

Efficiency of the Football Betting Market

Master's Thesis

Applied Economics and Finance
Copenhagen Business School, 2017

Ketil Albertsen

Student

Steffen Andersen

Supervisor

Number of pages: 54

Number of characters including spaces: 103,754

Date of submission: 2017-01-15

Abstract

This thesis looks into the efficiency of the football betting markets and tries to establish if there exist weak and semi-strong form inefficiencies.

Weak form inefficiency in the form of a home favorite-longshot bias is observed, where home wins with a market estimated probability of more than 75% produce abnormal returns.

It is shown that closing odds are more correct than opening odds, and that the market absorbs information when new information in the form of match results becomes available. Furthermore, it is demonstrated that for a subset of matches, the information contained in the most recent match results of the opposing teams is undervalued. Semi-strong inefficiency is observed through the use of a prediction model earning significant profits when taking this undervaluation into account.

Acknowledgments

This thesis completes my Master of Science degree in Applied Economics and Finance from Copenhagen Business School.

I would like to thank my supervisor Steffen Andersen for his help, support and ideas. Furthermore, I would like to thank Jo Saakvitne for being a valuable discussion partner throughout the work on this thesis.

A special thanks goes to my friends and family for all their support and input.

Table of Contents

List of Tables	4
List of Figures	5
1 Introduction	6
1.1 Motivation and research questions	7
1.2 Results and findings	8
1.3 Contribution to existing literature	8
1.4 Structure	9
2 Background	10
2.1 Market efficiency	10
2.1.1 Joint hypothesis problem	11
2.2 Betting markets	11
2.2.1 Bookmakers and prediction markets	12
2.2.2 Margin development	13
2.3 Betting market efficiency	13
2.3.1 Weak form efficiency	14
2.3.2 Semi-strong form efficiency	15
2.3.3 Strong form efficiency	16
2.4 Odds and probabilities	17
2.4.1 Odds notations	17
2.4.2 Margins	18
2.4.3 Adjusted probabilities	19
2.4.4 Types of bets	19
2.4.5 Staking strategies	20
2.5 Prediction models	21
2.5.1 Relevant factors	21
2.5.2 Rating systems	22
2.5.3 Evaluation criteria	22
3 Research methodology	24
3.1 Data	24
3.1.1 Data sources	24

3.2	Prediction model	25
3.2.1	Regression model	26
3.2.2	Pi-rating	28
3.2.3	Additional regressors	30
3.3	Market assessment	32
3.3.1	Evaluation criteria	33
3.3.2	Odds representation	33
3.3.3	Market bias	33
3.3.4	Prediction model profitability	34
3.3.5	Market development	35
3.3.6	Information uptake	35
4	Results and findings	37
4.1	Prediction model	37
4.1.1	Pi-rating	37
4.1.2	Regression model	38
4.2	Market assessment	40
4.2.1	Market bias	40
4.2.2	Prediction model	42
4.3	Information uptake	43
4.3.1	Price movement	43
4.3.2	Opening and closing odds	44
4.3.3	Information signal	45
5	Discussion	48
5.1	Results	48
5.1.1	Bias	48
5.1.2	Prediction model	49
5.1.3	Information signal	50
5.1.4	Information uptake	50
5.2	Efficiency	51
5.2.1	Weak form efficiency	51
5.2.2	Semi-strong form efficiency	52
5.3	Information overlap	52
5.4	Limitations and weaknesses	53
6	Conclusion	54
6.1	Main findings	54

6.2	Recommendations for further research	54
	References	55
	A Appendix	59
A.1	Pi-rating calibration	59
A.2	Regression model 1 – 2011	60
A.2.1	Initial specification	60
A.2.2	First stepwise model	60
A.2.3	Final model	61
A.3	Regression model 2 – 2011	62
A.3.1	Initial specification	62
A.3.2	Final model	62
A.4	Regression model evolution	64
A.5	Opening and closing odds	65
A.6	Price change/match result correlation	66
A.7	Odds change before and after new results	67

List of Tables

2.1	Probability and odds examples	17
2.2	Odds example comparing bookmakers and prediction markets	19
4.1	Evaluated independent variables	38
4.2	Mean profit from a bet on matches grouped by outcome	41
4.3	Mean profit from a bet on matches grouped by odds	41
4.4	Mean profit on prediction model based bets	42
4.5	Mean return using first available and closing odds on profitable strategies	44
4.6	Strategy summary	47
A.1	Mean squared error for pi-ratings	59
A.2	Model coefficient values based on training data from different seasons	64
A.3	Mean opening and closing probabilities formatted as odds	65
A.4	Mean opening and closing probabilities for home wins grouped by opening odds	65
A.5	Odds distribution for information based home win strategy	65

List of Figures

2.1	Development of margins in the English Premier League for selected bookmakers	13
4.1	Development of pi-rating for Manchester United	37
4.2	Cumulative probability distribution for pi.diff in model 1, season 2011	39
4.3	Margin and volume relative to their respective maximum/minimum values	43
4.4	Cumulative returns for selected matches	46

1 Introduction

The efficient market hypothesis states that “security prices fully reflect all available information” (Fama, 1970, p.1575). Fama outlined a weak form, semi-strong form and strong form definition of the efficiency of the markets. In a weakly form efficient market future prices cannot be determined by past prices, in a semi-strong form market prices reflect all available public information, and in a strong form market prices also reflect all available private information

The betting markets provide a unique avenue to research the efficiency of financial markets. It has “no relation to any aggregate risk”, short contract periods and a definite outcome allowing mispricing to be detected (Moskowitz, 2015, p.1). Betting can however also be viewed as a consumption of a service, and an efficient market is not necessarily expected. With increased trading opportunities however, it is reasonable to expect that the market exhibits the same behaviors as a regular financial market. As there is no obvious systemic risk in betting, there should not be any risk premia, and in an efficient betting market the expected negative value of investing (betting) should be equal to the transaction costs (the bookmaker’s margin).

With the rise of online betting, competition in the betting market has increased and margins have been decreasing. A lower margin increases the possibility of making a profit by betting in an inefficient market. While most online bookmakers place restrictions on customers consistently making a profit (Koning and van Velzen, 2009), there are betting markets without such limitations, turning betting into a financial market in its own right.

Following from the efficient market hypothesis by Fama forms of market efficiency can be described for the betting markets. Kuypers (2000) asserts that in a weakly efficient betting market there should be no difference in the expected value of bets based on prices; in a semi-strong no public information should be able to serve as criteria for selecting better than average bets; and in a market with strong form efficiency no private information should enable a bettor to make more favorable bets.

A great deal of research has been conducted on the efficiency of betting markets, and there are mixed results on whether the betting market is efficient, for a review see Williams (1999, 2005). Most of this research has however been conducted on traditional bookmakers, where bets from customers consistently making a profit could be refused. This could give rise to market inefficiencies as the bookmaker

could actively be limiting the information in the market.

1.1 Motivation and research questions

This thesis aims to answer whether the unrestricted betting markets are efficient. In order to answer this question a unique data set has been collected of historical odds set by the bookmaker Pinnacle, who has a “winners welcome” policy and do not place restrictions on customers who make a profit¹.

Historical football match results and the initial and final odds posted by Pinnacle have been collected. This enables an extensive comparison of initial and final odds, and an opportunity to assess efficiency at two points in time.

Additionally, detailed market movement for the 2016 seasons has been recorded. In this data set there are a multitude of matches where odds are posted for round n before round $n - 1$ is played. This offers an opportunity to see how new information in the form of match results affects odds in the market.

The research questions posed about the unrestricted football betting market are as follows:

- Is the market weakly efficient?
- Is the market semi-strongly efficient?
- Is new information absorbed by the market?

To answer these questions, strategies attempting to earn a profit in the market will be constructed. If a strategy can be constructed earning a profit, or yielding a favorable return compared to the expected value, this will be taken as evidence against efficiency. The null hypothesis is that the market is efficient, the market cannot be said to be inefficient unless it can be explicitly demonstrated.

In order to determine if there is weak-form inefficiency well known biases from the literature will be used as a selection criterion for bets to be placed. Semi-strong efficiency will be assessed by developing a prediction model based on team statistics that will place bets if predicted probabilities are higher than what the market estimates.

To see if new information is absorbed, the change in initial and final odds will be compared to final match result. Additionally, the prediction model will create predictions for round n at two points in time. One with the information available before round n is played, and one with the information available before round $n - 1$ is played.

¹<https://www.pinnacle.com/en/promotions/winners-welcome>

If there is a change in predicted probabilities this should be reflected by a change in market odds, provided that the market is semi-strongly efficient. There could however, be a lag in information uptake giving rise to temporary inefficiencies.

The main hypotheses based on findings by Vlastakis et al. (2009) and Constantinou and Fenton (2013b) is that weak form inefficiency will be observed in the form of biased odds, but that this is unlikely to be consistently profitable due to the competitive nature of the market.

Semi-strong efficiency is expected to hold for the prediction model constructed. New information is expected to be absorbed by the market, but a delay in information uptake is expected and could potentially be profitable and give evidence for temporary semi-strong inefficiencies. Lastly final odds are expected to be more correct than initial odds, demonstrating that the market absorbs information.

1.2 Results and findings

Evidence of weak form inefficiency is found through consistent bias in odds, demonstrated by the home favorite-longshot bias where home favorites are underestimated in the market. Additionally, there is indication home wins offer more favorable odds than away wins or draws.

Semi-strong inefficiency is found by the constructed prediction model consistently making a profit when determining both that the chance of a home win has increased, and that the market odds are underestimating the chance of a home win. This demonstrates that the market does not incorporate all available new information.

The market is demonstrated to absorb information. The first odds change after new information is available is in the same direction as the change in predictions made by the prediction model, but it takes a median of 3 hours before changes are observed. Additionally, the difference in initial and final odds are found to be significantly correlated with match results for all match outcomes. All demonstrated profitable strategies are also significantly less profitable when using final odds.

1.3 Contribution to existing literature

The results from this thesis further strengthen arguments made in literature that the betting markets are not efficient. It additionally provides evidence of inefficiency in a market where the bookmaker will not place restrictions on bettors making a profit.

An effect where the odds on a subset of matches do not incorporate the most recent information available is demonstrated. Furthermore, while market inefficiencies have been established, it has been shown that the market grows more efficient towards match start.

1.4 Structure

Chapter 1 – Introduction The introduction presents the research topic, motivations and research questions and includes a summary of the main findings.

Chapter 2 – Background Describes market efficiency in the context of betting markets, how betting markets work and provides an overview of research on the efficiency of betting markets. How betting odds function is then described, and an overview of statistical modeling of football match results is provided.

Chapter 3 – Methodology Gives an overview of the data set, how the prediction model will be constructed and how the market will be assessed.

Chapter 4 – Results Presents the results from the construction of the regression model, any pricing discrepancies found and the performance of the regression model. The results from new information become available in the market is then described.

Chapter 5 – Discussion The main results are explained, discussed and put into context of relevant literature.

Chapter 6 – Conclusion Gives the main conclusion to the research questions posed, and offers recommendations for further research.

2 Background

This chapter first describes market efficiency, and then introduces sports betting and betting markets, and how they are related to economic theory. Next it is described how odds are defined and margins are calculated. Forecast models aimed at predicting match results are then described.

