5	0.96705	0.94441
6	1.01990	0.99319

Table 4.9.: A comparison of the wealth relatives of bets calculated without taking account of commission, and the wealth relatives when commission is taken into account.

Die Roll	Back Stake	Back Odds	Lay Odds	Lay Stake	Wealth Relative
1		5.50	5.70	0.00258	0.96870
2		5.99	6.20		1.01719
3		5.90	6.10		0.98341
4	0.00265	6.40	6.60		1.03331
5		5.80	6.00		0.98341
6		5.90	6.10		1.01719
Odd		1.88	1.94		
Even	0.01652	2.10	2.17		

Table 4.10.: The staking required by the generalised Kelly for two markets bets and the capital multiplier for each outcome of the fair die, when commission is included in the optimisation. The wealth relatives have had commission deducted.

5. Analysing Kelly's Assumptions

When we compare the mean bet size available (Table 3.5) with bet sizes needed to produce the final wealth seen in Table 4.6 it is clear that the total wealth obtained is unrealistic. Even if we found a suitable model which gave good results it would not have been possible to place bets of the required size in many of the previous examples. In this chapter we examine the impact on the growth of wealth when using bet sizes available on Betfair. We also discuss another of Kelly's assumptions, namely that wealth is infinitesimally divisible. We investigate the probability that the Kelly Criterion fails, because the desired bet is lower than the minimum bet size.

5.1. Maximum Effective Bet Size

Account must be taken of the size of bets available in the exchange. Further developing our numerical solution we added a constraint that no bet may be larger than the size available in the market at that time. A summary of the results are given in Table 5.1.

The results indeed show a much lower terminal wealth. The lowest result is that of Over/Under 2.5 goals. This too is unsurprising as we can see in Table 3.5 Over/Under 2.5 goals had much lower average size than Match Odds or Correct Score. One further point to note is that as the available size is restricted, the number of bets placed increases. For example, the combined 4 strategy from Table 4.6 involved 5 714 bets, whereas the similar strategy in Table 5.1 involved 7 143 bets. The maximum possible would be 7 920. When size becomes a constraint previously less attractive bets are placed as they

Market	Back Only		Back and Lay	
	Final Wealth (£)	Growth Rate	Final Wealth (£)	Growth Rate
Match Odds	52 810.03	0.0141	87 627.02	0.0157
Correct score	118 051.10	0.0166	208 876.48	0.0183
Ov/Und 2.5	5 546.54	0.0073	15 963.97	0.0105
Combined 3	-	-	313 510.71	0.0195
Combined 4	-	-	546 655.60	0.0212

Table 5.1.: Results from using constrained sizes: otherwise as Table 4.6. The fraction staked at each stage is restricted to that available at the time. A substantial fall in terminal wealth and growth rates is seen.

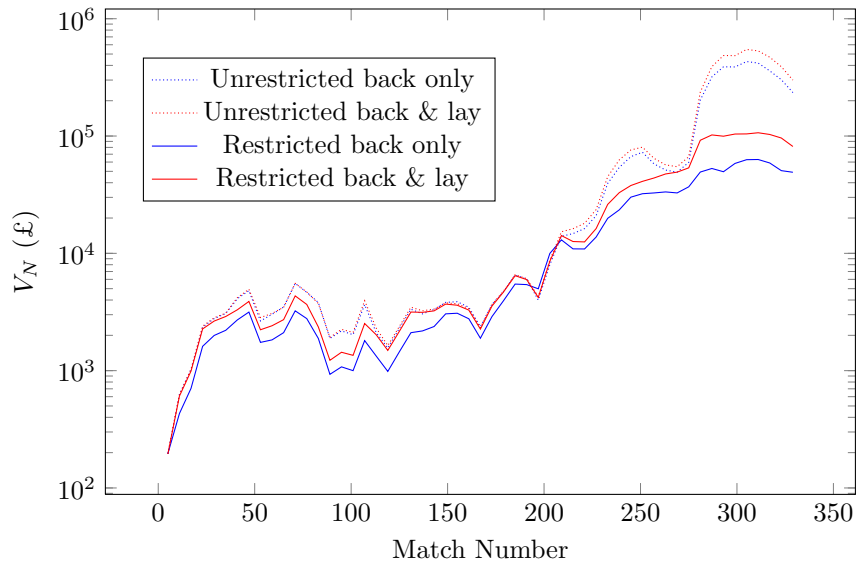


Figure 5.1.: Match odds market. When wealth goes above £10 000 market size appears to be an impediment to growth.

are available, once the maximum use is made of the best bets.

Figures 5.1, 5.2, & 5.3 are graphs of wealth against time for both the unconstrained and constrained (by market size) methods for match odds, correct score and Over/Under 2.5 goals respectively. These agree with the data in the table in that the greatest difference is encountered in the Over/Under 2.5 goals market. Figure 5.4 shows wealth for all four markets in the constrained and unconstrained cases. Overall the graphs suggest that wealth approaching £100,000 is the point at which market size becomes an issue when combining all markets. In practice this value might be higher. The data used was collected at 10:00 on the day of the match. There is evidence that liquidity increases towards the start of an event. Examining the data we collected confirms this – the mean size available to back in the hour before the match starts is just over £16,000. We captured the best back and the best lay price. The API provides, in addition, the second and third best of each. These could be added to the numerical optimisation with ease, as could prices from other exchanges or bookmakers. This should more than double the available size. Further work, or greater experience, would be needed to see if such a strategy was optimal, or whether it may be better to place bets at the best price now and then wait a while and then place more. We discuss placing bets after having already placed bets, in Section 6.1, but we think it unlikely that we could know the impact of placing large bets without doing so and measuring it. It does not imagine that anyone will do so in the near future and publish the results.




Figure 5.2.: Correct score market. When wealth goes above £50 000 market size appears to be an impediment to growth.

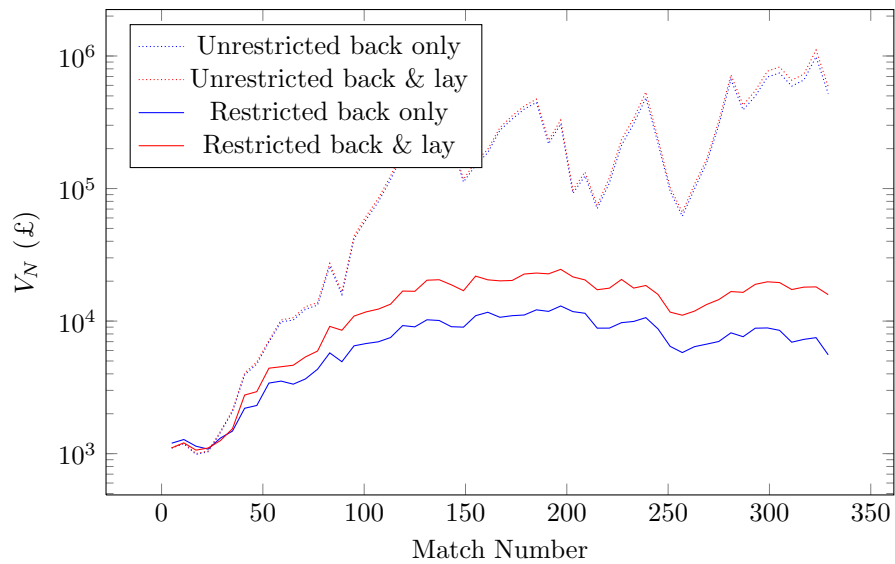


Figure 5.3.: Over under 2.5 goals market has lower volumes than the previous two. When wealth goes above £2 000 market size appears to be an impediment to growth.

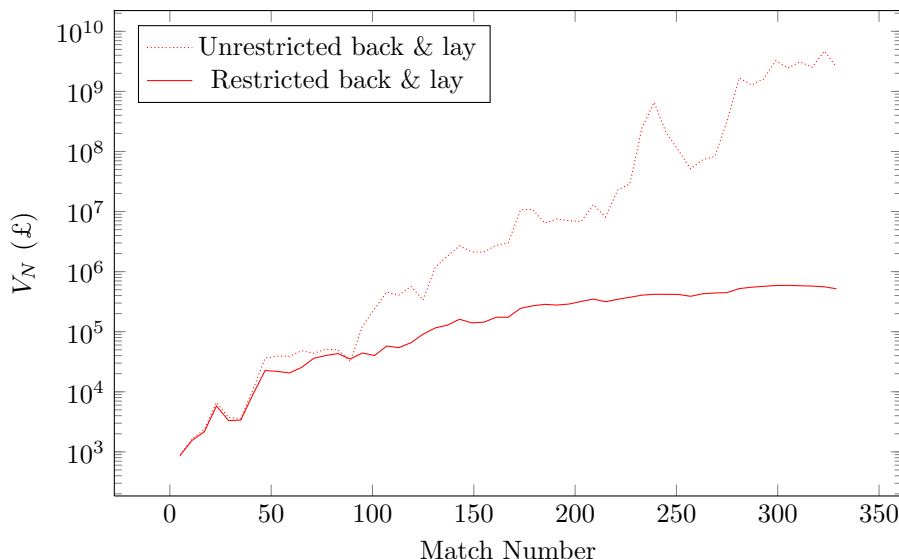


Figure 5.4.: All 4 markets combined. When wealth goes above £30 000 market size appears to be an impediment to growth.

5.2. Probability of Large Losses

Having examined the impact on the growth rate of the maximum bet size, above, in Section 5.3 we will turn our attention to the minimum bet size. One occasion when the minimum bet size will have an impact is the situation following substantial losses, in particular losses so great that the next required bet is below the minimum bet size. To aid with the analysis of that we shall first study the probability of large losses.

5.2.1. Kelly Cannot Go 'Bust', but ...

With the assumption of the infinitesimal division of money, and an investment with exponential growth, then eventually the Kelly Criterion will get ahead of any essentially different strategy [10]. There is no possibility of *going bust*, as no matter how much one loses there is always a tiny fraction available to remain in the game. We examine this in more detail below; here we consider if we are happy with possibly very large losses along the way. In this section we will define *going bust* as losing a very large fraction of the initial wealth, perhaps 90%, or 99%, typically the former. In each case we will define the probability of this loss.

We need a method of calculating the probability of wealth V at any stage being lower than ψV_0 , where V_0 is the initial wealth and ψ is the fraction of that initial wealth below which we consider losses too great. This problem looks somewhat similar to a random walk. However, we have not been able to fit our problem into standard random walk solutions. The issue we need to resolve is the drift of wealth higher with the expected gain. After every win wealth increases by more than it loses during the

equivalent number of losses. We expect that a better solution to the one below exists. This solution has poor convergence in some cases and the computational time needed to calculate each successive term increases as the terms get larger.

We consider P_k , the probability of wealth falling below ψV_0 after precisely k wins. If p is the probability of a win, o is the odds available on that favourable ($o \times p > 1$) bet, and if f is the optimal fraction of wealth staked on this bet given by,

$$f = p - \frac{1-p}{o-1}$$

then the wealth factor for a win W_{win} is $1 + (o-1)f$ and the wealth factor for a loss W_{loss} is $1 - f$. If wealth is to have fallen to ψ after k wins there must be L_k losses where L_k is the smallest integer such that

$$W_{win}^k \cdot W_{loss}^{L_k} < \psi, \text{ or}$$

$$L_k = \left\lceil \frac{\log(\psi) - k \log(W_{win})}{\log(W_{loss})} \right\rceil. \quad (5.1)$$

So P_k is given by

$$P_k = N_k p^k (1-p)^{L_k},$$

where N_k is the number of permutations of k wins and L_k losses. Naively one might consider $N_k = \binom{k+L_k}{k}$, but this is too large; it includes permutations which would result in the loss before the final win. For example, to be considered bust without a single win may take 67 consecutive losses, with a single win perhaps 70 losses are required, in such a case N_1 must exclude all combinations which have that win in the final 3 positions (as the player would already be considered bust). This leads to

$$N_k = \binom{k+L_k}{k} - \sum_{i=0}^{n-1} N_i \binom{L_k - L_i + k - i}{k - i}.$$

Hence the probability, p_ψ , of capital falling below ψ is given by

$$p_\psi = \sum_{k=0}^{\infty} N_k p^k (1-p)^{L_k}. \quad (5.2)$$

5.2.2. Probability of Substantial loss: Illustrative Examples

Figure 5.5 shows some examples of the probability, p_ψ of a 90% loss for various values of the probability of winning a bet, p and the expected value of that bet, $p \times o$. It is clear that in these cases as p increases the probability of substantial loss decreases,

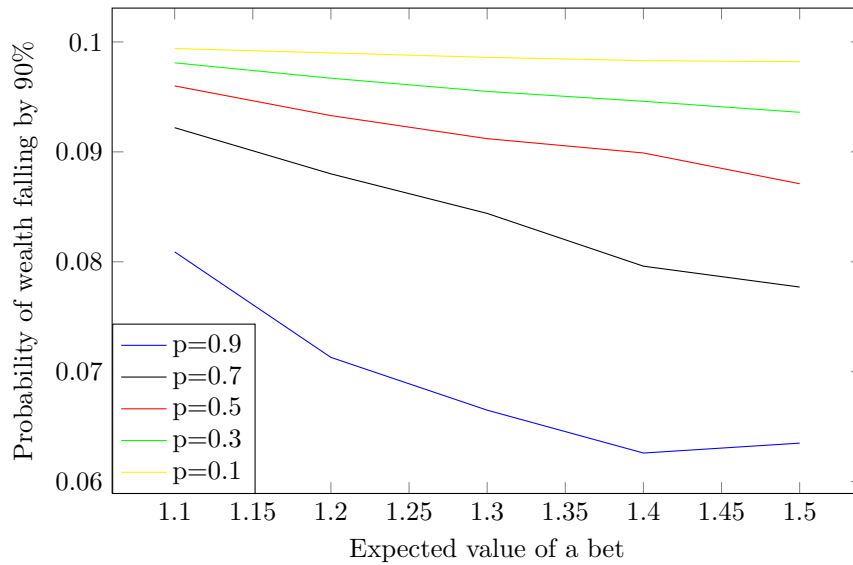


Figure 5.5.: A graph of probability of substantial loss versus expected value of the bet being wagered for different bets with different probabilities of winning. In each case, as the expected value increases the probability of large loss decreases. When more favourable bets are won a greater number of losing bets will be required to bring about a substantial loss.




Figure 5.6.: A graph of probability of substantial loss versus the underlying bet probability for different expected values. In each case, as the bet probability increases the probability of large loss decreases. When more favourable bets are won a greater number of losing bets will be required to bring about a substantial loss.