2.1 Market efficiency

The efficient market hypothesis can be described as “the simple statement that security prices fully reflect all available information” (Fama, 1970, p.1575). If there is information that is currently not incorporated in the price, traders will act on their knowledge and push prices until they fully reflect all available information. Taken into the context of betting markets, this means that on average one should not be better at predicting results than the market and thus one should not be able to make a profit by placing bets in the absence of risk premia.

The hypothesis of efficient markets can be divided into three main categories: weak, semi-strong and strong form efficiency (Fama, 1970). In weak form efficiency future prices cannot be determined by analyzing historical prices. Semi-strong form efficiency states that prices also incorporate all publicly available information, or will absorb any new such information rapidly. Strong form efficiency asserts that all information, both public and private, is reflected in the current prices.

Based on the arguments put forward by Kuypers (2000) this can be contextually defined for the betting market. Weak form efficiency would require all odds posted to be equally unprofitable, that there is no consistent bias in odds. Odds on an outcome can be thought to be analogous to the price of an asset, if some odds have a better expected return than others, above average returns can be made using only price information.

Semi-strong form efficiency would be upheld if it is not possible to pick better bets than average based on public information such as team statistics or injuries. If all public information is reflected in the prices, no analysis of this information should yield expected returns above average.

Strong form efficiency would imply a lack of ability to earn above average returns even on private information, such as unpublicized injuries. If all information is reflected in the prices, there should be no information that enables above average

returns.

Thaler and Ziemba (1988, p.163) puts forward a different but more concrete definition of market efficiency for betting markets. They postulate a weak and strong efficiency condition. Weak efficiency is that “no bets should have positive expected values”. Strong efficiency asserts that “all bets should have expected values equal to $(1 - t)$ times the amount of the bet”, where t is the transaction costs or bookmaker margin. Any reference to efficiency in general in this thesis will be using the definition put forward by Fama (1970).

2.1.1 Joint hypothesis problem

The joint hypothesis problem can be described as the problem that the efficiency of the market “must be tested jointly with some of equilibrium, an asset pricing model” (Fama, 1991, p.1576). This entails that when unexpected returns are found, it could be an indicator of an inefficient market, but it could also be an indicator of an insufficient modeling of the prices.

The betting market does not have any obvious systemic risk, and should therefore not have any risk premia in the prices (Moskowitz, 2015). In a market where this condition is met, returns deviating from the average should be evidence of inefficiency without being subjected to the joint hypothesis problem.

2.2 Betting markets

Sports betting is gambling where one can bet on some outcome of a match, and if successful receive a return based on the offered odds. Bets were typically placed over the counter at a bookmaker, but online sports betting has become increasingly more common, and bookmakers have followed suit with online offerings. The online gambling market has grown from around USD 25 billion in 2009 to an estimated USD 45 billion in 2016 (James Stocks & Co and KPMG, 2016), with sports betting accounting for more than 20% of the online gambling market share in 2012 (Crowdpark, 2012).

Traditionally brick and mortar bookmakers posted odds well in advance of a match and only rarely made adjustments (Forrest et al., 2005). With online bookmaking and the potential for instant updates there is however an increased flexibility in odds setting. On the customer side, information on all odds offerings from different bookmakers has become readily available.

New competitors with strictly an online presence have also appeared, some with

a focus on high volume, low margins and more frequent odds updates (Koning and van Velzen, 2009). Additionally, the emergence of prediction markets have enabled bettors to act as market makers, i.e. taking on the role as bookmaker, as they can now not only take odds but also offer odds, providing in aggregate more accurate odds than than traditional bookmakers (Franck et al., 2010).

These developments have turned online betting as a whole into a financial market in its own right. Bookmakers must harmonize their odds offerings with the odds available in the market in order to remain competitive and not open themselves up to arbitrage.

2.2.1 Bookmakers and prediction markets

Bookmakers post odds on events and offer a maximum amount that can be staked. A customer can then place bets based on these restrictions. Most online bookmakers act as traditional bookmakers and rely on casual gamblers, and will readily limit the amount that can wagered by customers consistently earning a profit (Koning and van Velzen, 2009).

Prediction markets, often referred to as betting exchanges, on the other hand work as an intermediary without taking on any risk. At prediction markets such as Betfair, customers can “back” or “lay” bets directly at an odds of their choosing, i.e. going long or short on a result. Backing a result is betting that a result will occur, as with a traditional bookmaker. Laying a result would be taking the bookmaker side, betting that a result will not occur¹. The prediction market will display offered odds and execute trades. A customer here is not restricted by limitations set by a bookmaker, but by what is being offered in the market. The prediction market generates a profit by commission on winnings, and has therefore no incentive to place limitations on customers making a profit.

A different kind of online bookmaker, referred to by Koning and van Velzen as “discount” bookmakers employ a hybrid model. They offer low margins based on high volume. In order to attract high volume customers, discount bookmakers such as Pinnacle offer extremely competitive odds and have a “winners welcome” policy². From a backer’s point of view the main difference between a prediction market and a discount bookmaker can be characterized by the fact that the discount bookmaker carries risk directly, which it tries to neutralize through volume.

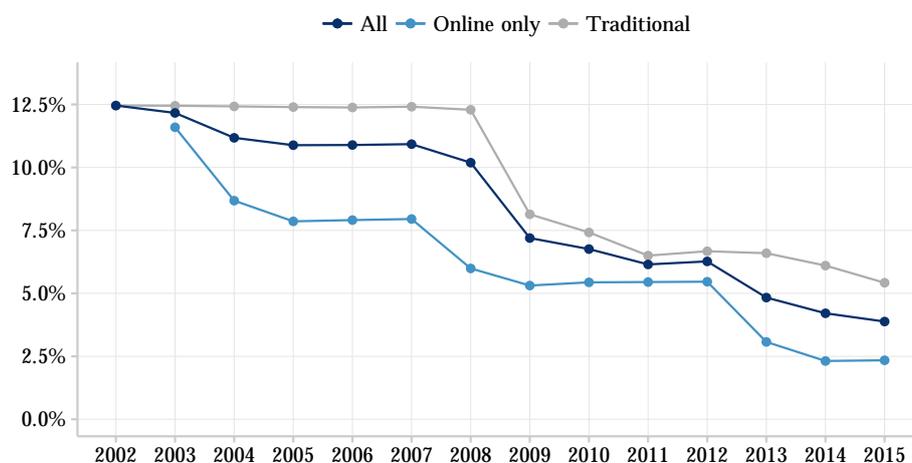
¹<https://betting.betfair.com/how-does-betfair-work.html>

²<https://www.pinnacle.com/en/promotions/winners-welcome>

2.2.2 Margin development

Using odds from football-data.co.uk it is clear that the competition has affected bookmaker margins. Exactly how these margins are calculated will be elaborated in section 2.4.2. Traditional bookmakers had on their online offerings for the English Premier League an average margin of 12.5% in 2002, which had decreased to around 5% in 2015. Their online exclusive competitors had even lower margins at about 2.5% in 2015.

Figure 2.1: Development of margins in the English Premier League for selected bookmakers



Lower bookmaker margins entails increased risk for bookmakers. High margins ensure any mispricing in bookmaker odds have to be relatively large in order to be capitalized on. Lower margins on the other hand make it easier to profit on imprecise odds setting. Forrest et al. (2005) find that in reaction to increasing market pressure bookmaker odds have become more accurate. This is however contested by Constantinou and Fenton (2013b), who find no indication of forecasting improvement in various European football leagues for the seasons 2006–2012.

2.3 Betting market efficiency

There have been a multitude of analyses on the validity of the efficient market hypothesis with regards to the betting markets. Snyder (1978) argues that in an efficient market the expected rate of return should be identical for all bets and simply equal the market maker's margin. Kuypers (2000) present a similar argument to Snyder that market efficiency entails that there should be no abnormal returns to either bookmaker or bettor. Abnormal returns are here defined as the bookmakers' margin, which would be positive for bookmakers and negative for bettors.

2.3.1 Weak form efficiency

Weak form inefficiency has been observed in betting markets both as different expectation values for different outcomes and for different levels of probability.

A review by Williams (2005, p.114) concludes that there is consistent bias in odds in the betting markets, and takes that as an indication of weak form market inefficiency. The direction of bias differs across markets however. Most sports and bets involving a bookmaker show a favorite-longshot bias, i.e. underestimation of high probability events. Parimutuel betting markets, where bets are pooled and odds are not determined until the pool is closed, showed indication of a bias in favor of longshots.

Kuypers (2000) argues that a profit maximizing bookmaker could set market inefficient odds if there are biases in bettors' expectations. Levitt (2004) finds that North American bookmakers in fact do set prices strategically based on biases held by their customers, and argues that bookmakers are in the business of both predicting game outcomes and bettor preferences. Graham and Stott (2008) find evidence in favor of this in the UK football betting markets.

Williams also notes that bias can be explained as a rational outcome depending on the utility function of bettors, and that while there exists bias that may yield abnormal returns, it is usually not profitable.

Paul and Weinbach (2010, p.128) find that "consumption plays a major role in the decision to gamble on sports", and that "pure investment-based gambling appears to be the exception, rather than the norm". Depending on the size of the investment-based participation in betting markets, and argument could be put forward that an inefficient betting market is a rational outcome.

2.3.1.1 Bias

As noted before by Williams the favorite-longshot bias is well known phenomenon in most betting markets. Cain et al. (2000) found clear confirmation of the bias in the UK football markets. Constantinou and Fenton (2013b) find additional confirmation of the favorite-longshot bias in European football markets, and are able to turn a profit by betting on favorites. Shin (1991) argues that a favorite-longshot bias is a consequence of the bookmaker protecting itself from insider trading by giving worse prices when the value of information is high.

Vlastakis et al. (2009, p.436) find that in the European football market "there is evidence that the home-field advantage is consistently overestimated", and argues that the home field advantage is mostly a conflation with the favorite-longshot bias.

Betting on home wins however generate considerably higher returns than betting on away wins in the research done by Constantinou and Fenton (2013b). Constantinou and Fenton also find that for England and Germany they would have made a profit by betting on the home team, when it was the least-likely result.

Gandar et al. (2004) find a reverse favorite-longshot bias in the American National Hockey League, where betting on the least likely team yields above average returns.

Forrest and Simmons (2008) find that odds in the Scottish and Spanish football leagues appear to be influenced relative to the number of fans. The odds for popular teams is found to be more favorable than for less popular teams. They further argue that this is to attract customers in a competitive market place.

2.3.2 Semi-strong form efficiency

There is evidence of semi-strong inefficiency, both through cross-market trading and statistical modeling generating abnormal returns.

Constantinou and Fenton (2013a) find that their statistical model based on team ratings is able to generate a profit in the English Premier league for the seasons 2008–2012. Goddard and Asimakopoulos (2004) was able to make a profit on their Poisson model based on past results as well as additional significant indicators.

Borghesi (2007, p.340) finds evidence for a weather bias in the NFL betting markets. He notes that “game day temperature significantly affects team performance and this information is not efficiently incorporated into bet prices”. This could be taken as a point against semi-strong efficiency, as this is public information that is undervalued.

2.3.2.1 Arbitrage

Arbitrage opportunities appears when it is possible to make a profit by exploiting the price difference between two markets without assuming risk. If the buying price is lower in one market than the selling price in another, one can simply buy in one market and immediately sell in the other. In the betting market an arbitrage opportunity occurs when one can bet on all possible outcomes for an event, and make a profit regardless of the outcome.

Franck et al. (2013) find a significant return on arbitrating bets between bookmakers and betting exchanges. They also find that arbitrage opportunities are frequently offered by bookmakers. They suggest that bookmakers set prices not only to optimize returns on a single bet, but taking into account the future behavior

and value of their customers. Constantinou and Fenton (2013b) also find numerous arbitrage opportunities in line with the findings by Franck et al..

Marshall (2009) finds that arbitrage opportunities are not exploited instantly, but nonetheless quite rapidly. He finds that arbitrage opportunities have a median profit of 1.5% and median duration of 15 minutes. He also notes that arbitrage opportunities that are difficult to find, for instance requiring posting at more than two bookmakers, take longer to disappear.

This behavior by bookmakers must be considered along with the reports by Koning and van Velzen (2009) that bookmakers are quick to put restrictions on customers making a profit. It also raises the question whether arbitrage can be executed consistently over time, and if there are frequent arbitrage opportunities between discount bookmakers and prediction markets where trading is unrestricted.

2.3.2.2 Price development

Brown (2012) finds evidence that informed traders make a profit in the prediction markets during tennis matches. He discovers that a subset of bettors observe match events before the wider public, forming a profitable trading strategy. This effect is possibly explained by a delay in television signal compared to watching the match live.

Gandar et al. (1998) find that the closing odds are significantly more accurate than opening odds in the NBA, and that the magnitude of change in odds is directly related to the bias exhibited in opening odds. Biased odds are moved sufficiently to remove bias by the closing of the market. Baryla et al. (2007) however find that opening odds in the NBA during the early season are more biased, and that prices are not moved sufficiently to completely eliminate bias. Croxson and Reade (2014) find that prices are updated swiftly in the live betting market upon goals being scored in football matches.

Moskowitz (2015) finds momentum effects, where bettors “chase” the change in odds, but the effects are not large enough to overcome the margin and form a profitable trading strategy.

2.3.3 Strong form efficiency

Strong form efficiency is concerned with whether private information is reflected in the market, which is difficult to measure as the potential information is unknown. An avenue that has been explored is measuring the amount of insider trading, suggesting that for the racing markets there is strong form inefficiency (Williams, 1999, p.26).

Private information is unlikely to be a significant factor for the top European football leagues, due their popularity and amount of public interest and scrutiny. Strong form efficiency is therefore not looked into any further. Cain et al. (2003, p.271) note that “bookmaker margins appear to be a good proxy for the degree of insider trading”. The margins and margin developments observed in section 2.2.2 supports the notion that the level of insider trading is low.

2.4 Odds and probabilities

2.4.1 Odds notations

Odds notations can be viewed as a stylistic representation of the probability that an event will occur. A bet will have a positive expected value if actual probabilities are higher than probabilities inferred by the odds.

Three main styles are used to quote odds in the west: European, English, and American (Cortis, 2015). Additionally, Cortis notes that there are three quoting styles mainly used in Asia: Hong Kong, Indonesian and Malayan. Any references to odds in general in this thesis will be using the European notation.

Table 2.1: Probability and odds examples

p	European	English	American
0.2	5.000	4/1	+400
0.4	2.500	5/2	+150
0.5	2.000	1/1	+100
0.6	1.667	2/3	-150
0.8	1.250	1/4	-400

The European odds notation, also called decimal odds, is simply the inverse of the probability. When staking on decimal odds, the bet multiplied by the odds is the total one receives back on winning bet. To calculate the profit one must subtract the stake. With o and p representing the odds and probability on some outcome j , this can be presented:

$$o_j = \frac{1}{p_j} \quad (\text{EU})$$

English odds, or fractional odds, are mostly used in the United Kingdom. It can be interpreted as profit divided by stake, so for a bet with odds 4/1 the potential

profit will be 4 for every 1 unit staked. For 1/4 one has the potential to win 1 unit by staking 4.

$$o_j = \frac{1 - p_j}{p_j} \quad (\text{UK})$$

American odds are sometimes referred to as moneyline odds. A positive figure denotes potential profit on a 100 unit bet, while a negative figure represents how much would have to be staked in order to win 100 units. Moneyline odds quoted +400 and -400 would be equivalent to 4/1 and 1/4 in fractional odds respectively.

$$o_j = \begin{cases} +100(\frac{1-p_j}{p_j}), & \text{if } p_j \leq 0.5 \\ -100(\frac{p_j}{1-p_j}), & \text{if } p_j > 0.5 \end{cases} \quad (\text{US})$$

2.4.2 Margins

Traditional bookmakers make their profit by ensuring the total probabilities of all outcomes end up being at least 1, i.e. $\sum_{i=0}^n p_i \geq 1$. If the sum of the probability space does not equal or exceed 1, the bookmaker can be arbitrated (Cortis, 2015). If the books are balanced the exceeding probability equals the bookmaker's margin, regardless of outcome. This can be expressed either as the margin, or conversely as the bookmaker's total payout rate to the bettors:

$$\begin{aligned} \text{margin} &= \sum_{i=0}^n p_i - 1 \\ \text{payout} &= 2 - \sum_{i=0}^n p_i \end{aligned}$$

Prediction markets however take a different approach more similar to the stock exchange in a traditional stock market. Customers offer odds on either an outcome occurring or not occurring. Offers that match are cleared and a bid/ask spread, or in Betfair terminology back/lay, remains available on the market.

The exchange makes its profits by taking a commission on winnings. One pays a commission typically between 2% and 5%³, depending on the volume one brings to the market. Actual odds can thus be adjusted to take this margin into account. With c denoting the commission, o^m denoting the offered European odds on the market, and o^e denoting actual European odds after fees, this can be expressed:

$$o^e = o^m - c(o^m - 1)$$

³<http://www.betfair.com/aboutUs/Betfair.Charges/>

Using adjusted odds, a total margin can be calculated in the same way for a prediction market as for a traditional bookmaker. As an example using the match Stoke – Manchester City, 20 August 2016, taking odds from `pinnacle.com` and `betfair.com`, one can see that the best odds for a given outcome vary with fee structure.

Table 2.2: Odds example comparing bookmakers and prediction markets

	Pinnacle		Betfair	
Commission	–	–	2%	5%
Margin	2.03%	0.07%	1.23%	3.03%
Stoke	5.340	5.900	5.802	5.655
Draw	3.980	3.950	3.891	3.803
Man. City	1.719	1.730	1.715	1.694

2.4.3 Adjusted probabilities

Since the probability space in odds offered by bookmakers exceeds 1, odds must be adjusted if one wants to interpret what the bookmaker believes to be the true probabilities. Franck et al. (2010) assumes the probability overround is distributed equally among all outcomes. With o_j representing the European odds offered for outcome j on an event, adjusted probability \bar{p}_j can then be expressed:

$$\bar{p}_j = \frac{1}{o_j \sum_{j=0}^m \frac{1}{o_j}}$$

2.4.4 Types of bets

One can bet on a multitude of outcomes in a football match, for instance the correct score, over or under 2.5 goals, number of corners, number of bookings or first team to score. When it comes down to final match result there are two main types of bets, match odds or Asian Handicap.

When betting on match odds, there are odds set for each of the three potential outcomes home, draw and away. The exception is in the prediction markets where one can also offer odds on an outcome, also known as a lay bet. As an example laying the draw at odds 4 for 1 unit, would incur a potential liability of 4 units for a potential gain of 1 unit if there is no draw. One could also look at it as staking 4

units at odds $(4 + 1)/4 = 1.25$, or interpret it as the UK odds 1/4, on not observing the draw.

In Asian Handicap one of the teams is given a handicap or a head start by a number of goals, typically to balance unequal teams. In the case of a draw, stakes are returned and the bet is void. If a team is given a head start or handicap by half a goal a draw is always avoided. Handicaps given in quarters are split in two bets, with one stake rounded up to nearest half goal handicap and the other rounded down to the nearest half goal handicap. E.g. a 2 unit bet at a handicap of -0.75 will be split into a bet with 1 unit at handicap -1, and 1 unit at handicap -0.5⁴.

In relation to full time results and scores Vlastakis et al. (2008, p.111) find that “the size of the Asian Handicap appears to be a significant predictor of both home and away scores”.

2.4.5 Staking strategies

A staking strategy is a set of rules for deciding how much money to place on a bet. The most straight forward staking strategy is simply a one unit stake on each bet placed. If the bettor is on average better at calculating odds than the market on the bets they place, this strategy should yield a profit, but will not necessarily maximize the profit made.

The Kelly criterion (Kelly, 1956) is a formula that determines the profit maximizing stake on a bet as a proportion of the current bankroll. In the follow equation d represents the fraction of the current bankroll one should bet, o the odds and p the actual probability of success. A negative fraction would mean that one should lay the bet. The Kelly staking system requires that one has the true probabilities or a very good proxy available.

$$d = \frac{po - (1 - p)}{o}$$

Hvattum and Arntzen (2010) also introduces a one unit win strategy, which entails staking such that the return of a successful bet equals one unit. The amount to be staked would then equal $(o - 1)^{-1}$. They find the method to outperform the one unit stake strategy, but as this strategy bets greater amounts on the most likely outcomes they interpret this as an indication of favorite-longshot bias.

⁴<https://www.pinnxacle.com/en/betting-articles/betting-strategy/betting-on-soccer-asian-handicap-markets>

2.5 Prediction models

There have been many attempts at modeling the result probabilities in football matches. From a statistical point of view two main methodologies have been applied: a goal scoring process following a Poisson distribution to predict goals scored, and logistic regression directly on the probability of a result Goddard (2005). Additionally, various machine learning models have been looked into for instance by Vlastakis et al. (2008) and Joseph et al. (2006), but are considered beyond the scope of this thesis.

A Poisson model attempts to model the number of goals scored by each team accounting for the attack strength of the team, and the defense strength of the opposing team Graham and Stott (2008). A match result can then be inferred by the probability distribution of goals scored. This method has been used in models developed by for instance by Dixon and Coles (1997) and Rue and Salvesen (2000).

An ordered probit or logit model attempts to model the observed result directly without a goal scoring process. The outcome is ranked as ordinal data for a home win, draw or away win and predicted by a proportional regression model. The obvious drawback of this methodology is that there is no prediction of goals. This method has been employed in research by for example Kuypers (2000), Goddard and Asimakopoulos (2004), Forrest et al. (2005) and Hvattum and Arntzen (2010).

Goddard (2005) evaluates the differences between bivariate Poisson regression and ordered probit regression. The article concludes that ordered probit models using goals based performance indicators might perform slightly better than Poisson models, but that ultimately there does not seem to be a large difference in forecasting accuracy between the models.

2.5.1 Relevant factors

In order to successfully predict match results, good indicators of team performance and ability need to be found. As mentioned Goddard (2005) used goals based performance factors, that is the number of goals scored or conceded by each team for up to nine of the previous home or away matches. Using goals scored and conceded as independent variables was found by Goddard to outperform using just match results.