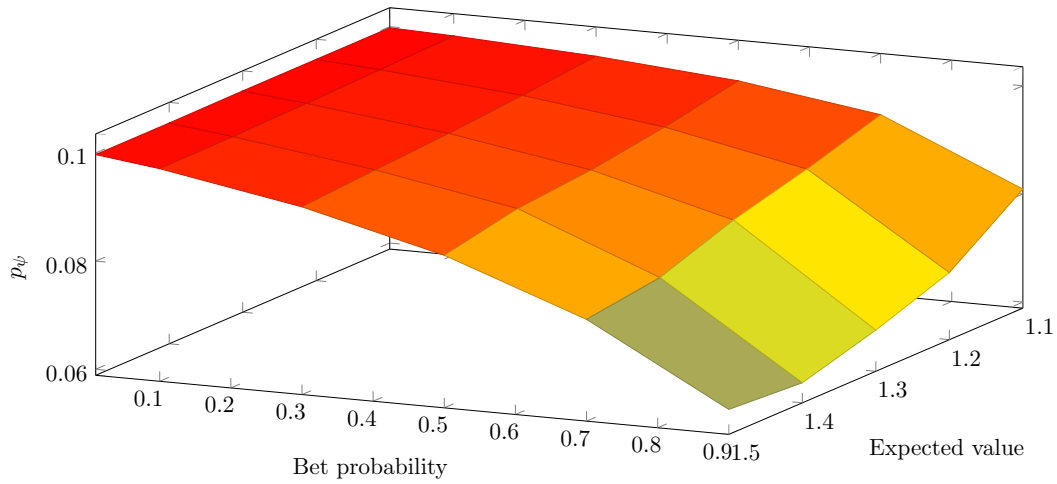


Figure 5.7.: A graph of probability of substantial loss, p_ψ , versus both expected value and probability of winning the bet being wagered.




Figure 5.8.: A graph of the number of iterations used to calculate the probability of substantial loss versus expected value and the probability of the bet being wagered. In each case, as the expected value increases the probability of large loss decreases. When more favourable bets are won a greater number of losing bets will be required to bring about a substantial loss.

p_ψ	ψ
0.25	0.232
0.10	0.0933
0.01	0.00932

Table 5.2.: The probability, p_ψ , of wealth falling below a fraction, ψ , of initial wealth appears to be close to ψ , at least for the moderate values of p and o chosen here ($p = 0.5, o \times p = 1.2$).

and, similarly, as $o \times p$ increases then p_ψ falls. What is striking is that many of these results are just below 10% and we chose $psi = 0.1$. Table 5.2 suggests that this result might not be a coincidence. It shows values of p_ψ for various values of ψ , with $p = 0.5$, and $o \times p = 1.2$. Figure 5.6 shows the same data as shown in Figure 5.5, but this time arranged as p_ψ versus p for various values of expected value. Again p_ψ is the probability of a 90% loss of wealth. As before we can see that as p increases the probability of substantial loss decreases, but now it is obvious that this relationship is not linear. Figure 5.7 combines the other two graphs into one (but note for ease of viewing the $o \times p$ axis has been reversed).

The results above suggest that there might be a simpler form for p_ψ than Equation 5.2. Whilst we have found several approximations which work well for $p < 0.8$ and $1.1 < p \times o < 1.5$ this range is too restrictive to be of much use. Common bets in this thesis involve laying events. As these might have probabilities below 0.1 then the probability of winning a lay bet is over 0.9. Similarly, Thorp [11] states that the expected value of an optimally played game of Blackjack (in 1969 without card counting) is about 1.001 per hand. It is when the expected value is close to 1 that we are in need of an alternative.

5.2.3. Probability of Substantial loss: Convergence

When p is large, and when $o \times p \gg 1$ Equation 5.2 converges reasonably quickly, but as p approaches 0, or $o \times p$ approaches 1 the number of iterations required for a given precision increases markedly. Not only does the number of iterations need to increase as p reduces, but each iteration increases in execution time. For example if we fix $\psi = 0.01$, and $o \times p = 1.1$, then calculating p_ψ for $p = 0.5$ needs 2853 iterations and takes 23 seconds, making only one change to $p = 0.01$ doubles the required number of iterations to 5735, but the time taken increases to over 7 hours. (This is not as fair a comparison as it could be: smaller values of the binomial coefficient are cached, but the machine used does not contain enough memory to cache all values. The speed gain will improve shorter runs disproportionately.) As p decreases the optimal fraction also decreases, so many more losses are needed to reduce wealth below ψ . Further compounding this is that the odds must have increased (as $o \times p > 1$) and so each win requires many more losses. For example if $p = \frac{1}{3}$ then when $k = 5000$ $L_k = L_{5000} \approx 10\ 050$. When $p = \frac{1}{10}$, however, if $k = 5000$ then $L_k \approx 55\ 000$. As the binomial coefficient looks like $\binom{k+L_k}{k}$ it

is easy to see the increased calculation time. Figure 5.8 graphs the number of iterations used to calculate the data in the previous subsection.

When calculating the values above we needed to decide where it was safe to terminate the series. Table 5.3 gives some intuition for one particular case. In that case the additional value of p_ψ provided by each iteration falls by a factor of 10 every 480 or so iterations. If this is true in general then we should be able to calculate p_ψ with arbitrary precision, by choosing appropriate k_{max} .

The lead term in Equation 5.2 is the Binomial Distribution,

$$Pr(X = k) = \binom{n}{k} p^k (1-p)^{n-k},$$

but as n increases, by Central Limit Theorem, this tends towards the Normal Distribution,

$$Pr(X = k) = \frac{1}{\sigma\sqrt{2\pi}} \exp^{-\frac{1}{2}\left(\frac{k-\mu}{\sigma}\right)^2}, \quad (5.3)$$

where $\mu = np$, and $\sigma^2 = np(p-1)$. In this case $n = k + L_k$, where L_k is given by Equation 5.1. In the Normal approximation we no longer need L_k to be an integer so we can rewrite Equation 5.1 as $L_k = a + kb$, where $a = \frac{\log(\psi)}{\log(W_{loss})}$, and $b = -\frac{\log(W_{win})}{\log(W_{loss})}$. Substituting these into Equation 5.3 we get

$$\begin{aligned} Pr(X = k) &\approx \frac{1}{\sqrt{2\pi p(1-p)(a+kb+k)}} \exp^{-\frac{1}{2}\left(\frac{k-p(a+kb+k)}{\sqrt{(a+kb+k)p(1-p)}}\right)^2}, \text{ rearranging} \\ &\approx \frac{1}{\sqrt{2\pi p(1-p)(a+kb+k)}} \exp^{-\frac{1}{2}\frac{(k(1-p-pb)-pa)^2}{p(1-p)(k(b+1)+a)}}. \end{aligned}$$

We are considering large values of k . p , a , and b are constant, and b is positive so eventually $\frac{1}{\sqrt{2\pi p(1-p)(a+kb+k)}}$ will reduce below 1 and be decreasing. Providing

$$|k(1-p-pb)| \gg pa, \text{ then}$$

$$Pr(X = k) < \exp^{-\frac{1}{2}k\frac{(1-p-pb)^2}{p(1-p)(b+1)}}.$$

For the data in Table 5.3 using the above equation we would expect the value to reduce by a factor of 10 every 542 iterations, which fits well with the observation.

The termination criteria that we use are based on these results. The program continues to iterate until:-

- $n = (k + L_k)$ is bigger than both $\frac{10}{p}$, and $\frac{10}{1-p}$ this ensures that the normal approximation is valid;
- $k > 4 \times \left| \frac{pa}{1-p-pb} \right|$, with the 4 being rather arbitrary; and

Scale	k	Difference
10^{-7}	1217	444
10^{-8}	1661	461
10^{-9}	2122	474
10^{-10}	2596	478
10^{-11}	3074	486
10^{-12}	3560	470
10^{-13}	4030	520
10^{-14}	4550	500
10^{-15}	5050	490
10^{-16}	5540	530
10^{-17}	6070	490
10^{-18}	6560	500
10^{-19}	7060	

Table 5.3.: Table showing the decline of the incremental value of each iteration of Equation 5.2, where $p = \frac{1}{3}$, and $o = 3.5$. The iteration number is the last time that scale is exceeded so, for example, all iterations with $k > 1217$ are smaller than 10^{-7} . This suggests that for large enough k the series becomes roughly geometric, dividing by 10 every 500 iterations.

- given a desired precision ϵ , the value of the k^{th} iteration it_k and the cumulative value of p_{ψ_k} are such that $\epsilon p_{\psi_k} > it_k \cdot 2^{\frac{p(1-p)(b+1)}{(1-p-pb)^2}}$.

It is possible that a more accurate or faster method would be needed by other researchers, and this is discussed in Further Research (Section 8.3). However, we needed this result to investigate the problem of the minimum bet size and the results above are adequate for our needs.

5.3. Minimum Bet Size

One of the fundamental assumptions that Kelly made is that money is infinitesimally divisible. Breiman also relied on this in his work on optimality. If it is always possible to place a bet, no matter how small then one can carry on following the Kelly Criterion and eventually the exponential growth will increase wealth substantially. However, it would be very surprising if you were able to place a bet, in the UK, below £0.01, and in many cases the minimum would be £1. On Betfair the minimum bet is £2, and that is the minimum we shall use for the rest of this section.

We shall consider three different cases when the minimum bet might be a problem. If the only bet, or the largest bet of several, is less than the minimum bet size, we shall examine what happens if we round up to the minimum. We will also examine the situation when the main bet is larger than the minimum, but an additional bet is required that is below the minimum. After that we shall build on the work of the previous section and examine initial losses, reducing wealth so much that a bet is no




Figure 5.9.: A graph of the minimum wealth (£) needed to be able to place a bet of £2 at the given odds and probability. The minimum probability plotted is 0.001, and the minimum expected value is 1.001.

longer possible.

5.3.1. Minimum Bet Size: Primary

Figure 5.9 shows the required wealth needed so that the first bet placed is larger than the minimum bet, for different probabilities and expected values. As the probability gets close to zero, or the expected value close to one the wealth required increases.

There are three causes of the required fraction being lower than the minimum bet: the available wealth might be low, such that even a very desirable bet with a high fraction equates to a small amount of money; the probability of a bet winning might be very low, creating a low fraction; or the expected value might be very low, again generating a small fraction. Obviously it might be a combination of two or three of these.

Ideally, the gambler/investor would have enough money so that small bets will not arise. Failing that what should one do? If the strategy is such that the probabilities and odds vary (perhaps betting on modelled football predictions, sometimes backing the home team with p over 0.5 and sometimes backing a weak away team with $p \approx 0.15$) then maybe the best strategy is to wait, and only place those bets naturally greater than the minimum bet. In time, wealth will grow and a higher proportion of bets will be included. If the nature of the underlying game is such that the probabilities do not change then the answer is less straightforward.

If faced with a bet below the minimum size and little chance that the next will be higher the obvious decision is whether or not to round up. If the requested bet is £1.99 and the minimum is £2, rounding up might seem like an obvious choice, but what if the

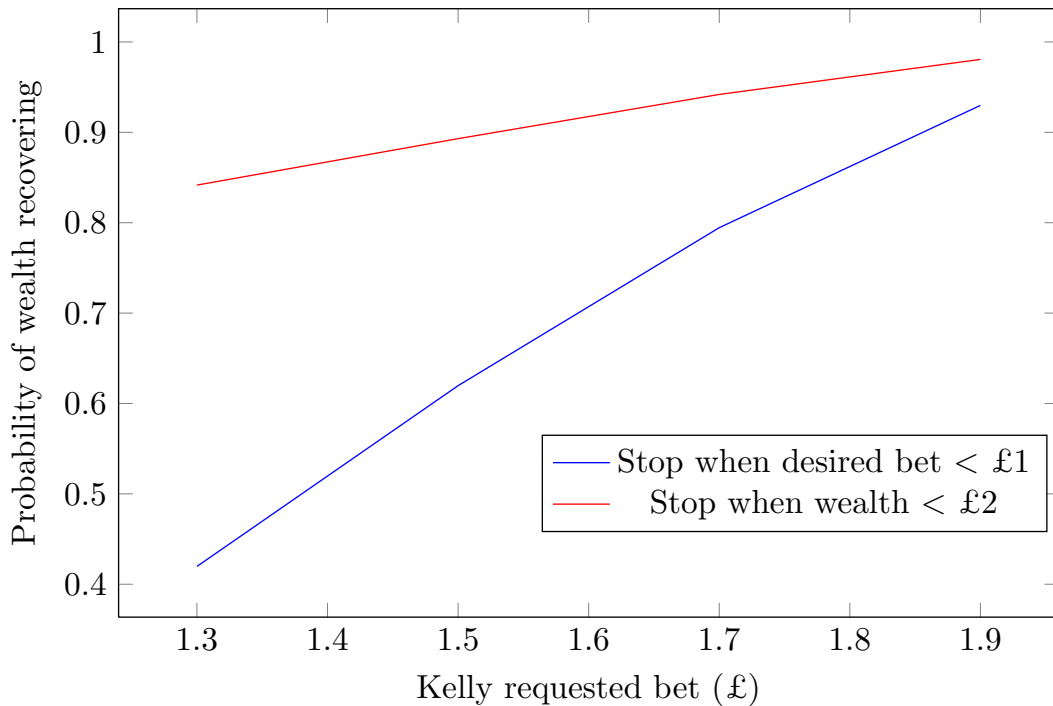


Figure 5.10.: A graph of the probability that wealth will increase sufficiently so that the Kelly bet is no longer below the minimum bet before the strategy goes bust. Starting wealth is £100, and $p = \frac{1}{2}$. There are two possible definitions of bust here. One is that it is no longer possible to place any bet, i.e. wealth has fallen below the minimum bet, and the other is that the required bet is less than half the minimum bet.

bet should be £1.01? Breiman [10] tells us that we should never bet more than the Kelly fraction, but if we bet twice the Kelly fraction we would expect, eventually, to go bust. In our case we are not doing this until our wealth has increased enough so that the minimum bet size no longer has an impact. Obviously our starting point is important. Figure 5.10 shows probabilities generated by a simulation in which success is the growth of wealth so that the desired fraction is the minimum bet. The simulation was performed twice, on the first occasion failure was wealth falling so far that the desired bet was less than half the minimum bet, and the second time failure was wealth dropping below the minimum bet.

The success probability is not very sensitive to any of the other parameters. Figure 5.11 shows this for a range of bet probabilities and values of starting wealth. We have fixed the size of the Kelly bet; if we have a large wealth then the expected gain on each bet must be small. If wealth is greater than £100 then the probability of a round up strategy being successful depends only open how much of a round up is needed. The gambler will need to decide if he is happy to do so.