Additionally, Goddard includes other factors not related to past team performance. Significance is found for a variable indicating whether the match is important for the team, that is if the result could influence relegation, promotion,

or qualification for European championship participation. Geographical distance, cup participation and size of team following was also demonstrated to be significant predictors. These factors will be expanded upon in section 3.2.3.

2.5.2 Rating systems

As an alternative to creating models based on lagged goals and results as performance indicators, rating systems as a proxy for team strength has been developed. Elo ratings were originally developed by Elo (1978) for ranking chess players. It assigns an initial numerical score to every player. After a match, points are added to the score of the winner and subtracted from the score of the loser. The amount of points is based on the pre-match rating of the players. A highly rated player losing to a low rated player will incur a larger shift in scores than if the opposite result was observed. A draw would mean that the lower rated player increases their score, while the higher rated player's score is decreased.

In football, rating systems have been applied on teams rather than players. Hvattum and Arntzen (2010) developed a rating model based on Elo ratings, but modified by taking goal difference into account. As such a 3-1 win will be rated more highly than a 2-1 win. Comparing the rating system against a goal based model developed by Goddard (2005), they find that Elo rating might be more efficient when based on short time spans. Goddard finds that the optimal estimation period for his model is 15 seasons.

Constantinou and Fenton (2013a) developed a rating system similar to Elo ratings called pi ratings. The system builds on the model by Hvattum and Arntzen by including goal difference in the ratings. Additionally, in order to account for a difference in performance based on field, a separate score is used for home and away ratings. This method is found by Constantinou and Fenton to outperform Elo ratings, as well as actually yielding a profit.

2.5.3 Evaluation criteria

Evaluation of prediction models are two-fold. A statistical evaluation can be conducted based on forecasting accuracy, and an economic evaluation can be done through observing the returns of the model against market odds.

Rue and Salvesen (2000) uses a pseudo-likelihood measurement constructed as the geometric mean of the predicted probabilities of the actual results. More commonly used is the however the Brier score (Brier, 1950) or quadratic loss function. It was originally developed to assess the accuracy of weather forecast, and is a

commonly used criteria in the literature. The Brier score has been used to evaluate model accuracy by for instance Forrest et al. (2005), Franck et al. (2010) and Hvattum and Arntzen (2010).

With $\hat{p}_{j,i}$ being the estimated probability for an outcome j for event i , and $y_{i,j}$ taking on the value 1 if the outcome was observed and 0 otherwise, the Brier score b for n events with m outcomes can be defined as:

$$b = \frac{1}{n} \sum_{i=0}^n \sum_{j=0}^m (\hat{p}_{i,j} - y_{i,j})^2$$

The economic criteria assess whether the model can make abnormal returns. By simulating bets placed on past matches the economic viability of the model is determined. Different staking strategies as discussed in 2.4.5 can be employed, usually the one unit staking strategy is included. For instance, Hvattum and Arntzen (2010) evaluate the model against a one unit bet, one unit win and Kelly criterion staking strategies. Constantinou and Fenton (2013a) on the other hand apply only the one unit bet strategy for their economic evaluation.

3 Research methodology

The underlying approach of this thesis is based on constructing a prediction model on historical data. The model is used to estimate a reasonable probability for the outcomes of a match before and after new relevant information becomes available. The change in forecasted probabilities will be evaluated against price changes observed in the market. Additionally, odds will be grouped in accordance to previously demonstrated biases in order to assess their prevalence.

This chapter will first describe the data set, then the construction of the regression model will be outlined. Finally, the market assessment is described.

3.1 Data

The basic data set consists of matches for the seasons 2004 through 2016 in the top five European leagues: English Premier League, French Ligue I, German Bundesliga, Italian Serie A, and Spanish Primera Division. Available for each match is date, home and away team, and number of home and away goals scored.

Additionally, geographic coordinates, and average home attendance for each team and season have been collected. Using the geographic coordinates, weather conditions have been retrieved for all matches. Attendance figures have also been obtained from second tier division where available in order to account for missing data on promoted teams.

All countries have their respective top tier cup tournament: English FA Cup, French Coupe de France, German DFB-Pokal, Italian Coppa Italia, and Spanish Copa del Rey. Results from these tournaments have been gathered for all teams and seasons.

Pinnacle odds are available for matches taking place in the seasons 2012 through 2016. Initial odds (opening odds) and odds at the time of match start (closing odds) have been gathered. Detailed pre-match odds development is available for the 2016 season only.

3.1.1 Data sources

This unique data set has been collected programmatically using a variety of sources. Any data not readily available in an export format or through an Application Pro-

gramming Interface (API) have been retrieved using a custom parser. These raw data have further been aggregated to form a consistent data set.

- Historical match results have been downloaded from `football-data.co.uk`.
- Geographic coordinates have been collected based on stadium locations from `wikixedia.org`.
- Weather data have been collected from the API provided by `darksky.net` using the retrieved stadium coordinates.
- Attendance figures have been collected from `welfussball.de`.
- Cup matches have been collected from `nifs.no`.
- Odds posted by Pinnacle for the seasons 2012–2015 have been collected from `oddsportal.com`.
- Detailed market data for the 2016 season have been collected by recording data retrieved from the `pinnacle.com` API every 10 seconds throughout the season.

The data set provided by `football-data.co.uk` includes bookmaker odds and has been widely used in previous analysis, for instance by Forrest and Simmons (2008), Franck et al. (2010), and Constantinou and Fenton (2013b). The odds used in this thesis have been collected separately for three reasons: the aforementioned data set does not consistently provide Pinnacle odds for all seasons, odds are collected days in advance, and it does not have odds from multiple points in time. Archived odds information provided by `oddsportal.com` includes additional matches as well as both first and last odds posted. This ensures all final odds gathered are from the time of kick-off (match start), when information is thought to be as complete as possible. It also enables a comparison of initial and final odds offerings.

3.2 Prediction model

An ordered probit regression model is constructed using the pi-ratings described by Constantinou and Fenton (2013a). Potential additional regressors from the data set will also be evaluated.

The seasons 2004–2006 are used to estimate initial pi-ratings. Seasons 2007–2011 are used to calibrate the pi-rating parameters and additional regressors for the probit model, while seasons 2012–2016 are be used to evaluate the performance of the model. The model is refitted after each season, in an attempt to improve forecasting capabilities for subsequent seasons.

Two models are constructed to predict the outcome of a match n . One is using information available at kick-off of match n , and one uses the information available before the start of match $n - 1$. The models are initially explored using all available regressors, before using a stepwise selection process to select the best model.

3.2.1 Regression model

The regression model will be constructed as an ordered probit model in R using the `ordinal` package (Christensen, 2015). The methodology is based on the model developed by Hvattum and Arntzen (2010), but pi-ratings are used over Elo-ratings. Further the model differs by evaluating additional regressors besides rating difference in attempt to improve upon the model. A rating system is chosen over a lag in results such as the model constructed by Goddard (2005), due to the quicker learning rate on a limited history demonstrated by Hvattum and Arntzen.

Match result is the dependent variable and takes on the value 1 for a home win, 0.5 for a draw or 0 for an away win. As the model is concerned with relative strength disparity, the independent team specific variables will be transformed to represent the difference between the two teams. With x_j^H and x_j^A representing the values of the home and away team respectively for some feature (e.g. team rating), the independent variable will be defined as $x_j = x_j^H - x_j^A$.

3.2.1.1 Ordered probit

Ordered regression is type of regression model where the response variable is categorical. The most common type of models are probit and logit. The key difference between a probit and a logit regression is the assumption of the distribution of the error term and the link function.

In this study a probit model is chosen as its specification is considered more intuitive by the author. Its usage is more common in econometrics and have been used to estimate football results by for example Goddard and Asimakopoulos (2004), Forrest et al. (2005) and Franck et al. (2010). Hvattum and Arntzen (2010) on the other hand used a logit model. Verbeek (2004, p.191) states that both models “typically yield very similar results in empirical work”. The choice of one model over the other is not believed to influence the findings.

The most common probit model is a binary regression which provides as the name indicates an estimator of only two outcomes. Football matches have three potential outcomes and a model providing multiple estimators is necessary. In order to account for more than two outcomes, a multinomial regression model can be used.

An ordered probit model is chosen over a multinomial one due to the results being ordinal. A home victory is further away from a loss than a draw, e.g. a result of 1-0 is closer to 1-1 than 1-2. The information contained in the order of these outcomes would not be accounted for when using a multinomial model.

A probit regression model assumes the error term, or random noise, follows a normal distribution. The model further assumes that for an observed discrete outcome y there exists a latent continuous variable y^* . The value of y is determined by whether y^* crosses a threshold value γ (Verbeek, 2004, p.203-204).

$$y_i^* = \sum_{k=0}^K x_{ki}\beta_k + \epsilon_i = \mathbf{X}_i^T \boldsymbol{\beta} + \epsilon_i$$

$$y_i = j \quad \text{if } \gamma_{j-1} \leq y_i^* \leq \gamma_j$$

The probability that an outcome occurs is then specified using the normal distribution. More specifically the probability that outcome j of a total m of outcomes will be observed for observation i can be written as:

$$p_{i,j} = \begin{cases} \Phi(\gamma_0 - \mathbf{X}_i^T \boldsymbol{\beta}), & \text{if } j = 0 \\ 1 - \Phi(\gamma_{j-1} - \mathbf{X}_i^T \boldsymbol{\beta}), & \text{if } j = m \\ \Phi(\gamma_j - \mathbf{X}_i^T \boldsymbol{\beta}) - \Phi(\gamma_{j-1} - \mathbf{X}_i^T \boldsymbol{\beta}), & \text{otherwise} \end{cases}$$

3.2.1.2 Maximum likelihood estimation

The estimation of the coefficients in an ordered probit model is done by maximum likelihood estimation. Maximum likelihood estimation begins with the assumption that we know, or can reasonably infer, the true distribution of the data. One then calculate the likelihood of observing the sample for a set of given parameters. The resulting value of parameters are the ones that provide the highest likelihood of observing the sample data, given the assumed distribution (Verbeek, 2004, p.161-164).

If $y_{i,j} = 1$ when outcome j is observed for event i and 0 otherwise, the log likelihood function to be find the optimal γ_j and β_j can be given as:

$$\ln \mathcal{L} = \sum_{i=0}^n \sum_{j=0}^m y_{i,j} \ln(p_{i,j})$$

3.2.1.3 Proportional odds assumption

An ordered probit model estimates the same coefficients for all outcomes, and assumes the that “the relationship between each pair of outcome groups is the same” (IDRE, 2016).

The `nominal_test` function from the R package `ordinal` by Christensen (2015) will be used to test this assumption. It performs a likelihood ratio test, where the null hypothesis is that there is no difference between the coefficients in the model for different outcomes. A significant result for any coefficient will therefore indicate a violation of the proportional odds assumption.

3.2.1.4 Stepwise Regression

The stepwise model selection tries to optimize the goodness of fit of a model by adding or subtracting variables automatically in a stepwise fashion. Goodness of fit is measured by the Akaike information criterion (AIC) and is performed using the `step` function included in R's standard library. AIC does not evaluate the absolute quality of the model, it is strictly a measure of the goodness of fit vs complexity for comparable models on the same data set.

3.2.2 Pi-rating

The pi-rating developed by Constantinou and Fenton (2013a) will be used as an indicator of team performance. The pi-rating system assigns an initial rating of 0 to every team. After each match ratings are updated with a new value derived from the previous ratings and the match result. If a team gains n rating points after a match, the opposing team will see a change in $-n$ rating points such that the average rating across all teams in a league will always be 0.

Pi-rating resembles the Elo-rating system used by Hvattum and Arntzen (2010). It however uses separate ratings for home and away matches in order to account for different performance at home and away matches, and was demonstrated by Constantinou and Fenton to outperform Elo-ratings.

3.2.2.1 Definition

The updated rating for a home(away) team is based on their current home(away) rating, the opponents current away(home) rating and the outcome of the match. In order to calculate new ratings, the pi system also needs two parameters specified:

1. A learning rate λ to specify to what extent new information overrides old information
2. A catch-up learning rate γ , which influences how much home results affects away ratings and vice-versa

$R_{\pi G}$ represents the rating for team π at ground $G = \{H, A\}$. An expected goal difference \hat{g}_{DG} is calculated for each side, and the total expected goal difference \hat{g}_D is the difference between the two.

$$\hat{g}_{DG} = \begin{cases} +10^{\frac{|R_{\pi G}|}{3}} - 1, & \text{if } R_{\pi G} \geq 0 \\ -10^{\frac{|R_{\pi G}|}{3}} - 1, & \text{if } R_{\pi G} < 0 \end{cases}$$

$$\hat{g}_D = \hat{g}_{DH} - \hat{g}_{DA}$$

The observed error e is defined as the absolute difference in actual and expected goal difference. It is weighted by a function ψ and the sign is either negative or positive for each team depending on whether the goal difference was overestimated or underestimated.

$$e = |g_D - \hat{g}_D|$$

$$\psi(e) = 3 * \log_{10}(1 + e)$$

$$\psi_H(e) = \begin{cases} +\psi(e), & \text{if } \hat{g}_D < g_D \\ -\psi(e), & \text{if } \hat{g}_D \geq g_D \end{cases}$$

$$\psi_A(e) = -\psi_H(e)$$

Applying the learning rates, new ratings can be calculated for the teams updating both home and away ratings. With \hat{R} defining the new rating, and α and β denoting the home and away team respectively:

$$\hat{R}_{\alpha H} = R_{\alpha H} + \lambda(\psi_H(e))$$

$$\hat{R}_{\alpha A} = R_{\alpha A} + \gamma(\hat{R}_{\alpha H} - R_{\alpha H})$$

$$\hat{R}_{\beta A} = R_{\beta A} + \lambda(\psi_A(e))$$

$$\hat{R}_{\beta H} = R_{\beta H} + \gamma(\hat{R}_{\beta A} - R_{\beta A})$$

3.2.2.2 Calibration

Match results from the seasons 2004–2006 are used to estimate initial pi-ratings for use in calibration and the regression model. Constantinou and Fenton (2013a, p.45) note that “two seasons of relevant historical outcomes might be enough for it to converge”, but also that “a further season of historical match outcomes might be required for teams with the uppermost difference from the average team”. To ensure

ratings are properly stabilized for the regression model, three seasons are used to initialize the pi ratings.

In order to find the optimal learning rates, ratings for the seasons 2004–2011 are calculated for different learning rate values. An optimal value for λ and γ is found by selecting the value pair providing the smallest mean squared error for the seasons 2007–2011 defined as:

$$MSE = \frac{1}{n} \sum_{i=0}^n (g_{iD} - \hat{g}_{iD})^2$$

3.2.2.3 Season limitations

After each season there will be up to three teams relegated and promoted depending on league and year. Promoted teams will not have an established rating. In order to provide an initial rating to these teams the average rating of the relegated teams is assigned equally to the promoted teams. Newly promoted teams are expected to be relatively weak, but this method is naturally a simple approximation and not an accurate reflection of team strength.

Additionally, at the beginning of each new season the team ratings will not account for results and developments that have occurred during the off-season. There can have been injuries, roster, administrative or managerial changes that will affect team effort and performance. These factors may be larger than any inter-match factors during the season, but are not addressed.

After each season the regression model will be refitted with the new data, in an attempt to improve the model for the following season.

3.2.3 Additional regressors

Goddard and Asimakopulos (2004) constructs an ordered probit model based on earlier work by Dobson and Goddard (2001) including both result factors as well as additional regressors. Most of the included potential additional regressors is based on this work.

Geographical distance

Clarke and Norman (1995) demonstrated that geographical distance between two teams has an effect on match results in English football. Longer travel distances increase the likelihood of a home win. This effect also yields significant effects on the model constructed by Goddard and Asimakopulos. Possible explanations could

be that long distance journeys cause fatigue, or that local support to some degree cancel the home field advantage.

The distance factor is included in the model as the natural logarithm of the geographical distance between the grounds of the two opposing teams. Additionally, a dummy variable is constructed where the distance is zero, i.e. two teams share the same stadium, which is true for multiple teams in Italy.

Big team effect

Goddard and Asimakopoulos also demonstrate a “big team” effect, where teams with a large following are more likely to win after controlling for league position. The advantages could be explained by the psychological pressure from a larger following, or that popular teams are in a better position from a financial standpoint.

This effect is included by taking the residuals from a regression on the natural logarithm of average attendance on league position from the previous season. Additionally, Goddard and Asimakopoulos uses the change in this variable between seasons to reduce the effect of temporary variations.

Cup effect

An elimination from the cup could indicate a loss of morale, or in itself be an indicator that the team is weaker than it appears. The effect could also go the other way, leaving the team to focus on the league. Goddard and Asimakopoulos (2004) have found that the negative effect dominates in the English Premier League for teams that have been knocked out of the FA cup.

This variable is constructed for the top cup of all the respective countries. Data on participation in the European cups is not available.

Weather

Weather has been found to have an effect on team performance in American football (Borghesi, 2007), and could potentially have an impact on match results in soccer. The exact direction of the effect is uncertain. Bad weather could impact the field and as such favor weaker teams who do not rely as much on fine play. It could also reduce any home field advantage with less attendance and less enthusiastic spectators. Overall it is likely that poor weather conditions would favor the away team or the draw, reducing the home ground advantage. It can likewise also be argued that nice weather conditions would increase the home advantage.

The weather variable is not very fine grained. It classifies weather as either clear-day, partly-cloudy-day, cloudy, fog, rain, wind or snow. A vast majority of the games are either classified as clear-day or partly-cloudy-day. Weather is therefore transformed to a single dummy variable for whether conditions are considered not good, i.e. either cloudy, fog, rain, wind or snow.

Promotion

Newly promoted team could be expected to be weaker than their counterparts that have recent experience in the top league. They could however also be more confident after being promoted from the lower tier, or the promotion could itself be an indicator of a large amount of resources being put into the club.

It is included as a dummy variable for all countries. Which effect that dominates, if any, is an open question.

Country

The models developed by Hvattum and Arntzen (2010) and Constantinou and Fenton (2013a) focus exclusively on the English Premier League. This model attempts to assess performance across multiple leagues, which could differ in their characteristics. Country is therefore included as a dummy variable. Additionally, there could be interaction effects, but in order to not overcomplicate the model this is not investigated further.

3.3 Market assessment

There have been observed a number of biases occurring in the betting markets, such as the well known favorite-longshot bias. Odds will be evaluated against biases in pricing and outcome in order to see to what degree prices are weakly efficient. All odds will be initially assessed on the basis of opening odds, any change in bias will be examined in the next section.

Furthermore, the profit making ability of the prediction model will be assessed. Additionally, biases may serve as a selection criteria for a prediction model, which may improve performance by selecting optimal matches from a pre-screened selection.

3.3.1 Evaluation criteria

Market strategies and the efficiency of the market will be evaluated on the basis of profitability using a simple one unit staking system. That is, for every match included in a strategy, a bet of one unit is placed. With $y_{i,j}$ taking on the value 1 if an outcome j is observed for match i , return after betting on outcome j is then defined as:

$$\pi_{i,j} = o_{i,j}y_{i,j} - 1$$

Alternative staking models described in 2.4.5 have been considered. The single one unit bet strategy is deemed sufficient to establish whether a strategy is profitable, and is the only staking strategy employed in this thesis.

The Kelly criterion (Kelly, 1956) optimizes expected utility from a bet, given that calculated probabilities are correct. The criterion is however difficult to apply when evaluating market biases where there is no direct modeling of the “true” probabilities.

In addition to the one unit bet staking system Hvattum and Arntzen (2010) uses a one unit win staking strategy, that is betting such that a win nets a profit of one unit. Bet stakes are then modified depending on odds, which complicates comparison across odds intervals.

Statistical forecast assessment is considered beyond the scope of this thesis, as the goal is to build on previously demonstrated models and evaluate based on market effects.

3.3.2 Odds representation

Due to their non-linear nature, odds are better handled as their inverse or probability. Means, medians and predictions will be calculated on the basis of the probability. When presented as odds for context, this will be the inverse of the probability.

3.3.3 Market bias

Market odds will be assessed on basis of outcome and prices. As discussed previously the favorite-longshot bias is a well studied and established phenomenon. Vlastakis et al. (2009) finds that the home advantage might be over-rated when the favorite-longshot bias is taken into account. Constantinou and Fenton (2013b) however finds that betting on home wins yield better results. Furthermore, they find evidence that betting on a home win when it is the least likely outcome also yields profitable

results, but only in certain leagues. In general, they find that betting on the most likely outcome gives the best returns.

Market odds will therefore be grouped according to outcome probability and league, in case different biases dominate in different leagues. Odds will be assessed in the probability intervals of 0.25. Longshots are defined as results with a market probability of less than 25%, while favorites are defined as results with a market probability of more than 75%. A reasonable hypothesis in line with the observation by Constantinou and Fenton is that high and low probability home wins are priced more favorably, as well as high probability away wins.

3.3.4 Prediction model profitability

The prediction model will be evaluated on its ability to make a profit, potentially combined with any market bias effects found. Hvattum and Arntzen (2010) find that their Elo rating based model performs better at predicting home wins, even though it is still unprofitable. Constantinou and Fenton (2013a) find that their pi-rating model is able to generate profits via longshot bets.

Forecasted values by the prediction model for event y_i is given as a probabilities $\hat{p}_{i,j}$ for the discrete outcomes $y_{i,j}$. This could be classified into a single forecast, and bets placed on what is predicted as most likely. A more optimal strategy would be placing bets on outcomes where the odds are higher than predicted, i.e. $\hat{p}_{i,j} > p_{i,j}$, where $p_{i,j}$ is the inverse of the offered bookmaker odds.

A more risk-averse approach would be including allowance for a standard error in the predictions. If a 95% confidence interval on forecasted probabilities of outcome $y_{i,j}$ is given by the upper and lower bound $\hat{p}_{i,j}^{min}$ and $\hat{p}_{i,j}^{max}$, only matches where $\hat{p}_{i,j}^{min} > p_{i,j}$ would be selected.