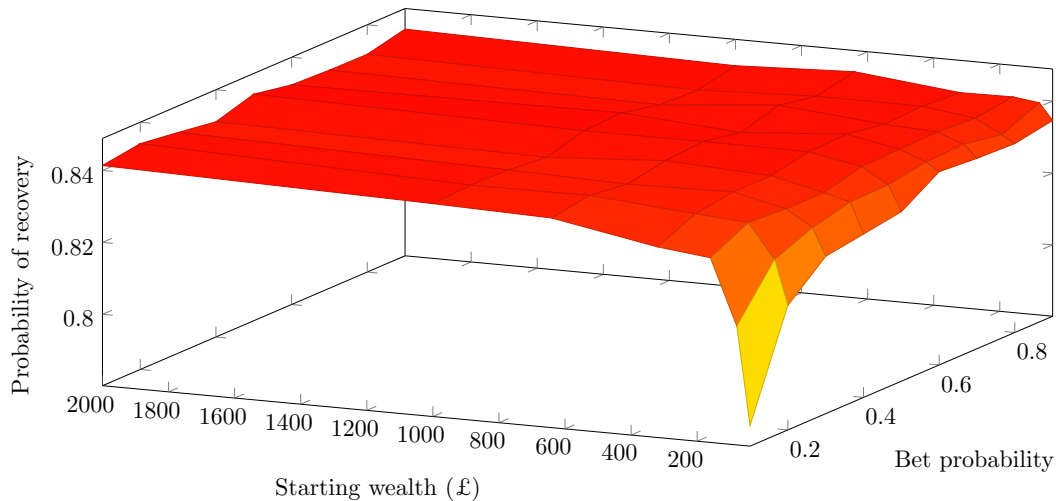


Figure 5.11.: A graph of the probability that wealth will increase sufficiently so that the Kelly bet is no longer below the minimum bet before the strategy goes bust. When starting wealth is larger than £200, and for most values of p this value is reasonably stable.

5.3.2. Minimum Bet Size: Secondary

When testing the code used to produce the examples in Chapter 4 we detected an error. The performance was more volatile than predicted. The prediction used the bets as fitted by our algorithm, but the simulation, correctly, used bets that were rounded to £2 or greater. This was somewhat unexpected as we had chosen a wealth great enough that the bets would be greater £2. What had happened was that, although some of the bets were greater than £2, many smaller bets were not. Increasing the wealth further increased these above £2, but also increased some other bets.

Here we consider a simpler situation. Having fitted our model we have a bet which is at least £2; we refer to this as the primary bet. The model also has another bet of size £1 which we refer to as the secondary bet. In this section we compare the impact on growth and volatility caused by not being able to make the secondary bet.

For example, consider an investor with wealth £10 000, examining a market which has one outcome with a probability of 0.1 winning and odds available of 10.09, and another outcome with probability of 0.01 and odds of 100.9. None of the other outcomes has useful odds. Following Kelly the investor would like to place a £10 bet on the first outcome and a £1 bet on the second outcome. This would give a growth rate of 1.20×10^{-3} . Unfortunately in this example, as is the case with Betfair, the minimum bet is £2. The investor cannot place the second bet, having to settle for a lower rate of growth, of 1.10×10^{-3} . Not only is the growth rate lower, but the risk is higher. The standard deviation of the return in the first case is £2.46 and higher (more risky) in the second at £2.58.

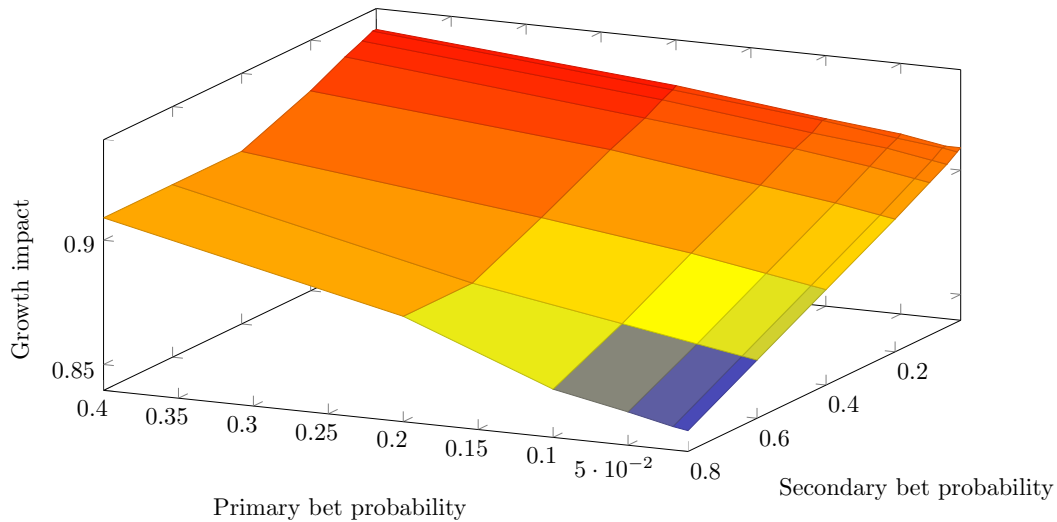


Figure 5.12.: A graph of the ratio of two growth rates in one case the smaller bet is ignored, in the other it is included. The ratio is calculated for different bets of differing probabilities, in each case the wealth is £10 000, the primary bet size is £10, and the smaller bet size is £1. A few of the points on the graph are not possible (when the sum of the probabilities is greater than one and have been filled using the nearest value).

Figure 5.12 extends the above example. In all cases the primary bet is still £10, but different probabilities are considered. It can be seen that as the primary bet becomes less likely, and as the secondary bet increases in probability then the ratio of growth rate decreases further. A lower number on this graph shows the lower growth rate given when smaller bets are excluded. In all cases the ratio of risk is independent of the probabilities.

The much larger variation is seen when we consider different primary bet sizes. Figure 5.13 shows the change in growth and risk for different values of primary bet size. Once the primary bet size reaches £200 the growth is reduced by less than 1% for all probabilities.

5.3.3. Minimum Bet Size: Large Loss

In Section 5.2 we have examined the probability of a large loss. We can use the results of that section to calculate the chance that a minimum bet will impact our strategy. For example, if we set $p = \frac{1}{3}$ and $o = 3.5$, then the optimal fraction of wealth to wager, f is $\frac{1}{15}$. If we continue using a minimum bet of £2 then as long as our initial wealth is at least £30 we can place the first bet. If it is exactly £30 and we lose that first bet then we will not be able to continue. In this case that will happen with probability $\frac{2}{3}$, and even if we win the first bet there is still a high chance that we lose later on. In fact the probability that we will have to stop because the minimum bet is too large is just over 94%. If we start with an initial wealth (greater than £30) then the probability that this




Figure 5.13.: A graph of the ratio of two growth rates and standard deviations for different sizes of primary bet. The growth ratio is the factor by which growth is lower when the secondary (£1) bet is excluded. The risk impact is the factor by which the standard deviation of return increases. In each case the wealth is £10 000, the primary probability is 0.1, the secondary bet probability is 0.01 and the smaller bet size is £1.

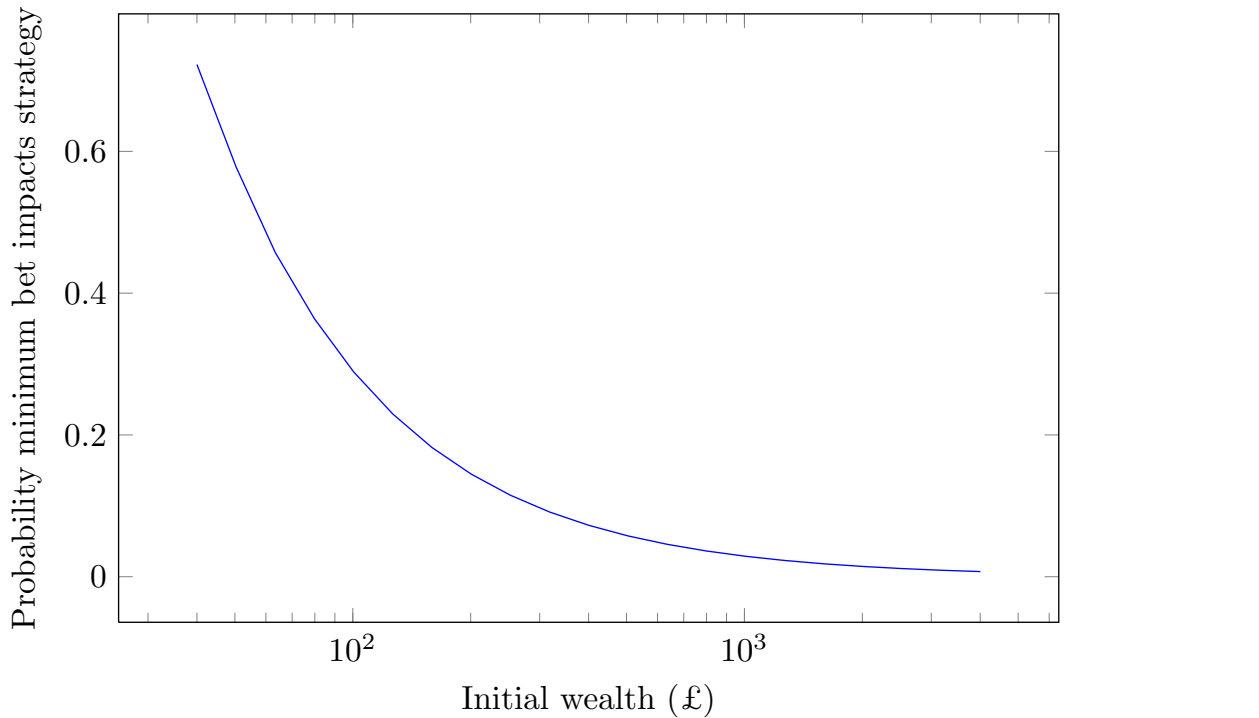


Figure 5.14.: A graph of probability that wealth falls so much that the required bet is below the minimum bet size by starting wealth. The minimum bet is £2, $p = \frac{1}{3}$ and odds $o = 3.5$.

strategy is impacted by the minimum bet size is shown in Figure 5.14. Selected values are also shown in Table 5.4.

There is no initial wealth large enough to ensure that a gambler who is following the Kelly Criterion may be certain that his strategy will not be impacted by a minimum stake (betting exchange or not). Here we consider the size of initial wealth so that there is only a small chance of a problem arising. In this section we fix the small chance at 1%. If the Blackjack table described in Section 5.2.2 had a minimum bet of £2 and if the chance of winning each hand is 50% then a gambler would need just over £200 000, to have 99% of not going bust. Meanwhile, a time traveller going back in time to just before Dixon and Coles published their paper on football modelling would need about £10 000, (assuming an average probability of 30%, and an expected value of 1.05). A gambler finding a favourable bet with a low probability of winning, and an expected value only slightly above 1 might do well to keep on searching for another opportunity.

5.4. Fractional Kelly

As we have shown in Section 5.2, following the Kelly Criterion can be an uncomfortable journey. For some the volatility of the returns are too great and therefore they adopt what is known as a fractional Kelly strategy, betting a fixed fraction of the Kelly bet.

Starting wealth (£)	Probability strategy impacted	
	$o = 3.5$	$o = 3.15$
32	92.4 %	
35	82.5 %	
90	32.3 %	94.7 %
100	29.0 %	85.1 %
500	5.81%	17.0 %
5000	0.581%	1.70%

Table 5.4.: The probability that a Kelly strategy is impacted by the minimum bet of £2, where $p = \frac{1}{3}$, and in the first case $o = 3.5$ and in the second $o = 3.15$. The second strategy cannot place a bet below £86.

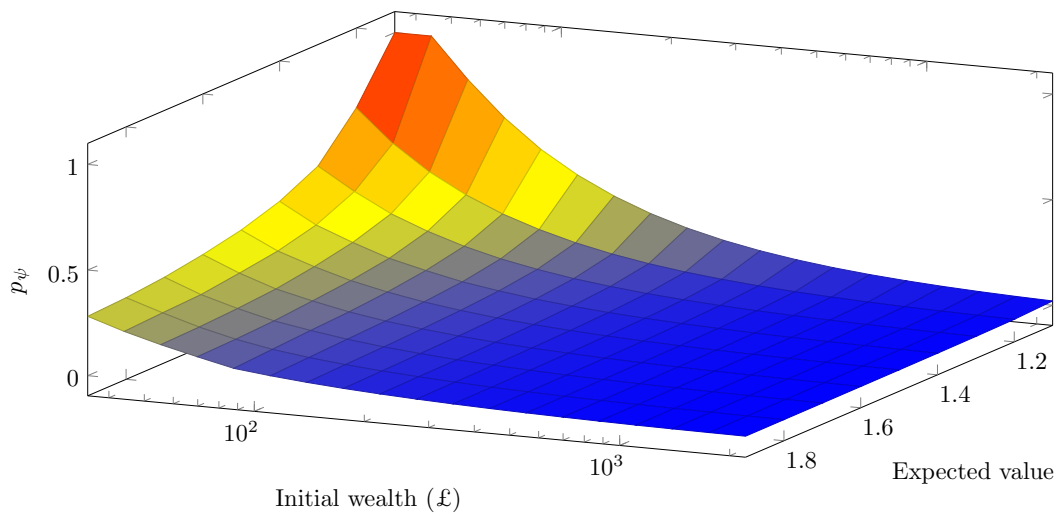


Figure 5.15.: A graph of probability that wealth falls so much that the required bet is below the minimum bet size by starting wealth. The minimum bet is £2, $p = \frac{1}{3}$. Note that the EV scale is reversed. Two bets, with expected value 1.1, have too little starting wealth and have been plotted with probability 1.

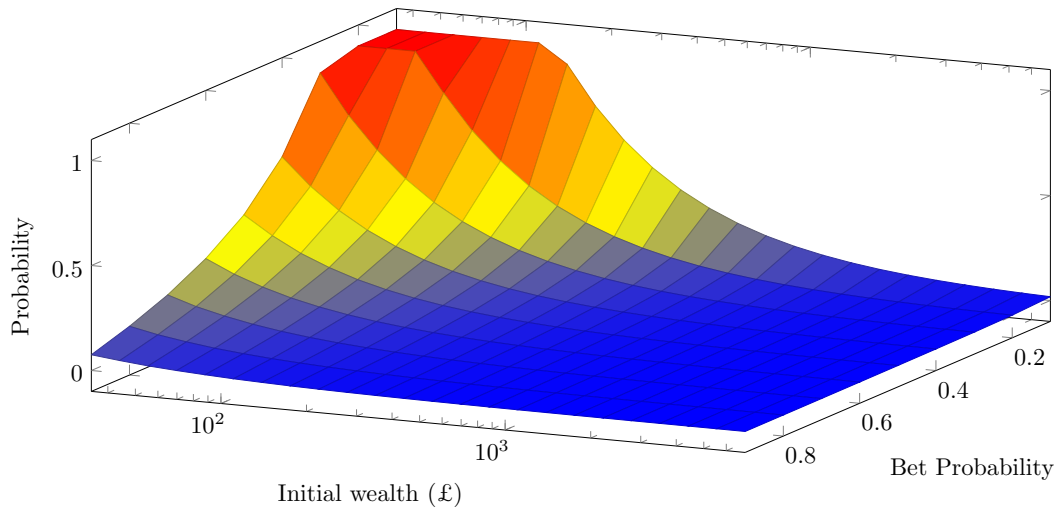


Figure 5.16.: A graph of probability that wealth falls so much that the required bet is below the minimum bet size by starting wealth. The minimum bet is £2, the expected value is $\frac{3.5}{3}$, which is the same as used in Figure 5.14. Note that the p scale is reversed. If the initial wealth is too low to place any bets, given the minimum bet size, then a value of 1 is used for the probability.

This can significantly reduce losses encountered on the path, whilst reducing growth, but by an acceptable proportion. Table 5.5 gives an example for fractions of $\frac{1}{2}$, and $\frac{1}{4}$, of the bet we have been using throughout this chapter, which has odds o of 3.5 and probability of winning of $\frac{1}{3}$.

We can see that using a fractional Kelly strategy will reduce volatility. This chapter, however, concerns the appropriateness of the assumptions made by Kelly in general, and in particular with regard to the minimum bet size. To be considered bust for this bet on the full Kelly strategy wealth needs to fall below £30, whereas the wealth needs to fall below £120 on the quarter Kelly strategy. Indeed, an investor starting with £150 would be more likely to be considered bust if he followed the quarter Kelly than the whole – although being considered bust when he still has £120 might be less painful to him. Increasing his starting wealth to £170 changes this, and with a starting wealth of £500, the chance of being unable to place the next bet falls below one in ten thousand. Table 5.6 has details for a range of initial wealth. When considering Figure 5.7, but with data calculated using a $\frac{1}{4}$ fractional strategy, all of the probabilities were below 10^{-6} .

As discussed above some gamblers and investors explicitly follow a fractional Kelly strategy. In addition to reducing volatility this strategy also provides some limited protection if probabilities or returns are incorrectly estimated. Without providing any evidence I claim that most people do so without thinking. Kelly asks us to bet a fraction of our entire wealth (if we follow Bernoulli that would include the net present value of future excess earnings). Most people allocate a fraction of their wealth to a particular

Fraction	Growth ($\times 10^{-3}$)
1.0	5.39
0.5	4.08
0.25	2.40

Table 5.5.: The logarithmic growth rates using fractional Kelly staking, where $p = \frac{1}{3}$ and $o = 3.5$.

Fraction	Growth ($\times 10^{-3}$)	Initial wealth			
		£150	£170	£500	£1000
1.0	5.39	2.08×10^{-1}	1.95×10^{-1}	6.27×10^{-2}	3.13×10^{-2}
0.5	4.08	7.95×10^{-2}	6.58×10^{-2}	2.28×10^{-3}	2.96×10^{-4}
0.25	2.40	3.40×10^{-1}	2.18×10^{-1}	8.95×10^{-5}	7.78×10^{-7}
0.1	1.05			2.82×10^{-3}	7.44×10^{-10}

Table 5.6.: The probability of being unable to continue with a fractional Kelly strategy for various fractions, where $p = \frac{1}{3}$ and $o = 3.5$.

betting or investing pool, and do not include the value of their house, pension and other savings, so by implication they are following a fractional Kelly strategy even if they claim that fraction is one. Of course there are some people, private equity and hedge fund managers such as Warren Buffet, who have such a large proportion of their wealth in one place and with which they take large risks, that their other assets are insignificant. I suspect that they are the minority.

6. Dynamically Hedging a Portfolio of Bets

Kelly considered betting on horse racing, where a bet was placed, the horses ran and winnings collected if successful. This would be repeated. For most races, much of the betting turnover occurred after the previous race had finished. Much gambling is still done in this way in much of the world. In some places (such as the UK) the rise of betting exchanges has led to increased betting several hours before a race. Imagine now having placed a back bet and then the odds change; what should the better do? The odds may have improved (increased). If originally faced with such odds Kelly would have told us to place a larger bet than currently placed. By how much should the bet be increased? If the odds were to worsen, and the corresponding lay bet odds reduce, should the gambler take a guaranteed profit now, or risk it in the race?

6.1. Increasing or Unwinding Previously Placed Bets

It is surprising that Thorp did not consider such activity until 2008 [40] when emailed by a reader of his regular column. It is especially surprising because Thorp wrote extensively about the application of Kelly Criterion to financial markets. The author of the email was Proebsting.

6.1.1. Proebsting's Paradox

Imagine being offered odds of 3 on heads appearing on the toss of a fair coin. Kelly would tell us to place 0.25 of wealth on such a favourable bet. Have placed 0.25 of our wealth on heads we are offered odds of 6 on heads. Assuming we are to place an additional bet of f^* at odds of 6 the two outcomes are:

$$\text{Heads } 1 + 2 \times 0.25 + 5 \times f^* \tag{6.1}$$

$$\text{Tails } 1 - 0.25 - f^* \tag{6.2}$$

This gives a growth rate G

$$G = 0.5 \log(1.5 + 5f^*) + 0.5 \log(0.75 - f^*)$$

Maximising G leads both Thorp and Proebsting to $f^* = 0.225$. They both remark in their email exchange that they find this surprising, at least at first. If we had been offered odds of 6 initially Kelly would have told us to bet 0.4 of our wealth, and yet we have just bet 0.475 on worse odds. This leads Proebsting to point out that when betting on the fair coin toss Kelly would never allow a bet of 0.5 or more of our capital, no matter how high the odds, but here it is easy to create a series of bets which leads to all wealth being bet on a coin toss.

This is a less significant issue for a betting exchange than it is in theory. Most exchanges have a finite number of available bets and maximum odds. Betfair, for example, has maximum odds of 1000. Aaron Brown, nonetheless offers us a satisfying solution to this paradox. If we consider odds that move, then before we place our second bet we should consider our wealth. The wealth is not unchanged, but has reduced because the value of our bet is reduced (because odds are now 6, odds of 3 are worth less). In this case the bet of 0.25 of initial wealth at odds of 3 revalued to odds of 6 leads to a loss of 0.125. The bet needing to be placed is also 0.125 of the original capital. The new wealth is 0.875 of the original wealth. Kelly tells us to bet 0.4 of this which is the same as a bet of 0.35 of the original capital. We do not actually unwind and then replace, we just bet a net 0.225 of the original capital ($0.225 + 0.125 = 0.35$) and paradox is resolved.

6.1.2. Previous Bet: Increasing Risk

More generally, consider that we have placed a back bet of fraction f , of wealth W , at odds o and probability p then as usual

$$f = p - \frac{1-p}{o-1}.$$

Assume now that the odds improve to o^* , where $o^* > o > \frac{1}{p}$. We need to find the fraction of W , f^* which maximises the growth of our wealth, G given by

$$G = p \log(1 + (o-1)f + (o^*-1)f^*) + (1-p) \log(1 - f - f^*)$$

$$G' = 0 \Rightarrow f^*(o^*-1) = p(o^*-1) - pf(o^*-1) - (1-p) - (1-p)(o-1)f$$

but we know f so

$$f^*(o^*-1) = po^* - 1 - [p(o^*-1) + (1-p)(o-1)] \frac{po-1}{o-1}.$$

Re-arranging gives

$$f^* = \frac{po(o^*-o)(1-p)}{(o-1)(o^*-1)}.$$

This is the additional fraction of the original wealth which needs to be staked at the

new odds. The total wealth staked is $f + f^*$. Assume that instead of increasing our bet when offered the increased odds we decided to remove it instead. Assume that we could place a lay bet at the same odds o^* of size f^L , such that

$$1 + (o - 1)f - (o^* - 1)f^L = 1 - f + f^L,$$

re-arranging to give

$$f^L = \frac{fo}{o^*}.$$

If we revalue our existing bet at the current market price we see that our wealth has fallen by $f - f^L$, giving the new wealth relative of $1 - f + f^L$. Now offered odds to back of o^* , as before we calculate the fraction of this new wealth f^N which we would like to stake, which is the simple Kelly fraction, therefore,

$$f^N = p - \frac{1 - p}{o^* - 1}.$$

As a fraction of our original wealth f^{OW} this would be $f^N(1 - f + f^L)$, substituting for f^N, f , and f^L , gives

$$f^{OW} = \left(p - \frac{1 - p}{o^* - 1}\right) \left(1 + \left(p - \frac{1 - p}{o - 1}\right) \left(\frac{o}{o^*} - 1\right)\right)$$

The change in wealth (Δ_W) due to the new bets is $f - f^L + f^{OW}$. Re-arranging gives

$$\begin{aligned} \Delta_W &= \frac{po - 1}{o - 1} \cdot \frac{o^* - o}{o^*} + \left(p - \frac{1 - p}{o^* - 1}\right) \left(1 + \left(p - \frac{1 - p}{o - 1}\right) \left(\frac{o}{o^*} - 1\right)\right) \\ \Delta_W &= \frac{1 - o^* - op - oo^*p + o^*po - oop + oopp - po^*po}{(o - 1)(o^* - 1)} \\ &= \frac{(po - 1)(o^* - 1) + op(o^* - o)(1 - p)}{(o - 1)(o^* - 1)} \\ &= f + f^* \end{aligned}$$

So the paradox is resolved. If we place a bet and reconsider it at a later stage we need to revalue our capital. If the odds on a back bet have risen then we need to accept that we have lost money. One of the criticisms of expected log investing is that an investor does not know what is coming next and a better opportunity may present itself later. We discuss this further in Section 8.2.

6.1.3. Previous Bet: Reducing Risk

Whilst the problem of what to do when odds improve has been discussed, the similar problem of what to do when odds worsen does not appear to have been considered. Perhaps this is because much of this work has been done by Americans and there are no

betting exchanges legally available in the United States. Therefore they have not had to consider lay bets. As in the previous section we will assume that we have already placed a bet of size f at odds of o and probability p with f given by the Kelly fraction

$$f = p - \frac{1-p}{o-1}.$$

This time odds fall so far that it is possible to place a lay bet with odds o^* where $o^* < o$. We will place an additional bet of size f^* , chosen to maximise the growth rate G , where

$$G = p \log(1 + (o-1)f - (o^*-1)f^*) + (1-p) \log(1 - f + f^*)$$

$$G' = 0 \Rightarrow f^*(o^*-1) = (1-p)(1 + (o-1)f) - p(o^*-1)(1-f)$$

$$f^*(o^*-1) = 1 - po^* + f(o-1 - po + po^*)$$

but we know f so

$$f^*(o^*-1) = 1 - po^* + \frac{po-1}{o-1}(o-1 - po + po^*).$$

Re-arranging gives

$$f^* = \frac{po(o-o^*)(1-p)}{(o-1)(o^*-1)}.$$

If $o^* = 1/p$, then $f^* = of/o^*$. That is a complete unwind, or in exchange talk *greening off*, on Betfair outcomes with profits are shown in green and losses are shown in red. If all outcomes are green it is said to have been greened off.

6.2. Including Previous Bets

Let us consider a simple coin tossing example. The coin is fair ($p = \frac{1}{2}$) and before each toss odds, o of 2.05 are available. If we follow the standard Kelly method we would stake a fraction, f , of just under 2.4% of our wealth, and if repeated this would eventually tend towards a growth rate, G_1 , of 2.98×10^{-4} . Now imagine that having placed this bet the odds, o^* improve to 2.10. Using

$$f^* = \frac{po(o^*-o)(1-p)}{(o-1)(o^*-1)}$$

we would stake an additional 2.2% of our original wealth giving a new growth rate, G_2 of 5.68×10^{-4} . Of course, if we know that the odds will improve every time, we should avoid placing the first bet altogether, and place the bet when the odds are 2.10. Then growth rate, G_3 would be much better at 11.4×10^{-4} , even if the odds did not increase every time, merely most times, we should still wait, rather than place the first bet. In general if probability that the odds will increase is π_{inc} the growth rate of placing the

first bet and placing the second where possible will be

$$\pi_{inc}G_2 + (1 - \pi_{inc})G_1$$

and the growth rate of placing no initial bets and placing only the higher bet if possible is

$$\pi_{inc}G_3.$$

From these it is easy to calculate whether or not we should place the first bet. In this case we should not place the bet 2.05 if we think that the probability that odds will improve to 2.1 is greater than 34.4%.

If the probability does not increase then perhaps it will decrease. As above the first bet is placed at odds of 2.05, but this time the odds available fall so far that it becomes possible to lay this bet at 2.00. In this case the fraction of the original wealth will be 2.44%. In this case this is the same as $\frac{of}{o^*}$, which is the equation for the size of a bet required to unwind a previous bet. In this case the fraction by which wealth increases on a win is the same as the fraction for a loss (1.00059) giving growth of 5.95×10^{-4} . This is twice the growth of leaving bet to run, not only that but also this bet now has no risk. If we were sure that the odds would fall we might have placed more wealth on the first bet. We will refer to such activity as trading which we define as backing something in the expectation that the odds will fall in the future, or laying something in the expectation that the odds will rise. The activity we have been following so far is more akin to the investor who buys because the fundamental value is good. For example, when betting during a race, perhaps the odds on a horse fall whenever the commentator mentions the horse favourably. If a trader thought he could predict which horses might be mentioned, or could trade much faster than everyone else he might back horses which have odds he thought would fall, irrespective of which horses he thought would win. Whilst such activity has much promise it is beyond the thrust of this thesis.

6.3. Moving Probabilities

This chapter has, so far, concentrated on moving odds. We shall now consider a change in probabilities. Imagine we have developed a football model which predicts a team's performance based upon the starting players. If a player is suspended or injured the team's probability of winning the next match will be reduced, and similarly if he recovers from injury more quickly than previously expected the probability will increase. Of course, the odds in this case will also move as shown in Figure 6.1. This shows the match odds of an English Premiership match between Manchester United and Chelsea for the week before the start of the match on Saturday 3rd April 2010. A week before the match Manchester United, playing at home, are the clear favourite. However, on




Figure 6.1.: A graph showing the changing match odds for the Premiership match between Manchester United and Chelsea on 3rd April 2010. Manchester United, the home team, is the clear favourite, with odds shown in red and blue, but then United’s odds increase suddenly on Tuesday evening. At this moment Wayne Rooney, one of Manchester United’s star players, limps off the pitch during a match against Bayern Munich and there is concern that he might not be able to play against Chelsea. Later in the week it is announced that his injury is serious enough to cause him to miss this match. Data taken from Betfair API.

Tuesday evening Manchester are playing Bayern Munich and Wayne Rooney is injured and limps off the pitch. There is concern that he will not be able to play on Saturday. Sometime on Friday it is announced that, whilst Rooney’s injury is not as bad as some had feared, he will miss the match against Chelsea. Strangely the odds of Manchester United winning improve. This fits nicely with the old City adage “Buy on rumour, sell on fact”. Perhaps the market participants are convinced by Sir Alex Ferguson, the Manchester United manager, as he discusses his plans for the team and Rooney’s replacement. United lost.