For every match a bet is made on the outcome for which the difference in forecasted probabilities and offered odds is the greatest. For example, if forecasted values are $\{0.6, 0.3, 0.1\}$ and offered probabilities are $\{0.56, 0.28, 0.20\}$, a home bet will be placed as the discrepancy is larger than for the draw.

Additionally, home, draw and away predictions will be evaluated separately, with a bet being place if forecasted probabilities are higher than offered propabilities. This evaluation might place multiple bets for different outcomes on the same match.

A reasonable expectation based on the findings by Constantinou and Fenton (2013a, 2013b) and Hvattum and Arntzen (2010) would be that model will be most successful at picking home wins, and especially home longshots. Additionally it should be noted that the simple classifier will place a bet on almost every match.

The stricter selection criteria should produce a better average return, but this will come at the expense of a lower volume.

3.3.5 Market development

Parts of this analysis will be based on the detailed market data which is only available for the 2016 season. Price formation will be explored to see whether there are more opportune moments to place bets and if there is predictive power in the price movements.

If information is absorbed by the market, closing odds should be more accurate than opening odds. This entails that opening odds should on average be further from true odds, and be preferable for a statistics based betting model. On the other hand since private information is higher at the time opening odds are posted, this could mean that margins are higher to protect the market maker.

Average returns in opening and closing odds price will be assessed to see which factor potentially dominates. Any found biases and the prediction model will be also be measured against both opening and closing odds.

3.3.6 Information uptake

In the data set there are posted odds for matches where at least one of the opposing teams have a match to be played in the interim. To illustrate with an example, the match Newcastle – Stoke was played 31 October, and odds on that match were first posted 19 October. However, Stoke is playing Watford on the 24th, and Newcastle is playing Sunderland on the 25th. This gives us a window where new objective information is provided through match results on already published market odds.

The prediction models will forecast probabilities of the respective outcomes based on the information available before the first and after the last interim match is played. In our example forecasts will be generated before the Stoke - Watford, and after the Newcastle – Sunderland match has taken place.

In this example there are actually two points of information reveal, but this methodology only measures the changes after the last one. The reasoning behind this is that there is limited price movement and trading on matches long before kick-off, and that most of these events will be taking place on the same day, narrowing the time window available for any price movement to occur.

Time of the final whistle is not available, but is defined as kick-off time plus 2 hours, which should give ample room for stoppage time. Price movements could occur within this window. Match results could to some degree be determined before

full time if one team has a considerable lead, or odds movements could be observed just after the match has ended. It is assumed that these effects will have an impact on odds observed after the match has ended as defined by kick-off plus 2 hours. An alternative would be observing first or any odds movements after kick-off of the interim match, but these changes could be influenced by noise unrelated to the interim match and would considerably complicate the analysis.

Odds before and after the new information will be evaluated. Potentially there will be a delay in information uptake, or new information might not be sufficiently accounted for, which could be exploited. The hypothesis is that odds will move in the same direction as the prediction, but that there might be a time lag. Pre-match odds and closing odds will also be compared to determine if the information signal can be detected in the final odds posted.

4 Results and findings

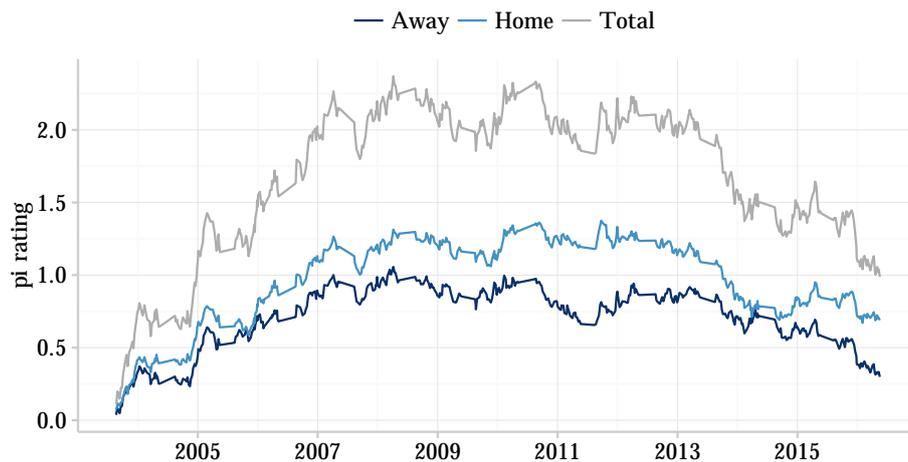
This chapter will first present the results from the construction of the regression model. Next the results from the market assessment will be presented, detailing any bias found in pricing and the profitability of the constructed prediction model. Third market development will be assessed, including development in bias and market reaction to an information signal.

4.1 Prediction model

4.1.1 Pi-rating

As described in 3.2.2.2 the learning rate λ and the catch-up learning rate γ need to be calibrated for the pi-rating model. This is estimated by iterating through multiple values for the two parameters based on the calibration performed by Constantinou and Fenton (2013a). The two parameters are essentially off by an order of magnitude, i.e. the catch-up rate is about 10 times larger than the learning rate. λ values are tested in the range 0.025 – 0.075 with a step of 0.005. γ values are evaluated in the range 0.25 – 0.75 with a step of 0.05

Figure 4.1: Development of pi-rating for Manchester United



After iterating through different value pairs of λ and γ , the optimal values for this data set based on seasons 2007–2011 are found to be $\lambda = 0.035, \gamma = 0.5$. (A.1). A slight adjustment from the original values $\lambda = 0.035, \gamma = 0.7$ found by Constantinou and Fenton. The lower catch-up rate could be explained by the initial calibrations

which have somewhat stabilized a rating baseline, but it could also be an artifact of optimizing values across different leagues.

The maximum home and away pi ratings observed are 1.72 and 1.46, the minimum -0.92 and -1.19, while the mean is 0.18 and -0.18. This is a pretty clear indication of the difference in home and away performance, even among the bigger teams.

For an illustration in the development of pi-rating one can look at the pi-ratings for Manchester United in figure 4.1. Ratings look to be somewhat stabilized by the 2007 season. One can also see how home and away rating to a large degree mirror each other, but not perfectly. Manchester United has apparently underperformed at home in beginning of the 2014 season, to the extent that the home rating was almost overtaken by the away rating.

4.1.2 Regression model

The model is estimated using stepwise regression. The preliminary specification takes on all potential regressors, and eliminates and adds regressors until yielding the model with the lowest AIC score. Two models are estimated, one using the most recent pi rating difference, and one using the lagged difference.

Table 4.1 gives a summary of evaluated regressors, and whether they have been included in a final model specification.

Table 4.1: Evaluated independent variables

Variable	Description	Model
pi_diff	Difference in pre-match pi-rating between the teams	1
pi_diff_1	Lagged difference in pre-match rating	2
ap_diff	Difference in “big team” effect	1, 2
dap_diff	Difference in change in “big team” effect	–
cup_diff	Difference in cup elimination	–
promoted_diff	Whether one of the teams have been promoted	–
weather_bad	If the weather is considered not good	–
distance	Natural logarithm of the geographical distance	–
derby	Whether two teams share the same stadium	–
country	One dummy variable for each country	–

4.1.2.1 Model 1 – Most recent rating difference

This model uses the most recent pi rating difference. A complete specification and regression estimations can be seen in the appendix (A.2). The model with the lowest AIC score after stepwise iteration includes the rating difference, the big team effect difference, geographical distance, derby and country.

After running a likelihood-ratio test, it becomes clear that both the country and derby variables in this specification violate the proportional odds assumption. Potential alternatives to handle this problem could be to construct a model per country, use a different model that does not assume proportionality, or ignore it and look into the effects of the violation. The decided upon solution was to drop the offending variables in order to keep the model specification uncomplicated.

Dropping these variables from the initial model and re-running a stepwise selection renders a model which includes only the rating difference and big team effect. Both coefficients are as expected positive, i.e. if the rating or big team effect is larger for the home team it favors the home team and vice versa. No proportionality issues are apparent after this adjustment.

Figure 4.2: Cumulative probability distribution for pi_diff in model 1, season 2011

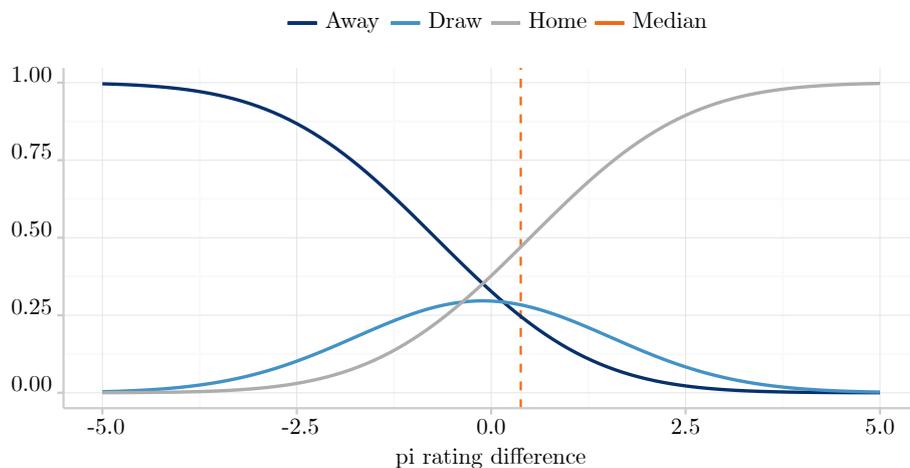


Figure 4.2 illustrates the impact of pi-ratings on the predicted result. The median difference in home and away ratings is 0.18, which can be interpreted as the home field advantage.

4.1.2.2 Model 2 - Lagged rating difference

This model is used to estimate the probability of a given result before new information appears. Complete specification is in the appendix (A.3). All factors except the rating difference and cup elimination is known ahead of time. Cup elimination

would only happen once per season, and the effect would likely not be large. The implication of this additional information is therefore not taken into account.

The initial specification excludes country and derby due to the violations of the proportional odds assumption when specifying model 1. After running a stepwise selection procedure one can observe the same result as for the first model, only the rating and big team effects are accounted for in the final model. There is not found any proportional odds violations in this final model.

4.1.2.3 Model evolution and differences

The models are refitted after every season in an effort to incrementally improve forecasting ability. An interesting result is an increasing magnitude in the pi rating coefficient, with a decreasing big team effect. This suggests that the pi ratings are improving slightly with additional seasons. A table detailing coefficients for the various seasons can be found in the appendix (A.4).

The difference in magnitude between the two models is small, and model 1 could probably have been used as a forecasting model with lagged pi ratings. This indicates that the pi rating algorithm is quite robust.

Both models have nearly identical mean prediction values. A paired t-test concludes that the two models give different prediction values, but the magnitude is low. This is mostly an indicator that the pi-ratings remain relatively stable. Additionally, for the most part teams play at home and away every other match, the average impact of a change in pi-rating is therefore softened by the catch-up learning factor.

4.2 Market assessment

4.2.1 Market bias

Home field advantage

When looking at all countries under one, there seems to be a clear prevalence of home bias. Betting consistently at opening odds on a home win on all matches yield an average return of -1.3%, while the equivalent numbers for a draw or away win are -3.6% and -3.4%. The overall results are mostly consistent across seasons as well. Betting on home matches would turn a profit in 2 out of 6 seasons, while draw and away bets are consistently negative.