6.4. The Uncertain Biased Coin

In an imaginary federal country, made up of two states each with its own president, the national game is football. Many things, including the minting of coins, are devolved to the states in this country. However, the football league is not, and teams from both states take part. In order to satisfy national pride the referee has two coins in his pocket, one from each state. At the start of the match the away team is invited to call heads or tails and the referee randomly takes one of the two coins out of his pocket, tosses it and shows the result to both captains. Life is somewhat dull in this country and

consequently the people have taken to gambling on everything, including the toss of the coin. There is an active market in bets on this coin toss, with odds for both teams being close to 1.98 to back and 2.02 to lay. Unfortunately, one of the presidents is a paranoid narcissist, (perhaps the surprising fact is that one of the presidents is not). He was convinced that the image of the other president on his state's coin was larger than his own. Last year he ordered a new coin with a much larger and thicker head on the head side of the coin. The tails side remains unchanged.

Mr R. B. Trage did some tests on this new coin. He has flipped it many times and was convinced that, because of the weight of the head, it will land tails up 52% of the time. Last year he followed a strategy of backing tails every match. As a Kelly investor he placed 1% of his wealth on tails, for a logarithmic growth of 4.90×10^{-5} . This is small, but there are lots of football matches each season and he is happy with that. This year things have become more complicated.

The local television station has introduced high definition coverage of football matches. Having backed tails Mr. Trage can see on the television coverage that the coin chosen by the referee is the fair coin. What should he do? If the odds are unchanged he should do nothing; at 2.02 it is not worth reducing the bet. But if the odds have moved lower, perhaps because of Mr. Trage's previous bet, then he should lay tails. At 2.01 he would reduce his risk by about half, and at 2.0 he would remove all risk, but guarantee himself a small loss. (Remember, he bet tails in the hope that biased would be chosen, but it was not.) If, on the other hand, Mr. Trage sees the biased coin he should increase his bet to 0.0302. If he was to repeat this his asymptotic growth would be 14.90×10^{-5} .

Of course, he can do even better. If he is sure to observe the coin in the referee's hand, he should not place his bet until after he has seen the biased coin, and do nothing if he sees a fair coin. Doing so would offer growth of 22.36×10^{-5} .

In a more realistic scenario it would not always be obvious which coin is chosen: perhaps only one side of the coin is visible and the tails sides of both coins are identical. The best strategy to follow in this case would be to bet slightly less than half the ideal fraction, and increase it if shown the biased coin, giving growth of 11.55×10^{-5} . If the situation is even worse and the television coverage only shows the coin half of the time, and as before it is only possible to identify which coin has been chosen half of those times then three quarters of the time we are uncertain which of the two coins has been chosen. In this case the best strategy is to bet about 0.007, initially and increase to 0.0302 when certain the biased coin has been chosen, giving growth of 7.76×10^{-5} . Following the greedy-maximise-at-each-stage method performs nearly as well, in this case, with growth of 7.40×10^{-5} , outperforming the wait-and-see strategy, which has growth of 5.59×10^{-5} .

It is clear from this example that knowledge of how the model might change in the future becomes very important. The above example is slightly skewed: it either became a bet

to be increased, or an unfair bet, but too close to fair to unwind. In many examples it would be possible to unwind and go further.

7. Market Making with Inverse Kelly

Having extended Kelly to modern betting exchanges it appears natural to want to test these extensions. To do so requires a model with greater predictive accuracy than the betting market. The betting markets on leading sports appear to be very efficient. It would take an exceptional model to consistently beat the market. The markets of both lesser sports and the lower divisions of mainstream sports would appear to have better prospects for testing the new developments. Unfortunately, these markets tend to have substantially wider prices; indeed, sometimes no prices until a few minutes before the match or race starts. It is possible that the prices available to match with, in the market, will be so poor (too low to back and too high to lay) that it may be necessary to place one or more bets in the market in the hope that they will match later.

7.1. Market Making

Market making is a financial services activity where the market maker (often an investment bank) quotes prices (usually buy and sell) for a particular financial product on a given exchange. Here we use the term in a similar manner, but for betting exchanges. The market maker hopes that without taking a large amount of risk he can profit by selling products at a higher price than he buys them. And this will be true for the betting exchange market maker. He will need to back the runners at higher odds than he lays them.

Kelly, and every other gambler until about ten years ago faced a simple choice when placing a bet; if he would like to back a particular outcome, but the odds were too low he could back them anyway or do nothing. With betting exchanges we now have an additional option; we can offer a bet at our choice of odds. In time that bet might be matched (traded) or it might expire without trading.

In Table 3.5 we saw that the average over-round for correct score bets was about 2% for English Premiership football in the 2009/10 season. On a recent Wednesday in February 2013 the over-round for betting on the correct score for the following Saturday's League One matches was about 12%, and there was no realistic price in any market for matches on the following Tuesday. (English League One is the third level, after the Premiership and the Championship.) Anyone planning to place a bet would need to choose a price, rather than merely deciding to trade, or not to trade with an existing price. We suggest

choosing a notional sum of money, C , and then selecting the decimal odds o such that Kelly's fractional stake f gives the desired bet size b . Then if p is the probability of the backed event from Kelly we have:

$$f = p - \frac{1-p}{o-1}.$$

Rearranging this gives:

$$o = 1 + \frac{1-p}{p-f}. \quad (7.1)$$

From this we see that we need $p > f$. If offering lay bets as well as back bets there is another important consideration. Each of the entries is considered on its own, but it is possible that more than one will trade at once. For example, if offering back and lay prices on a tennis match it is possible that someone may match the back on one player and the lay bet on the other player. With some exchanges this is in fact likely; Betfair's web client automatically does this calculation and a client will see the size which is generated by combining a direct transaction with appropriate combinations of other bets. So b should be chosen to allow for this.

To be practical this method combines existing bets into the calculation and at the same time we consider multiple markets. So, if an event has n possible outcomes and the probability of outcome i is p_i , and if there are m possible bets which depend solely upon those outcomes, then bet j has odds o_j and traditionally we are choosing the fraction of our wealth x_j to wager on bet j . For each of the j we can take into account previously made bets with size s_j and odds e_j . \mathbf{M} is a matrix with $m_{ij} = 1$ if bet j is considered winning if outcome i occurs, and $m_{ij} = 0$ otherwise. Previously, we would have wanted to maximise G where:

$$G = \sum_i p_i \log(1 + \sum_j m_{ij}(o_j x_j + e_j s_j - x_j - s_j)).$$

For market making we find for each j in turn the value of o_j such that maximising G gives $x_j = f$. Again we must ensure that $f < p_i$ for all i . Other than this constraint we are free to choose f as necessary to control the desired over-round. From Equation 7.1 we can see that

$$\frac{1}{o} = \frac{p-f}{1-f},$$

so the lay under-round if we place m back bets is

$$\frac{1-mf}{1-f}.$$

The choice of C becomes a decision about the confidence of the model used. If making markets on one side only (all backs, or all lays) we might choose it such that $f = \frac{b}{C}$;

Result	Winnings	Probability
Home Win	£- 3.00	0.50
Away Win	£0.00	0.25
Draw	£10.00	0.25
Expected value	£1.00	

Table 7.1.: Payout value for each of the possible match results for the four bets in the example. We have £100 lay bets of 3.9 for the draw, 4.0 for the away win and two for the home win of 2.04 and 1.99.

or conversely, when offering bets on both sides $f = \frac{2b}{C}$, where b is the size of the bet placed in the market in both cases. If a bet is matched at size f the next market prices generated will be such as to attempt to unwind this bet at fair value, or the middle of our previous market. With greater confidence in a model a higher value of C would appear suitable as prices would change more slowly with trading activity.

7.2. Example

One of the attractions of this method of market making is that it deals well with complicated sets of dependent markets, such as those discussed in Chapter 4. In this section we discuss a simplified example.

Consider a football match with $\mathbf{p} = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})^T$. This is close to the average of a well matched league, such as the English Premiership. There is a 50% chance of a home win, and a 25% chance of an away win or draw. Suppose that we set wealth to be £5 000, and $f = 0.02$. This would create bets of £100. Figure 7.1 shows the initial bets which would be created by this example. There are three back bets below fair value and three lay bets above fair value. We now assume that another member of the exchange places a trade, in this case backing the draw, at 3.9; so we have laid the draw. Figure 7.2 shows the situation having recalculated our prices including the new bet. The odds on the draw have fallen, and the odds on the other two outcomes have increased. Had we been making markets in the correct score, odds on 0-0, 1-1 and so on would also have fallen, with odds on other markets rising. The red square in Figure 7.2 indicates a matched bet. We now assume that a customer backs the away win (at 4.0), giving the results shown in Figure 7.3. This is followed by two more bets, of a home win (at 2.04), shown in Figure 7.4, and after the odds have adjusted another home win bet (at 1.99) giving the final situation shown in Figure 7.3. The value of this position is £1.

It is easy to see the biggest problem with market making. Everything went well in our example. We matched four £100 bets and yet the profit is only £1. This is true with most forms of market making.

HomeWin	1.95	2.06
AwayWin	3.90	4.30
Draw	3.90	4.30

Figure 7.1.: Initial prices in football market making example. Prices on the left (in blue) are those that the market can back (i.e. we would lay) and those on the right (in pink) are those that the market can lay. The yellow dotted line indicates fair value given by $p = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$, and $f = 0.02$. All prices are rounded to Betfair valid prices.

HomeWin	2.00	2.10
AwayWin	4.00	4.40
Draw	3.65	4.10

Figure 7.2.: Prices in football market making example after first bet is matched. The red square indicates that the customer has backed the draw (so we have laid it). Prices have now adjusted. The yellow dotted line indicates fair value given by $p = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$, and $f = 0.02$. All prices are rounded to Betfair valid prices.

HomeWin	2.04	2.14
AwayWin	3.75	4.10
Draw	3.75	4.10

Figure 7.3.: Prices in football market making example after second bet is matched. The second red square indicates that the customer has backed the away win (so we have laid it). Prices have now adjusted. The yellow dotted line indicates fair value given by $p = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$, and $f = 0.02$. All prices are rounded to Betfair valid prices.




Figure 7.4.: Prices in football market making example after third bet is matched. The third red square indicates that the customer has backed the home win (so we have laid it). Prices have now adjusted. The yellow dotted line indicates fair value given by $p = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$, and $f = 0.02$. All prices are rounded to Betfair valid prices.




Figure 7.5.: Prices in football market making example after fourth bet is matched. The fourth red square indicates that the customer has backed the home win for a second time (so we have again laid it). Prices have now adjusted. The yellow dotted line indicates fair value given by $p = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$, and $f = 0.02$. All prices are rounded to Betfair valid prices.

7.3. Small Scale Live Test

In Chapters 3 and 4 we were able to create a test using historical data. Suitable assumptions, such as a test size much smaller than the market so as not to impact it needed to be made. In this case, the whole point is that we are having an impact on the market. Prices are not available at the time the bets are placed into the exchange. Whether the bets will be matched immediately, improved upon, or ignored is not immediately obvious. We would like to create a market making simulation tool. We have not done so, but as a tentative preliminary examination we have created some data which will help with design and calibration. One of the advantages of studying betting markets rather than financial markets is that it is possible to make tests in real markets with a considerably smaller capital outlay; the minimum bet on Betfair is £2.

We used a model based upon Dixon and Coles [82], fitted to the previous matches of the current season and the two prior seasons, but with a value for tau of one. We estimated the probabilities of final scores for English Football League One and Two (the third and fourth level, respectively). On Betfair it is typically possible to bet on football matches a week or so in advance. When a match was added to Betfair the market making tool randomly assigned it a category. One category was to ignore the market, placing no bets, offering these markets as controls. Other categories had various choices for b , C and f . One category (8) assumed that confidence in our probabilities increased with the number of bets matched. As this confidence grew the prices quoted narrowed (f reduced) and the sizes quoted (b) increased.

We had a suitable model and data to quote prices in four markets for each match: Match Odds; Correct Score; Over/Under 2.5; and Over/Under 3.5. Unfortunately the cost of quoting in Correct Score was too high. The amount of capital needed to quote back and lay bets of size b in each of the m runners of a market is $b(m + l - 2)$, where l is the highest lay odds placed. The Correct Score market has 17 runners and it is quite common for the the highest lay odds to be over 50. As the smallest bet allowed is £2 this would require over £130 for each game compared to about £10 for each Match Odds market.

The market making tool ran, with some interruptions, from 11th February to 28th March 2013, a period of 215 football matches. Table 7.2 gives a breakdown of the matches over this period. Of the 215 matches, we placed no bets on 91 of them. This was to enable us to see if our activity made any other difference. We had speculated that quoting prices in a market might increase the total activity in that market (beyond our own activity). At first it appeared as though we had made an impact, at least with category 8 matches. The average volume was double that of the others. The strategy that category 8 deploys increases the size and reduces the over-round as confidence in the prices increases. It was not a surprise that this category had the most impact. Performing an F-test on the

	Fallow	Cat 1 - 7	Cat 8	Total
League One	45	49	13	107
League Two	46	47	15	108
Market Size	£39 682	£40 689	£84 669	

Table 7.2.: Details of football matches which occurred during the market making test. Fallow matches were left untouched by the market making software.

	Match Odds	Over or Under 2.5	Over or Under 3.5	Correct Score
Back	1065	577	216	0
Lay	2864	749	232	19

Table 7.3.: Details of the bets matched during the market making test. Only one match, only, had placed Correct Score bets.

data shows that the variances are significantly different with a high degree of confidence. However, it became clear that this was not due to our involvement. One of the category 8 matches (Tranmere vs Stevenage) happened to be on TV during a weekend when there was no Premiership or Championship football. This one match had over 10 times the normal turnover. This together with another unusual match account for all the increase.

It is worth pointing out here that the distribution of the fallow matches is not uniform across this period. In the first half of the time period there is a higher proportion of ignored matches than there is later in the period. As it became clear that the strategy was not bleeding money the ratio was changed. It is possible that this decision has introduced bias.

Over this time we placed 6 415 bets. In several cases the software was interrupted and in some others there were human interventions which introduced errors. Sometimes the thread detecting new markets in Betfair was too slow and by the time it placed prices in the market there were already prices there. Occasionally these matched immediately. One of the parameters that our future simulation will need to fit is the inter-arrival times for the bets. These almost instant matching bets skew the data significantly. Some of these we can easily detect and have excluded, reducing our total bet count to 5722 bets. Table 7.3 shows the breakdown of these bets. As previously stated, quoting in the Correct Score market was particularly expensive and was, therefore, assigned a low probability. This happened only once, matching 19 lay bets not including the winning score. Unfortunately we have been unable to draw any strong, significant conclusions from the data collected.

It is immediately obvious for markets with more than two outcomes that customers prefer to make back bets (so we make lay bets). Perhaps for binary markets customers still favour back bets, but the Betfair web client converts our back bets into lay bets on

the other outcome.

We had expected that bets which matched quickly might be more likely to be losing bets. To estimate this we calculated for each market a money weighted average odds, normalised so that the reciprocal gives a market implied probability. This is our fair market price. We use this to calculate an expected profit on each bet. We anticipated a positive correlation between this and the time between placing a bet and it being matched. Surprisingly, the calculated correlation is -0.07. See Figure 7.6. It is possible that when the data has been further cleaned this will change, but it is surprising. It is possible that we need to consider how far in advance of the match we are placing bets. If we place bets too far ahead it may be that no-one notices these bets for many hours, or days, and this will probably include the bet with the highest negative expected value.

7.4. Market Making: Final Comments

More work is needed on this method before any significant conclusions can be made. Anyone undertaking market making needs to realise that in general when things go well small profits will accrue, and yet when things go wrong the losses will be large. In the example above the system worked perfectly and made an expected profit of £1. Had only the first bet been placed (because the system crashed during calculation, or ran out of money, or a power cut, or lack of interest...) and the match been drawn the loss would have been £290. Clearly one needs to be sure that these events are indeed rare. Nassim Taleb argues [162] that often such events are more common than people realise. The strategy above that showed some promise increased in size over time. This helped to overcome some of the worry about model accuracy, but would not have helped with any of the the possible issues mentioned here.

Once convinced that market making is the correct strategy and whilst choosing a method for calculating prices there is an obvious further choice. Should the method use the model or should it be independent from it. This method is independent. If there is a systematic problem with the the model, there is a chance with this method that some danger may be avoided. As a consequence, no model knowledge is used. If the probability of the home team winning is set too high at the outset, bets laying this will be matched and eventually the prices will rise (implied probability fall). This method will, at the same time, adjust all other prices down. This might not be the best strategy. Examine again Figure 6.1, the graph of match odds changing over time for Manchester United verses Chelsea. When Manchester United's odds spiked up, Chelsea spiked down, but the draw didn't move much.

Similarly in other sports there may be correlations captured in the model, but not used by Kelly market making. For example, in Formula One motor-racing in the 2013




Figure 7.6.: Time taken before the matched bet was matched, compared to expected value of trade, measured using trade weighted average price.




Figure 7.7.: Length of time between matching and start of match, compared to expected value of trade measured using traded weighted average price.

season, tyre degradation is crucial. Imagine having placed some bets the day before the race the temperature is much hotter than expected. Perhaps the favourite team, say Red Bull, is impacted by this more than the second favourite team. The odds of Sebastian Vettel winning the race increase. This model would respond by reducing all other drivers including Mark Webber, Vettel's Red Bull team-mate. A method which used the underlying model would probably not make this mistake.

8. Conclusion

8.1. Achievements

This thesis has extended Kelly's original method for choosing a stake size, bringing it up to date with betting available in the UK (and many other countries). Whilst there has been much research on Kelly's work, and on Maximising the Expected Log (MEL) more generally (see Section 2.1), most of it has concentrated on financial markets rather than betting. The research that has focused on betting has predominantly been authored by Americans (such as Ed Thorp) and as betting exchanges have largely avoided America their work has been limited to what we (in the UK) would call traditional bookmakers. There has been no research on laying bets or on unwinding bets, and very little on increase bets, trading bets and market making. This thesis fills in that gap. The principal achievements of this thesis are:-

Extended Kelly to lay bets In those countries which allow betting exchanges the members are quite used to both backing horses and laying them. If the odds on one horse are too low, rather than backing many of the other horses it is more convenient to lay just that one. This thesis allows such a user to calculate the appropriate size of the bet. This is calculated explicitly.

Extended Kelly to sets of bets (numerically) Breiman posed a more general form of Kelly's original problem (see Section 2.1.2) and stated that it was difficult to solve generally. This thesis solves such a generalisation numerically. In practical cases the solution is found quickly with only basic PC hardware. We show an example using historical football data and the return is higher than betting on each market individually.

Incorporated commission Betting exchanges' fee structures are different from bookmakers: they charge a fee on winning bets, as opposed to adjusting the original odds. Our solution incorporates commission. This may be set to zero, and so can solve traditional bookmaking bets too.

Examined assumptions Kelly assumes that money is infinitesimally divisible. In research on financial markets it is usually assumed that the portfolio is so large that such an assumption need not be of concern. In betting, placing very large bets is not possible and there is a relatively large minimum bet size. We examine

the probability that the minimum bet size will cause problems for various sizes of initial wealth.

Examined dynamic hedging One of the attractions of betting exchanges when compared to traditional bookmakers is the ease with which positions can be traded in and out. We studied which additional bets should be placed having first placed a back bet, both when odds increase (which has been considered before) and also when they decrease (which we have not seen researched).

Inverted Kelly for market making When looking for opportunities to test some of these ideas it became apparent that in leading markets the prices are very competitive, certainly better than the available models. In the less competitive markets - lower football divisions, for example - there are often no, or very few prices. It becomes necessary to make prices in the hope that others will trade them. We use Kelly, to choose the odds we wish to trade.

The main applications of this work will be for gambling. Indeed the paper [14] which we published in 2013 has been picked up by SmarterSig. SmarterSig is a subscription based website for people interested in gambling research. They have adopted the work from Chapter 3 and put it into a spreadsheet [13]. They are currently working on incorporating Chapter 4 into their spreadsheet. In addition to consumers the research presented in this thesis may also be of interest to bookmakers.

The economic similarity between betting and insurance is a very old idea. Bernoulli highlighted it at least 275 years ago (see Section 2.1.1). The reality for most of us, however, is that insurance is a game we have to play, or at least cannot afford not to: we have to buy buildings insurance, because the bank demands it, and car insurance because it is a legal requirement. In any case, for most types of insurance we do not have the requisite information to calculate the risks, and obtaining it is beyond our means. However, it is possible to think of insurance applications which might lend themselves to more detailed analysis. Perhaps the organiser of outdoor events is considering rain insurance. It should be possible to use some of the ideas contained in this thesis to assist with making such a decision. We do not claim to have any knowledge of the underwriting industry, but it is possible that these ideas might be useful for underwriters.

There are possible financial applications too, although we would expect most banks to have proprietary systems specifically designed for financial markets. It is possible that the market making which is independent from the model may be combined with other risk systems to provide a safer risk framework.

8.2. Should We Follow Kelly?

Anyone considering using the ideas presented here, whether for gambling, insurance or something else needs to be happy to Maximise the Expected Log (MEL) and so be considered a Kelly investor. This can be a very volatile strategy. Many Kelly adherents actually bet only a fraction of the Kelly stake. We include fractional Kelly betters in what we term Kelly betters. Even following a fractional strategy can be volatile so understanding is required, and there are many other reasons not to follow such a strategy.

One of Samuelson's criticisms is that an investor does not know what is going to happen next. In Chapter 6, we consider this point. We assume that a back bet has been placed and then allow the odds to rise. We are able to calculate which bet should be placed next. Of course if we knew that the odds were going to rise the correct strategy would have been not to place the first bet, but place a larger bet once the odds have risen. Market knowledge and experience are valuable and we believe that these combined with Kelly will outperform either individually. It is possible that a completely different opportunity will present itself, and having invested a sizable proportion of wealth in the first there is little available for the second. Finding one good opportunity is hard enough, we would be truly fortunate to find ourselves in the position of having two, but if one is expecting a better trade to occur in the near future it might be prudent to wait.

Most of the activity described in this thesis is akin to investing. The probabilities are calculated and odds are examined and those that seem to be cheap are backed. When we use knowledge of how prices move over time our activity becomes increasingly like trading: we buy something now in the hope that we can sell it at a better price in the future. Anyone following a trading strategy might be able to justify betting much greater amounts than the Kelly fraction as they do not intend to continue to hold such a position at the time of the match.

It is often hard to assess with a high degree of accuracy the probabilities needed by the Kelly criterion. A small error in calculating these could lead to significant over betting and ruin. The betting market (including betting exchanges and bookmakers) on leading sports is active and in the vast majority of cases any improvements will be very small. In some cases calculating the probabilities will be much easier. When Thorp used Kelly staking for his Blackjack strategy it was easy to know the probabilities of a certain outcome. Even then he found that some casinos cheated. A recent MSc thesis [41] which tries to recreate Dixon and Coles' paper from 1997 but taking odds from betting exchanges, found that when the difference in price between the market and the model was greatest the market was more likely to be correct. It is these very times that the size of the bet would have been largest. This leads us to potential future work.

8.3. Future Work

One of the shortcomings of Kelly is the level of accuracy needed in calculating the probabilities, it would be interesting to refine it still further in order to take into account some uncertainty in the probability distribution. The system would then place relatively smaller bets (possibly zero in size) if the uncertainty is high for a given probability.

It is possible that having invested in one opportunity another will present itself or more likely, having invested a small amount in a diverse range of opportunities some more will appear. Across Europe many football leagues play at the weekend. How much money should a gambler on football matches allocate on a Friday? Will improved situations be available on Saturday morning. Others have done work on independent bets [158], [40]. It should be possible to combine their work with the work here. Further, as the betting exchanges are quite liquid, it should be possible to reduce existing bets (at a small cost) to place money into any new opportunity. This would remove Samuelson's criticism in the case of liquid markets.

When commission has been included in Chapter 4 we no longer claim that the problem is convex and so we no longer claim that a unique global maximum will be found. Having run some tests and simulations nothing has indicated that a sub optimal local minimum is ever found. Perhaps the problem is still convex, or some other second order sufficient conditions exist.

An improved equation to calculate the probability that Kelly does lead to ruin in practice might be found, and certainly it should be possible to improve the termination criterion.

Several improvements could be made to the main method, each small but together significantly improving the usefulness. Extra constraints should be added to allow for other issues that a gambler may have, such as the amount of cash deployed. Perhaps there should be a cost of cash included. Additional betting exchanges should be added. Allowance should be made for a match ending early, or in an unusual way. For example, if a football match ends early because of crowd trouble, or bad weather, some bets will be settled and some will be voided. If we had laid $0-0$, $1-0$, $0-1$, $2-0$, $1-1$, and $0-2$ and backed *under 2.5 goals*, and the floodlights fail at half-time with the score 3-0, we would lose a significant amount of money. The (winning) correct score bets would be voided, and the over 2.5 goals bet would be lost. A final useful addition would be to include previous bets in different markets which it might be worth unwinding or reducing if new opportunities present themselves.

These further refinements would provide an exceptional tool-kit for a gambler. Even without these improvements we feel that following our improved version of fractional Kelly will often prove optimal. As an addition to market knowledge rather than a substitute the work presented here will enhance the returns of a systematic gambler betting on an exchange.

A. Glossary

A.1. Guide to Betting Terminology

Using the language from Chapter 3 we bet on the events of a game. The game ends with a final outcome. Each event is a set of outcomes. If an event contains the final outcome we call the event winning. When placing a bet, one party places a back bet, and the other places a lay bet. If the event/bet is winning the back bet wins, otherwise the lay bet wins. Excluding commission the sum of winnings and losses of the back bet and lay bet is zero. The person placing the back bet on an event is often called the backer. He might be said to be backing an event. Similarly his counter-party is the layer and might be said to be laying the event.

In a simple example the game might be a horse race. In this case each event would contain one outcome – a horse. The final outcome would be the winning horse. If you back a horse, e.g. *Frankel* you are betting that it wins. You might be said to be backing Frankel. The person you place that back bet with, whether a bookmaker, an exchange or a friend in your office is laying the bet. He is laying Frankel, or equivalently backing all other horses. To make it slightly more complicated we can include place bets. To win a place bet your horse needs to finish in the first few, typically 2, 3 or 4 depending upon the number of runners. For this example we will assume the first two places. Now the outcomes are all ordered pairs of horses. The first horse is the winner and the other named horse finishes second. Now when betting on Frankel to win, the event contains all pairs of horses with the first horse named Frankel. If betting on Frankel for a place these events would contain all pairs of horses that contain Frankel as either first or second named horse.

In addition to winning or losing there is sometimes a third possibility. A bet may be voided. If a horse does not start, or if a football match is postponed before kick off then in many circumstances all bets will be voided. Thus, there will be no winners or losers and all stake money is returned. Ante post bets on horse races are unusual in this regard as non running horse count as losing bets. For most other sports there are some common themes, but specific rules do vary. And even within a sport some of the bets specifically exclude voiding in certain cases. For example in a tennis match between Player A and Player B, there will be many bets, but exchanges might list Player A to win bet, and Player A to get through to the next round. If Player B withdraws before

the start of the match with an injury the first of those two will be voided, but not the second. Some researchers [163] are examining the differences in odds in such markets to infer the (market's view of the) probability of withdrawal.

In the previous paragraph we personified the betting exchanges or bookmakers' prices. One of the founders of Betfair was a trader in the financial markets [164], and, therefore it is not surprising that some of the language of financial markets had transferred to betting exchanges. In a financial market there is normally a buy, or bid price, and a sell, or ask price. Similarly the betting exchange has back prices and lay prices, where the term price is used interchangeably with odds. In financial markets it is often common to discuss the market's view of something – meaning the probability inferred by the money weighted average of relevant prices. We often talk of financial markets in general without specifying an exchange. In some cases it will be obvious which exchange we intend and in some cases there will be a choice of exchange. The same is true of betting markets. In this thesis if we do not otherwise specify we are discussing Betfair, which is the current market leader in terms of turnover.

In betting markets, as with financial markets friction costs can be significant. Friction costs include commission, other fees and the bid-ask spread, which is the difference between the sell price and the buy price. The bid-ask spread in financial markets is sometimes used as a measure of liquidity [165]. In betting markets we have an advantage because we know that the fair odds would produce implied probabilities that sum to one. So in addition to the bid-ask spread we can examine the sum of the implied probabilities. The distance from one is the over-round. In decimal odds the implied probability is the reciprocal of the odds. The state when the bid-ask spread turns negative is termed a backwardation. On many exchanges it is not allowed to happen directly, but it may be that the bid odds on one exchange are higher than the lay odds on another exchange. Sometimes it is possible to create a bet from a combination of other bets. In two runner markets, for example, it is possible to create lay odds on one runner using the back odds of the other.

Glossary

Back To bet for something, or to place a back bet. As opposed to Lay.

Back bet A bet that the named event or runner will win. Traditionally the only option available in the UK.

Backwardation In financial markets a backwardation is said to occur when a financial asset can be bought and simultaneously sold for a profit. In betting we use it when an event can be backed at one price and effectively laid at a lower price.

Balanced book A book is balanced, for a particular race, if its value is independent of the winner of that race. A bookmaker will try to create a balanced book or failing that reduce the biggest possible loss.

Better A person who places a bet. We use this interchangeably with gambler.

Commission A betting exchange does not provide odds and generates revenue by charging users commission. Betfair charges a fraction of winnings on each market.

Decimal odds Decimal odds are common on betting exchanges and with bookmakers in some countries outside the UK. These odds refer to the amount of money returned to the gambler per unit stake on a win. As decimal odds are higher (by one unit) than traditional odds for a given return the bookmaking industry is considering switching to decimal odds.