If one looks directly at the different leagues, a more nuanced picture is painted.

The English Premier League actually shows a profitable away bias, while there is a profitable home bias in Spain. The Spanish home bias is most consistent and has been observed in 5 out of 6 seasons.

In summary there's a clear overall bias against draws and aways, and all else being equal betting on home wins will yield considerably better results.

Table 4.2: Mean profit from a bet on matches grouped by outcome

	ENG	ESP	FRA	GER	ITA	Sum
Home	-0.035	0.028	-0.021	-0.009	-0.033	-0.013
Draw	-0.029	-0.071	-0.035	-0.046	-0.004	-0.036
Away	0.033	-0.062	-0.056	-0.011	-0.067	-0.034
Sum	-0.010	-0.035	-0.037	-0.022	-0.035	-0.028

Favorite-longshot bias

A bias in favor of both longshots and favorites seem to exist. This is in combination with the home field advantage for most leagues. The effect is especially strong in Spain, reinforcing the found home effect in Spanish football. There also seems to be somewhat of an away bias in the English Premier League.

While betting on home wins is preferable overall, there appears to be a distinct skew towards either high or low probability wins. Odds found in the probability range 0.25 – 0.75 yield a distinctly negative return, while odds outside this range are profitable. This result is also consistent across seasons.

Table 4.3: Mean profit from a bet on matches grouped by odds

	$p_j \leq 0.25$		$0.25 < p_j \leq 0.5$		$0.5 < p_j \leq 0.75$		$0.75 < p_j$	
	Home	Away	Home	Away	Home	Away	Home	Away
ENG	0.131	0.108	-0.125	-0.023	-0.019	-0.019	0.024	-0.133
ESP	0.343	-0.088	-0.030	-0.036	-0.035	-0.017	0.064	-0.098
FRA	-0.034	-0.106	-0.013	-0.023	-0.038	0.045	0.031	-0.360
GER	0.054	0.027	-0.010	-0.035	-0.035	-0.047	0.000	0.068
ITA	-0.220	-0.103	-0.001	-0.061	-0.015	0.077	-0.027	-0.040
Sum	0.073	-0.038	-0.034	-0.036	-0.028	0.003	0.024	-0.059

4.2.2 Prediction model

Using the criteria $\hat{p}_i > p_i$ bets are placed on 8,419 of 9,099 potential matches, with an average return of -0.016. The more discriminate $\hat{p}_i^{min} > p_i$ fares a little bit better with an average return of -0.012, placing bets on 6,925 matches.

From table 4.2 one can see the expected return on randomly placed bets is -0.028. The models consequently perform better than chance, but not well enough to cover the margin.

The negative return is present in all countries, except for Germany where both models turn a profit. There is also no improvement to be found across seasons, in fact season 2012 turns a small profit with 2016 yielding the greatest loss.

4.2.2.1 Bias impact

When betting only on one outcome, the strategy is simply that a bet is placed if the forecasted probability exceeds the offered odds on that outcome. A bet can therefore be placed on multiple outcomes.

Table 4.4: Mean profit on prediction model based bets

		Home	Draw	Away	Max
\hat{p}	μ	0.004	-0.054	-0.022	-0.016
	n	6,488	3,575	3,087	8,419
\hat{p}^{min}	μ	0.024	-0.059	-0.037	-0.012
	n	3,858	2,399	2,383	6,925

As expected the model performs significantly better at home wins than away and draws. The simple model manages to turn a tiny profit on home wins, while the conservative model is making a considerable profit.

Accounting for the favorite-longshot bias and only including matches with an odds of $p \leq 0.25$ or $p \geq 0.75$ would rise the mean return on home wins to 0.095 and 0.142 for the \hat{p} and \hat{p}^{min} strategies respectively. The extra matches included in the \hat{p} strategy contribute a mean of 0.004, so the increased volume does not provide a substantial increase in profits, but mostly serves to increase variance.

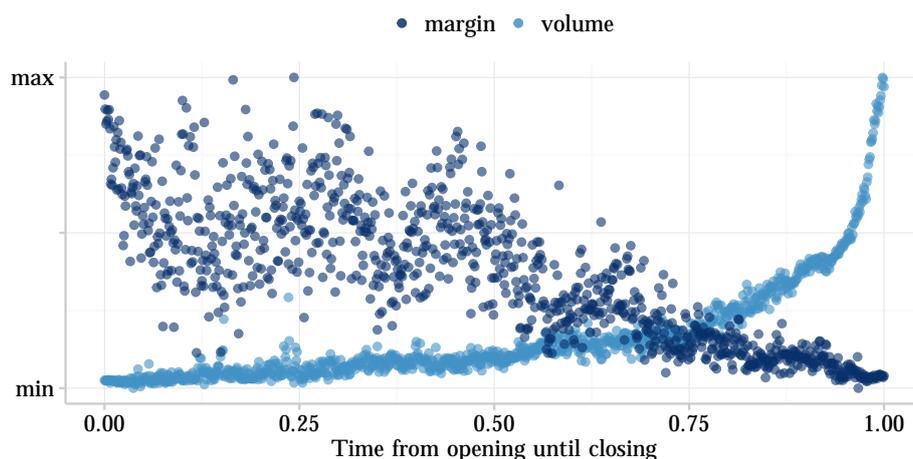
4.3 Information uptake

4.3.1 Price movement

Odds for the 2016 season are posted typically 10 days before an event takes place, but with much lower volume and higher margins than what is available at closing. As time progress margins steadily decrease, while volume and activity increase. More than 60% of price movements take place at match day, 25% within two hours before kick-off.

Initial postings generally offer a volume of EUR 1,000 and an average margin of 2.81%, while closing odds are available at a volume of EUR 10,000 and a 2.03% margin. The magnitude in price changes also notably decrease, mirroring the volume being offered. The relationship between margin and volume over time can be observed in figure 4.3.

Figure 4.3: Margin and volume relative to their respective maximum/minimum values



The mean difference in the highest and lowest observed market probabilities for home and away wins are 5.97 pp and 5.37 pp respectively, while there is much less movement in the draw price where the difference is only 2.33 pp.

These magnitudes are considerably larger than the margin, and it is therefore possible to make a profit by betting at optimal points in time. If a bet is placed on a home win when the odds is at its highest, this position can be closed with a profit by betting on a draw and away win when the home odds is at its lowest.

As a consequence of the decreasing margins, mean offered probabilities are lower (i.e. higher odds are offered) at closing than opening. All else being equal, that is if the market were efficient, closing odds should be preferred over opening odds.

A Pearson correlation test finds that the difference in adjusted opening and clos-

ing probabilities is significantly correlated with match result for all three outcomes (A.6). This indicates that closing odds are more correct than opening odds and that the market is absorbing information.

4.3.2 Opening and closing odds

With closing being more accurate than opening odds, one would expect the observed biases to either decrease in magnitude or evaporate fully. The overall most consistent bias observed is the favorite and longshot bias for home wins.

Closing odds for home wins is lower than opening odds indicating some correction. More specifically the home favorite and longshot observe decreasing odds, while the odds for the middle probability outcomes is slightly raised. Mostly decreasing margins however seem to benefit the draw and away.

Table 4.5: Mean return using first available and closing odds on profitable strategies

Strategy	Description	First	Close
$p_{i,H,t=0} \leq 0.25$	Home longshot bias selected by opening odds	0.073	0.027
$p_{i,H,t=1} \leq 0.25$	Home longshot bias selected by closing odds	–	0.005
$p_{i,H,t=0} > 0.75$	Home favorite bias selected by opening odds	0.024	0.024
$p_{i,H,t=1} > 0.75$	Home favorite bias selected by closing odds	–	0.019
$\hat{p}_{i,H}^{max} > p_{i,H,t=f}$	Predicted home win selected by first odds	0.024	0.014
$\hat{p}_{i,H}^{max} > p_{i,H,t=1}$	Predicted home win selected by closing odds	–	-0.011

Matches can potentially change odds brackets from opening until closing, so both opening and closing odds is used as a selection criteria. It’s worth noting that for the prediction model “first” odds refer to the first available odds when the information incorporated in the pi rating is available. With overlapping results in past seasons this will be the closing odds, as for the 2016 season it will be the first odds observed after all previous results for both teams are observed.

All strategies are more profitable using first available odds, except the home favorite bias which does not seem much affected. Using first odds as a selection criterion however remains a superior criterion for all strategies. This is probably due to odds without the inherent bias moving brackets, and odds with observed bias being pushed out of the bracket.

The results are consistent with closing odds being superior to opening odds, and the market being more information complete.

4.3.3 Information signal

The data set contains a subset of 1,266 matches where there is an overlap in odds postings and an interim match for at least one of the two opposing teams. Outcomes for the match are estimated by the regression models before and after the new results become available, denoted as \hat{p} and \hat{z} respectively.

Selecting matches where $\hat{p}_{i,j}^{min}$ differs from $\hat{z}_{i,j}^{max}$ or vice-versa turns out to be impractical as they overlap too much and such a criterion would yield a single digit number of matches. A simple criteria of $\hat{p}_{i,j} > \hat{z}_{i,j}$ will therefore be applied.

Draws are ignored as it appears mostly to serve as an indicator for an increased chance of either a home or away win. There is a large degree of overlap of matches with an increase in drawing with matches that have an increase in either a home or away win.

4.3.3.1 Odds difference

One would expect to see a decrease in odds on a home win, i.e. increase in offered probabilities, for matches where $\hat{p}_{i,H} > \hat{z}_{i,H}$. Similar effects on the away odds is expected for $\hat{p}_{i,A} > \hat{z}_{i,A}$. Adjusted probabilities will be used for comparison to account for any changes in margin.

A paired t-test is conducted showing a significant change in adjusted probabilities before and after the last interim match (A.7). Matches where $\hat{p}_{i,H} > \hat{z}_{i,H}$ see an increase in offered home probabilities, while matches where $\hat{p}_{i,A} > \hat{z}_{i,H}$ observe an increase in offered away probabilities.

The effect persists when using unadjusted probabilities, there is in other words a consistent, predictable change in offered odds. The median time before odds are adjusted is about 3 hours after the interim match is finished, indicating a certain lag in information uptake.

Even though a change in odds can be predicted, the magnitude is not large enough to eliminate the bookmaker margin and can as such not be directly use to profit from arbitrage. Additionally, only the next tick in odds changes after a result is significant. Subsequent ticks are not significantly different from pre-result observed odds.

The difference in adjusted probabilities between pre-match and closing is not significant either. This could be due to other information of higher importance closer to match start dissipating the effect, or it could be due to an undervaluation of the information signal. The last effect seems like a plausible as hypothesis due to the effect of the signal disappearing after a single tick.