Event An event is a set of outcomes.

Fair odds Odds offered without an over-round, such that the sum of the implied probabilities is one.

Fractional Kelly Kelly staking can be a very volatile strategy and so many implicitly or explicitly bet only a fraction of their wealth. We include fractional Kelly when we refer to Kelly betting.

Gambler A person who places a bet. We use this interchangeably with better.

Game We define a game as a future contest.

GOP Growth Optimal Portfolio, is another name for a Kelly portfolio (see also MEL).

Growth rate Throughout this thesis, we use growth rate to be the asymptotic, per game, (natural) logarithmic growth rate of wealth. (Kelly used base 2 logarithms as did a few of the other writers in this field, who were also interested in the relationship with information theory.) If V_0 is the initial wealth and V_N is the wealth after N games (matches,races,etc.), then as $N \rightarrow \infty, G \rightarrow \frac{1}{N} \log \frac{V_N}{V_0}$. Alternatively if a game has T possible outcomes and outcome j has probability p_j of occurring and changes wealth by a factor WF_j , then $G = \log \prod_T WF_j^{p_j}$.

In play Betting on a match or race whilst it is in progress. This is possible on many of the more popular games. Also called in running.

Lay To bet against something, or to place a lay bet. As opposed to Back.

Lay bet A bet that the named event or runner will lose. Nearly every bet that a traditional bookmaker writes is a lay bet, as their customers place back bets.

Matching When a bet is placed into a betting exchange if the opposite bet is present both will trade. This is called matching. A bet may match in full, partially, or not at all. A bet which did not match at first may match later.

MEL Maximising the Expected Log, the name more often used for Kelly portfolios by economists (see also GOP).

Odds The odds are the ratio of the winning payout to the losing payout. There is more than one way to calculate them. In the UK two methods are common, see fractional odds and decimal odds.

Offset It is only possible to place a bet on a betting exchange if the member has enough money in his account. It is possible that less money will be needed for subsequent bets if the exchange allows offsetting. For example, placing a £10 lay bet in a horse race at odds of 4 requires £30, placing a second £10 bet on a different horse at odds of 5 requires no additional cash. Betfair allows offsetting within markets, but not between markets.

Outcome We define outcome as a possible result of a game.

Over-round The difference between a bookmaker's odds and the fair odds, often expressed as a %. Historically between 10% and 15%, now as low as 2%.

Spread betting A spread better pays or receives the difference between some pre-arranged value and the final value. He might *buy* total goals in a football match at 2.5 in £10 per goal. If the final score is 4-1 than he wins £25, if the final score

is 1-1 he loses £5. Had he *sold* instead of buying the win and loss would have been reversed. Spread betting existed in the UK long before betting exchanges and has many of the advantages of betting exchanges: it is possible to buy and sell, it is possible to unwind before the start of the match, and betting was allowed during the match long before it was introduced by betting exchanges to fixed odds betting..

Trader We use the term trader to describe someone who trades something at one price in the hope that it can be later reversed at a better price. Contrast that with the Kelly bettor who is more akin to a value investor; he backs something because he thinks it is more likely to win than the probability implied by the odds.

Under-round The difference between an exchange's lay odds and the fair odds, often expressed as a %. Unlike over-round this is not a universally accepted term.

Unwind A person who places a bet, or a collection of bets, in order to reduce the risk of previous bets is unwinding his book. Having placed a back bet, he might place a lay bet on the same runner, or back all other runners.

B. Common Football Bets

B.1. Bets Determined by Final Score

The following is a list of some of the football bets available on betting exchanges. This thesis refers to the first three only. The rest are provided to demonstrate the diversity of bets available.

Match Odds

This is a bet on the result – the winning team. As a draw is common, one of the available bets is the draw (about $\frac{1}{4}$ of league football matches end in a draw). On Betfair, bets usually refer to the state at the end of normal play (after 90 minutes). If the match is a cup game, or some other type of match which has extra time in the event of a draw, and the match goes into extra time, the bet is settled as a draw.

Correct Score

This is a bet on the final score. As with most other bets this is usually a bet on the score after the normal 90 minutes. There is often a catch all bet for high scores. On Betfair, for example, there are 17 possible bets, 16 bets with specified scores, plus *other*, which wins when either team scores 4 goals or more.

Over/Under XX Goals

This is a bet on the total number of goals scored by both teams in the 90 minutes of play. The XX is normally a number which cannot be achieved such as 2.5. The average number of goals in league football is between 2 and 3 and so Over/Under 2.5 is often the most active of these bets, although bets can exist from 0.5 to higher than 9.5. If it is possible to place bets match during the match (often termed in-play or in-running betting), then as the score rises more markets are introduced.

Both Teams to Score

This is a bet on whether or not both teams will score. Similarly there is usually a bet on whether or not the home team will score, will not let in a goal and so on.

Asian Handicap XX

A bet on which team will win given XX is added to one of the team's final scores. XX can be a whole number or a whole number and a half, in the first case a draw

is possible, but not in the second. The most active market is often that which has both teams equally likely to win. Sometimes bets can be combined as XX & $XX\frac{1}{2}$, in which case a bet is actually a pair of bets.

Draw No Bet

Similar to Match Odds, except that in the case of a draw the bet is voided and stakes are returned.

Total Goals

A bet on the total number of goals, the bets are of the type, 3 or more, 4 or more, and so on.

B.2. Bets not Determinable by Final Score

Bets in Section B.1 can be determined by knowing the final score. There are many other bets, some of which are very esoteric. Some would require a very small change in a model. Some of these require knowledge of the score at half time, whilst others require the order of the goals. Other bets need the name of the player who scored the goal. Yet other more bets require details beyond goals, such as corners or yellow and red cards. Examples are as follows.

Half Time

Just like Match Odds above, but for the state of the match after 45 minutes.

Half Time Score

A Correct Score market for the first 45 minutes only.

Half Time / Full Time

A bet with 9 possible outcomes, one for each of the three first half states, combined with each of the three final results.

Next Goal

Before the start of the match this is a bet on the team which will score the first goal, or possibly none. After the match has started this becomes a bet on the next goal.

Sending Off

Will a player (on either team) be sent off.

Penalty Taken

Will either team take a penalty.

Multi Corners

The number of corners in the first half multiplied by the number of corners in the second half. This is a very volatile bet which some gamblers seem to like.

Bibliography

- [1] U.S. Census Bureau. *Statistical Abstract of the United States 2011*. Government Printing Office, 2011.
- [2] Gambling Commission. Industry statistics April 2009 to September 2012, June 2013.
- [3] National Lottery Commission. *National Lottery Commission Annual Report and Accounts 2010/2011*. London: The Stationary Office, 2011.
- [4] National Lottery Commission. *National Lottery Commission Annual Report and Account 2011/2012*. London: The Stationary Office, 2012.
- [5] John Ashton. *The History of Gambling in England*. Duckworth & Co., London, 1898.
- [6] Nerina J. Calabiano. From antiquity to Australia: A brief account of gambling. *International Journal of Mental Health and Addiction*, 1, 1:16–19, 2003.
- [7] Charlotte Fabiansson. *Pathways to Excessive Gambling: A Societal Perspective on Youth and Adult Gambling Pursuits*. Ashgate Publishing Limited, 2010.
- [8] Daniel Bernoulli. Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1):23–36, 1954 (Translation into English).
- [9] John L. Kelly Jr. A new interpretation of information rate. *The Bell System Technical Journal*, July:917–926, 1956.
- [10] Leo Breiman. Optimal gambling systems for favorable games. *Proceedings Fourth Berkeley Symposium*, 1:65–78, 1961.
- [11] Edward O. Thorp. Optimal gambling systems for favorable games. *Revue de l'Institut International de Statistique / Review of the International Statistical Institute*, 37(3):273–293, 1969.
- [12] Edmund Noon, William J. Knottenbelt, and Daniel Kuhn. Kelly's fractional staking in the twenty first century. In David Percy, J. James Reade, and Phil Scarf, editors, *3rd International Conference on Mathematics in Sport*. Institute of Mathematics and its Applications, June 2011.

- [13] John F. Jackson. Kelly: Selection in same part iii. *www.SmarterSIG.com*, August, 2013.
- [14] Edmund Noon, William J. Knottenbelt, and Daniel Kuhn. Kelly's fractional staking updated for betting exchanges. *IMA Journal of Management Mathematics*, 24:283–299, 2013.
- [15] Edmund Noon and William J. Knottenbelt. Market making with an inverse Kelly strategy. In *4th International Conference on Mathematics in Sport*, 2013.
- [16] Oystein Ore. Pascal and the invention of probability theory. *The American Mathematical Monthly*, 67(5):pp. 409–419, 1960.
- [17] Frank C. Duckworth and Anthony J. Lewis. A fair method for resetting the target in interrupted one-day cricket matches. *The Journal of the Operational Research Society*, 49(3):pp. 220–227, 1998.
- [18] William Poundstone. *Fortune's Formula: The Untold Story fo the Scientific Betting System that Beat the Casinos and Wall Street*. Hill and Wang, 2005.
- [19] Claude E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423 & 623–656, 1948.
- [20] Harry M. Markowitz. Investment for the long run: New evidence for an old rule. *The Journal of Finance*, 31(5):1273–1286, 1976.
- [21] John Haigh. *Taking Chances : Winning with Probability (2nd Edition)*. Oxford University Press, 2003.
- [22] Leo Breiman. Investment policies for expanding businesses optimal in a long-run sense. *Naval Research Logistics Quarterly*, 7(4):647–651, 1960.
- [23] Edward O. Thorp. *Beat The Dealer*. Random House, 1962.
- [24] Lester E. Dubins and Leonard J. Savage. *How to Gamble If You Must: Inequalities for Stochastic Processes*. McGraw-Hill, 1965.
- [25] INFORMS John von Neumann prize winners. <https://www.informs.org/Recognize-Excellence/INFORMS-Prizes-Awards/John-von-Neumann-Theory-Prize>. Accessed 22/07/2013.
- [26] Nobel prize (economics) winners. http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/1990/markowitz-facts.html, 1990. Accessed 29/08/2013.
- [27] Harry M. Markowitz. Portfolio selection*. *The Journal of Finance*, 7(1):77–91, 1952.
- [28] Harry M. Markowitz. *Portfolio Selection: Efficient Diversification of Investments*. John Wiley & Sons, Inc., 1959.

- [29] William F. Sharpe. Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3):425–442, 1964.
- [30] Merton H. Miller and Franco Modigliani. Dividend policy, growth, and the valuation of shares. *The Journal of Business*, 34(4):pp. 411–433, 1961.
- [31] Eugene F. Fama. The behavior of stock-market prices. *The Journal of Business*, 38(1):pp. 34–105, 1965.
- [32] Mike Dempsey. The capital asset pricing model (CAPM): The history of a failed revolutionary idea in finance? *Abacus*, 49:7–23, 2013.
- [33] Tom Smith and Kathleen Walsh. Why the CAPM is half-right and everything else is wrong. *Abacus*, 49:73–78, 2013.
- [34] Henry Allen Latané. Criteria for choice among risky ventures. *The Journal of Political Economy*, 67(2):144–155, 1959.
- [35] Harry M. Markowitz. Investment for the long run. *unpublished*, 1972. Obtained from <http://finance.wharton.upenn.edu/~rlwctr/papers/7220.PDF> Accessed 29/08/2013.
- [36] Morten Mosegaard Christensen. On the history of the growth optimal portfolio. 2005.
- [37] Edward O. Thorp. The Kelly criterion in blackjack, sports betting and the stock market. In *The 10th International Conference on Gambling and Risk Taking*, 1997.
- [38] Louis M. Rotando and Edward O. Thorp. The Kelly criterion and the stock market. *The American Mathematical Monthly*, 99(10):922–931, 1992.
- [39] Robert C. Merton. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3:373–413, 1971.
- [40] Edward O. Thorp. Understanding the Kelly criterion – part ii. *Wilmott Magazine*, 2008.
- [41] Edmund Noon. A framework to evaluate models of association football, statistically and with market data. Master’s thesis, Imperial College London, 2009.
- [42] Wikipedia. Proebsting’s paradox. http://en.wikipedia.org/wiki/Proebsting%27s_paradox. Accessed 29/08/2013.
- [43] Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637, 1973.

- [44] Lixin Zeng. Pricing weather derivatives. *The Journal of Risk Finance*, 1:72–78, 2000.
- [45] Peter Alaton, Boualem Djehiche, and David Stillberger. On modelling and pricing weather derivatives. *Applied Mathematical Finance*, 9:1–20, 2002.
- [46] Paul A. Samuelson. The "fallacy" of maximizing the geometric mean in long sequences of investing or gambling. *Proceedings National Academy Sciences USA*, 68, No. 10:pp 2493–2496, October 1971.
- [47] Robert C. Merton and Paul A. Samuelson. Fallacy of the log-normal approximation to optimal portfolio decision-making over many periods. *Journal of Financial Economics*, 1(1):67 – 94, 1974.
- [48] Nils H. Hakansson. Capital growth and the mean-variance approach to portfolio selection. *The Journal of Financial and Quantitative Analysis*, 6(1):517–557, 1971.
- [49] Nils H. Hakansson. Comment on Merton and Samuelson. *Journal of Financial Economics*, 1(1):95 –, 1974.
- [50] Edward O. Thorp. Portfolio choice and the Kelly criterion. In *Business and Economics Statistics Proceedings, American Statistical Association*, 1971.
- [51] Edward O. Thorp and Robert Whitley. Concave utilities are distinguished by their optimal strategies. In Grani, Sarkadi, and Vincze, editors, *Proceedings of the European Meeting of Statisticians, Budapest 1972*, pages 813–830, 1974.
- [52] Robert M. Bell and Thomas M. Cover. Competitive optimality of logarithmic investment. *Mathematics of Operations Research*, 5(2):161–166, 1980.
- [53] M. Barry Goldman. A negative report on the near optimality of the max-expected-log policy as applied to bounded utilities for long lived programs. *Journal of Financial Economics*, 1(1):97–103, 1974.
- [54] David G. Luenberger. A preference foundation for log mean-variance criteria in portfolio choice problems. *Journal of Economic Dynamics and Control*, 17(5-6):887 – 906, 1993.
- [55] Paul A. Samuelson. Risk and uncertainty: A fallacy of large numbers. *Scientia*, 98(4-5):108–113, 1963.
- [56] Robert Bell and Thomas M. Cover. Game-theoretic optimal portfolios. *Management Science*, 34(6):724–733, 1988.
- [57] Paul A. Samuelson. Why we should not make mean log of wealth big though years to act are long. *Journal of Banking & Finance*, 3(4):305–307, 1979.

- [58] Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. John Wiley & Sons, 1991.
- [59] Thomas M. Cover. Universal portfolios. *Mathematical Finance*, 1:1–29, 1991.
- [60] David P. Helmbold, Robert E. Schapire, Yoram Singer, and Manfred K. Warmuth. On-line portfolio selection using multiplicative updates. *Mathematical Finance*, 8(4):325–347, 1998.
- [61] Avrim Blum and Adam Kalai. Universal portfolios with and without transaction costs. *Machine Learning*, 35(3):193–205, 1999.
- [62] Burton Malkiel. *A Random Walk Down Wall Street*. W. W. Norton & Company, Inc., 1973.
- [63] Warren Buffett. Letter to partners: Performance in 1965. Letter, January 1966.
- [64] Phellm Boyle, Lorenzo Garlappi, Raman Uppal, and Tan Wang. Keynes meets Markowitz: The trade-off between familiarity and diversification. *Management Science*, 58 no. 2:253–272, 2012.
- [65] Benjamin Graham and David Dodd. *Security Analysis*. McGraw-Hill, 1934.
- [66] Michael J. Moroney. *Facts from figures*. Pelican, 1951.
- [67] Ladislaus von Bortkiewicz. *Das Gesetz der kleinen Zahlen [The law of small numbers]*. B.G. Teubner, 1898.
- [68] Joan del Castillo and Marta Prez-Casany. Overdispersed and underdispersed poisson generalizations. *Journal of Statistical Planning and Inference*, 134(2):486 – 500, 2005.
- [69] Bobby McMahon. The most important soccer performance analyst you have never heard of. *Forbes Magazine*, 21 October, 2012.
- [70] Oyvind Larson. Charles Reep: A major influence on British and Norwegian football. *Soccer and Society*, 2(3):58–78, 2001.
- [71] Charles Reep and B. Benjamin. Skill and chance in association football. *Journal of the Royal Statistical Society. Series A (General)*, 131(4):581–585, 1968.
- [72] Charles. Reep, Richard Pollard, and B. Benjamin. Skill and chance in ball games. *Journal of the Royal Statistical Society. Series A (General)*, 134(4):623–629, 1971.
- [73] J. Hennessy. Pins are better than form. *The Times*, March 21st:10, 1969.
- [74] I. D. Hill. Association football and statistical inference. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 23(2):203–208, 1974.

- [75] Michael J. Maher. Modelling association football scores. *Statistica Neerlandica*, 36 nr. 3.:109–118, 1982.
- [76] Mark Thompson. On any given Sunday: Fair competitor orderings with maximum likelihood methods. *Journal of the American Statistical Association*, 70(351):536–541, 1975.
- [77] David Harville. The use of linear-model methodology to rate high school or college football teams. *Journal of the American Statistical Association*, 72(358):pp. 278–289, 1977.
- [78] David Harville. Predictions for National Football League games via linear-model methodology. *Journal of the American Statistical Association*, 75(371):516–524, 1980.
- [79] Stephen R. Clarke and John M. Norman. Home ground advantage of individual clubs in English soccer. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 44(4):509–521, 1995.
- [80] Richard Pollard. 69.9 goal-scoring and the negative binomial distribution. *The Mathematical Gazette*, 69(447):45–47, 1985.
- [81] Richard C. Griffiths and Richard K. Milne. A class of bivariate Poisson processes. *Journal of Multivariate Analysis*, 3:380, 1978.
- [82] Mark J. Dixon and Stuart G. Coles. Modelling association football scores and inefficiencies in the football betting market. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 46(2):265–280, 1997.
- [83] V. Barnett and S. Hilditch. The effect of an artificial pitch surface on home team performance in football (soccer). *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 156(1):pp. 39–50, 1993.
- [84] Richard Pollard and G. Pollard. Long-term trends in home advantage in professional team sports in North America and England (18762003). *Journal of Sports Sciences*, 23(4):337–350, 2005. PMID: 16089179.
- [85] John Goddard. Regression models for forecasting goals and match results in association football. *International Journal of Forecasting*, 21:331–340, 2005.
- [86] David Forrest and Robert Simmons. Forecasting sport: the behaviour and performance of football tipsters. *International Journal of Forecasting*, 16(3):317–331, 2000.
- [87] David Forrest and Robert Simmons. Focus on sport: Making up the results: the work of the football pools panel, 1963-1997. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 49(2):253–260, 2000.

- [88] Rob Smyth and Georgina Turner. Has anyone ever gone down with a positive goal difference? *The Guardian*, 5th May, 2004.
- [89] Michael Lewis. *Moneyball: The Art of Winning an Unfair Game*. 2003.
- [90] Andy Hunter. Manchester City to open the archive on player data and statistics. *The Guardian*, 16th August, 2012.
- [91] F. Bonomo, G. Duran, and J. Marenco. A virtual coach for fantasy soccer using mathematical programming. In *Proceedings of the 4th International Conference on Mathematics in Sport*, 2013.
- [92] J. Beliën, D. Goossens, and D. Van Reeth. Optimization modeling for analyzing fantasy sport games. In *Proceedings of the 4th International Conference on Mathematics in Sport*, 2013.
- [93] Iain McLean. A non-zero-sum game of football. *British Journal of Political Science*, 10(2):253–259, 1980.
- [94] Lawrence Booth and Robert Smyth. What’s the dodgiest game in football history? *The Guardian*, 11th August, 2004.
- [95] BBC. Bookmaker may stop football bets. http://news.bbc.co.uk/sport1/hi/front_page/8093133.stm, June 2009. Accessed 27/07/2013.
- [96] Steven D. Levitt and Stephen J. Dubner. *Freakonomics*. Penguin, 2005.
- [97] Mark Duggan and Steven D. Levitt. Winning isn’t everything: Corruption in sumo wrestling. Working Paper 7798, National Bureau of Economic Research, July 2000.
- [98] Wikipedia. History of match fixing. http://en.wikipedia.org/wiki/Match_fixing. Accessed 29/08/2013.
- [99] Declan Hill. How gambling corruptors fix football matches. *European Sport Management Quarterly*, 9(4):411–432, 2009.
- [100] Raul Caruso. The basics economics of match-fixing in sport tournaments. *Economic Analysis and Policy*, 39:355–377, 2007.
- [101] David Forrest and Robert Simmons. Sport and gambling. *Oxford Review of Economic Policy*, 19(4):598–611, 2003.
- [102] F. Carmichael, G. Rossi, and D. Thomas. Team performance in the Italian Serie A: 2000/01 - 2009/10. In *Proceedings of the 4th International Conference on Mathematics in Sport*, 2013.
- [103] J. James Reade and Sachiko Aike. Using forecasting to detect corruption in

- international football. In *Proceedings of the 4th International Conference on Mathematics in Sport*, 2013.
- [104] David Forrest. The threat to football from betting-related corruption. *International Journal of Sport Finance*, 7(2):99–116, May 2012.
- [105] Margaret H. Wright. The interior-point revolution in optimization: History, recent developments, and lasting consequences. *Bulletin (New Series) of the American Mathematical Society*, 42, No. 1:39–56, 2004.
- [106] Narendra Karmarkar. A new polynomial time algorithm for linear programming. *Combinatorica*, 4:373395, 1984.
- [107] George B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, 1963.
- [108] Victor Klee and George J. Minty. How good is the simplex algorithm? In Oved Shisha, editor, *Inequalities III (Proceedings of the Third Symposium on Inequalities held at the University of California, Los Angeles, Calif., September 19, 1969, dedicated to the memory of Theodore S. Motzkin)*, 1972.
- [109] Leonid G. Khachiyan. A polynomial algorithm in linear programming. 1979.
- [110] Andreas Wächter and Lorenz T. Biegler. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, 2006.
- [111] Andreas Wächter. *An Interior Point Algorithm for Large-Scale Nonlinear Optimization with Applications in Process Engineering*. phd, Carnegie Mellon University, Pittsburgh, PA, January 2002.
- [112] Andreas Wächter and Lorenz T. Biegler. Line search filter methods for nonlinear programming: Motivation and global convergence. *SIAM Journal on Optimization*, 16(1):1–31, 2005.
- [113] Andreas Wächter and Lorenz T. Biegler. Line search filter methods for nonlinear programming: Local convergence. *SIAM Journal on Optimization*, 16(1):32–48, 2005.
- [114] David Forrest and Robert Simmons. Globalisation and efficiency in the fixed-odds soccer betting market. Centre for the Study of Gambling and Commercial Gaming, University of Salford, (Working paper), December 2001.
- [115] David Forrest. Online gambling: An economic perspective. In *Routledge International Handbook of Internet Gambling*. Routledge, 2012.
- [116] Roger Munting. *An Economic and Social History of Gambling in Britain and the USA*. Manchester University Press, 1996.

- [117] Bruce C. Daniels. *Puritans At Play: Leisure and Recreation in Early New England*. Palgrave Macmillan, 1995.
- [118] UK Parliament. *Gaming Act 1845*. London: The Stationary Office, 1845.
- [119] UK Parliament. *Betting Act 1853*. London: The Stationary Office, 1853.
- [120] UK Parliament. *Betting and Lotteries Act 1934*. London: The Stationary Office, 1934.
- [121] UK Parliament. *Betting and Gaming Act 1960*. London: The Stationary Office, 1960.
- [122] Stan Hey. Our national love affair: a history of the betting shop. *The Independent*, 05/04, 2008.
- [123] Ladbrokes. Ladbrokes Plc. History.
<http://www.ladbrokesplc.com/about-ladbrokes/history.aspx>. Accessed 29/07/2013.
- [124] Gambling Consultant (website).
<http://www.gamblingconsultant.co.uk/the-gaming-act-1968-executiv/>. Including content from the PhD thesis of Dr Seamus Murphy. Accessed 29/07/2013.
- [125] UK Parliament. *Gaming Act 1968*. London: The Stationary Office, 1968.
- [126] House of Commons Culture, Media and Sport Committee. *The Gambling Act 2005: A bet worth taking?* London: The Stationary Office, 2012.
- [127] Wray Vamplew and Joyce Kay. *Encyclopedia of British Horse Racing*. Routledge, 2005.
- [128] Gambling Compliance Ltd. European regulated online markets: Data report, 2012.
- [129] UK Parliament. *Financial Services and Markets Act 2000*. London: The Stationary Office, 2000.
- [130] UK Parliament. *Gambling Act 2005*. London: The Stationary Office, 2005.
- [131] George Curry. The national lottery a critique of state sponsored gambling. The Christian Institute, 1996.
- [132] Betfair annual report 2012.
- [133] The World Bank Group. World data: World development indicators.
<http://databank.worldbank.org/data/views/reports/tableview.aspx>.
- [134] National Centre for Social Research. British gambling prevalence survey 2010. The Gambling Commission, 2011.

- [135] Douglas Seay. Gambling and crime. In *National Gambling Impact Study Commission*, On-site meeting in Biloxi, Mississippi and New Orleans, Louisiana, 1998.
- [136] Haringey Overview and Scrutiny Committee. Scrutiny review of the clustering of betting shops in Haringey.
http://www.haringey.gov.uk/final_report_pdf-5.pdf, April 2011.
 Accessed 27/07/2013.
- [137] Nick Harris. Exclusive: Gambling is ‘bigger threat to sport than doping’. *The Independent*, 06th Mar, 2009.
- [138] Gambling Commission. The year in focus: Annual review 2012/13, 2013.
- [139] Amelia Gentleman. Roulette machines: the crack cocaine of gambling. *The Guardian*, 27th May, 2013.
- [140] Adam Daniel Rivers. The fool and the thief: A study of the UK gambling industry. Norwich Economic Papers (volume 2), UEA, 2011.
- [141] Home Affairs Committee. Knife crime. London: The Stationary Office, June 2009.
- [142] Mike Godwin. *Cyber Rights: Defending Free Speech in the Digital Age*. Times Books, 1998.
- [143] Emma Hern, Will Glazebrook, and Mike Beckett. Reducing knife crime: We need to ban the sale of long pointed kitchen knives. *BMJ: British Medical Journal*, 330(7502):1221, 2005.
- [144] Alice K. Ross. Why self-exclusion is not an answer to problem gambling. *The Bureau of Investigative Journalism*, 24th Jul, 2012.
- [145] Dan Robson. Facial recognition: a system problem gamblers can’t beat? *Toronto Star Newspapers Ltd.*, 12th Jan, 2011.
- [146] Alistair Osborne. Bookies shops are not a ‘blight on the high street’. *The Telegraph*, 12th Feb, 2012.
- [147] Niall O’Connor. A history of the betting exchange industry. *Bettingmarket.com*, 2013.
- [148] Des Laffey. Entrepreneurship and innovation in the UK betting industry: The rise of person-to-person betting. *European Management Journal*, 23(3):351 – 359, 2005.
- [149] Betfair case study: From startup to world’s largest betting exchange in ten years. *Marketline*, December:1 – 18, 2011.

- [150] Greg Wood. Flutter’s departure leaves bitter taste. *The Guardian*, 16th January, 2002.
- [151] Steve Easton and Katherine Uylangco. An examination of in-play sports betting using one-day cricket matches. *Journal of Prediction Markets*, 1(2):93–109, July 2007.
- [152] Steve Wilson. Matthew Le Tissier admits being part of attempted betting scam at Southampton. *The Telegraph*, 03rd Sep, 2009.
- [153] Betting On. History of bookmakers. <http://www.bettingon.com/bookmakers>. Accessed 27/07/2013.
- [154] John Gandar, Richard Zuber, Thomas O’Brien, and Ben Russo. Testing rationality in the point spread betting market. *The Journal of Finance*, 43(4):995–1008, 1988.
- [155] Mark J. Dixon and Michael E. Robinson. A birth process model for association football matches. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 47(3):523–538, 1998.
- [156] Harold W. Kuhn and Albert W. Tucker. Proceedings of the second Berkeley symposium on mathematical statistics and probability. 1951.
- [157] Dimitris Karlis and Ioannis Ntzoufras. Bayesian modelling of football outcomes: using the skellam’s distribution for the goal difference. *IMA Journal of Management Mathematics*, 20 (2):133–145, 2009.
- [158] Robin Insley, Lucia Mok, and Tim Swartz. Issues related to sports gambling. *The Australian and New Zealand Journal of Statistics*, 46:219–232, 2004.
- [159] Stephen Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge University Press, 2006 reprint with corrections, first published in 2004 edition, 2006.
- [160] Andreas Wächter and Lorenz T. Biegler. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106:25–57, 2006.
- [161] Betfair. www.betfair.com.
- [162] Nassim Nicholas Taleb. *The Black Swan*. Allen Lane, 2007.
- [163] Adam C.J. Cutmore and William J. Knottenbelt. Quantitative models for retirement risk in professional tennis. In *Proceedings of the 4th International Conference on Mathematics in Sport*, 2013.
- [164] Jody Clarke. Andrew Black: punter who revolutionised gambling. *MoneyWeek*, 21/08/2009, 2009.

- [165] Yakov Amihud and Haim Mendelson. Asset pricing and the bid-ask spread.
Journal of Financial Economics, 17(2):223 – 249, 1986.