4.3.3.2 Prediction strategy

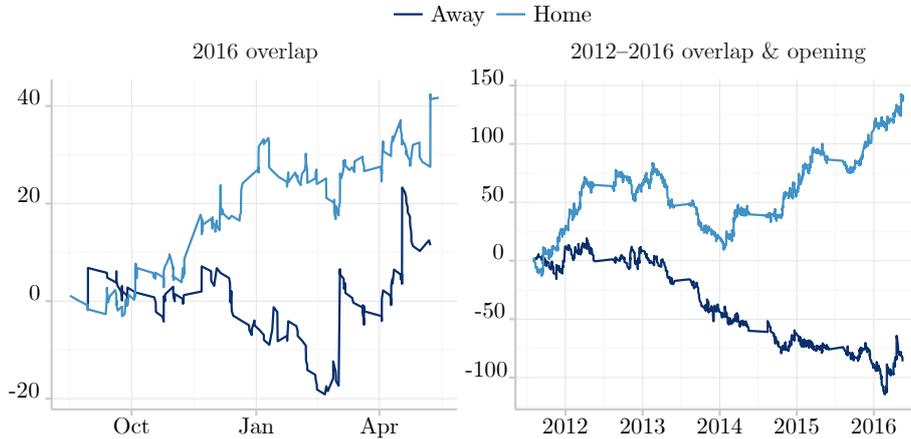
A strategy is constructed for home and away wins. Home and away outcomes where the significant model indicates an underestimation in the market, and there is an increase in the predicted likelihood of the given outcome will be selected. Bets will be placed on these matches using the odds available on first tick after the appearance of new information. The selection strategy can be expressed as $\hat{p}_{i,j} > \hat{z}_{i,j} \cap \hat{p}_{i,j}^{min} > p_{i,j}$.

Both strategies generate positive returns, but the returns for the away strategy are highly volatile and a positive return is only observed due to two extreme long shots with odds at 20.75 and 24.9 going in. The impact of these two matches can be seen clearly as a spike in the away line for 2016 in figure 4.4.

The home strategy looks stable, and generates an average return of 0.173, netting a profit of 41.71 over 240 matches during the 2016 season. An interesting observation is that matches that do not have an overlapping result also generate a profit using the home win strategy. Combining all home matches for the 2016 season matching the criteria yield a mean return of 0.152 and total of 53.16 over 348 matches. Running a two-tailed t-test on the returns one can reject the null hypothesis of a zero mean at a 90% level.

This effect is not present for away matches, but instead generates a negative return further underlining the outlier effect observed.

Figure 4.4: Cumulative returns for selected matches



4.3.3.3 Historical results

In the 2016 season a majority of matches in the data set had overlapping results. For the seasons 2012–2015 however there is only between 42 and 95 matches with overlapping odds per season. Possible reasons for this will be mentioned in section

5.3, but it also impacts the selection criteria when testing the historical results of the strategy. For an equal comparison between opening and closing odds these overlapping matches have been excluded.

A summary of the strategies is shown in table 4.6. The away strategy observes significant losses when including matches before 2016. This again argues for the 2016 season profits in the away match strategy being a fluke.

There are significant positive returns for home strategy also when including previous seasons. Using first available odds the home strategy generates a considerable profit of 0.074 per match. A two-tailed t-test rejects the null hypothesis of a zero mean at a 95% level.

As a control matches where the prediction model indicates an increase in the probability of a home win, but no indication of underestimation in the market is tested. This strategy gives a clear negative return for both home and away matches.

Table 4.6: Strategy summary

Strategy	Count	Mean		Sum	
		First	Close	First	Close
Home overlap	242	0.173	0.155	41.71	37.43
Away overlap	158	0.074	0.087	13.51	11.51
Home all	1,877	0.074	0.057	139.69	106.88
Away all	1,009	-0.085	-0.100	-86.25	-100.56
Home control	2,701	-0.046	-0.056	-124.88	-151.79
Away control	3,255	-0.037	-0.036	-121.77	-117.91

More favorable odds are observed earlier in line with previous findings. At closing odds the profitability of the strategies is notably decreased. The mean volume observed at the optimal staking time in 2016 is about 1,500 EUR, compared to a mean volume of around 9,000 EUR at closing time.

The profitable home betting strategy both for matches with and without an overlapping result have substantially larger returns when using first possible odds than closing odds. They are however still profitable, this suggests there is some market correction, but that the full effect is not picked up.

A closer look at the odds distribution of matches included in the strategy reveals that a majority of the profit is generated by longshots. Matches where $p < 0.25$ have a mean of 0.242 and account for about 60% of the total observed profits, but only 18% of the matches (A.5).

5 Discussion

The aim of this thesis was to explore the efficiency of the unrestricted betting markets. Weak form efficiency has been investigated through the assessment of known biases in odds setting. Semi-strong efficiency has been assessed by creating a prediction model based on historical data and examining the impact of new information through match results becoming available.

5.1 Results

5.1.1 Bias

A consistent home bias is observed, where betting on home wins only yield more favorable results than betting on draws or away wins. A favorite-longshot as well as a longshot bias is also observed for home wins. There is not found a consistent bias in the pricing of away wins or draws, apart from being less favorable than home wins.

The observed bias varies some by league, and a home favorite seem especially strong in Spain. The Spanish favorites are likely to be dominated by the big teams, and confirms the effect found by Forrest and Simmons (2008). A more thorough analysis broken down by league is not performed as the differences do not seem that significant, and to avoid any data mining.

This is consistent with research by Constantinou and Fenton (2013b) who find considerably higher returns from home wins than away wins or draws. Additionally, the results support their findings of home favorites and longshots being priced more favorably. They have also described an away favorite bias which can not be established from the data used in thesis.

These results however directly contradict the findings by Vlastakis et al. (2009) who find that the home-field advantage is overestimated by bookmakers once the favorite-longshot bias is accounted for. In this thesis a favorite-longshot bias is on the contrary only observed for home wins.

This appararent discrepancy could lie in a different definition of favorite. Constantinou and Fenton define favorites as having a probability higher than 60%, 70% or 80% when testing the effect. Vlastakis et al. on the other hand define the favorite as the team assigned highest probability of winning. In this thesis a simple criterion

of a probability higher than 75% is used, but the effect is also seen using 70% or 80%.

What is new in this thesis is both the observation of bias in an unrestricted football betting market, but also the results from observing odds at two points in time. The home favorite bias is robust and observed in both opening and closing odds without any large difference in profitability. The home longshot bias however is greatly reduced by closing odds.

It must also be noted that while these biases are observed to be profitable, and are consistent with theory and previous findings, the variance is high and one cannot conclude that they are statistically significantly profitable by themselves. In fact only the returns from the home favorite bias is statistically different from the returns from placing other bets.

A bias observed with a traditional bookmaker can be explained by the bookmaker optimizing on customer preference. With an unrestricted bookmaker this explanation is not as likely as they face larger market pressure. A point in favor of this behavior however is that opening odds which also include home longshot biased odds are set by Pinnacle themselves, and have not been subject to market pressure. Another potential reason is that the betting market is actually dominated by bettors and not professional actors as observed by Paul and Weinbach (2010). Protection against insider information as noted by Shin (1991) does not seem like a likely explanation for the top European leagues where most information should be widely available.

5.1.2 Prediction model

The prediction model used the pi-rating system developed by Constantinou and Fenton (2013a) in a probit model. In addition, the “big team” effect described by Goddard and Asimakopoulos (2004) was added. The model was not able to make a profit when betting on the best predicted outcome against market odds, draws or away wins. When betting on the best predicted outcome the model incurred a loss smaller than the bookmaker margin, and it was able to make a profit by predicting home wins.

Hvattum and Arntzen (2010) also observed that their logit model based on Elo-ratings was significantly more successful at predicting home wins than other results. While the pi-rating system uses separate ratings for home and away matches, the effect persisted in the model constructed in this thesis.

A potential explanation could be that home wins are easier to predict than away wins or draws. A different explanation lies in the economic evaluation criteria, where

observed margins were higher for away wins and draws. The margin is however far lower than the profits gained by the optimal home betting strategies.

The home profit is increased when focusing only on home favorites and longshots, and is then statistically significantly profitable. This is in line with the findings from Constantinou and Fenton (2013a) who observes that their pi-ratings based betting strategy generates profits from long-shots.

There is a selection problem when concluding that the biased based home win strategy is profitable. While the observed home biases are in line theory and expected findings, the decision to select on this bias criteria for the prediction model was taken after observing it in the results.

5.1.3 Information signal

An effect on odds by new information in the form of results becoming available has been demonstrated, it however takes some time before it is observed in the market. Some time lag is in line with the previously mentioned findings on arbitrage by Marshall (2009). While the majority of mispricing events disappear in minutes in his research, the observed lag in information uptake found in this thesis is measured in hours.

A tick is an observed odds change in the market. This signal disappears after a single tick suggesting that the information is either not very important and overshadowed by other information, or not incorporated fully in pricing. A strategy is therefore constructed using the prediction model, selecting home matches where the model observes both an increased chance for a home/away win as well as a higher probability for a home/away win than the market. This strategy is when placing bets on the home win, able to earn a statistically significant profit.

These results suggest that the most recent information about a team is under-rated by the market for a subset of matches. The prediction model was needed as an additional qualifier, a change in predicted probability was in itself not a good selection criteria for which matches to bet on. This could be due to the weak information value of the signal by itself. For most matches there will be a change in predicted probability, and the strength of the signal was not measured.

5.1.4 Information uptake

Closing odds are more accurate than opening odds, and trading volume is significantly higher at closing than opening. Most price movements occur at match day. Strategies constructed both using biases and the prediction model are substantially

less profitable using closing odds than first available odds.

All strategies earning a significant positive return using first available odds, become less profitable and statistically not significantly different from 0. Additionally match results are significantly correlated with the change in odds for all outcomes. This indicates a clear movement towards a more informed market, and is in line with the findings from the NBA betting market by Gandar et al. (1998) and Baryla et al. (2007).

That the strategies remain profitable suggests that the bias is not fully accounted for, but merely partially corrected.

5.2 Efficiency

To a large degree betting is a kind of consumption. Bettors buy a form of entertainment by placing bets on football matches, an efficient market is therefore not necessarily expected. The bookmaker can also select and price discriminate its customers. The market examined in this thesis however does not exclude customers making a profit, is of an online nature, and provides programming interfaces to facilitate computer trading and should attract professional investors.

In an efficient market one should not be able to make a profit in the absence of a risk premium, which is a profit made for carrying systemic risk. In betting there is no obvious systemic risk, and market inefficiencies must be seen in light of this.

5.2.1 Weak form efficiency

Weak form inefficiency in the market is observed in the form a home bias on favorites and longshots. For the home favorite bias, it also persists from opening until closing. While it cannot be concluded that it is profitable on a significant level, it can be concluded it offers more favorable odds than other outcomes.

This violates weak form market efficiency as defined by Fama (1991), or strong form betting market efficiency as defined by Thaler and Ziemba (1988), that all outcomes should have an equal expected value. This is in line with multiple findings of weak market inefficiencies in the football betting market, such as Forrest and Simmons (2008), Vlastakis et al. (2009), Franck et al. (2013), and Constantinou and Fenton (2013b).

What is new in this thesis is that bias is observed in an unrestricted betting market, and that some bias is corrected for by the market, while the home favorite-longshot is not fully corrected for.

5.2.2 Semi-strong form efficiency

Semi-strong form inefficiency in the market is observed by the prediction model earning a statistically significant profit when observing an increased chance for a home win along with favorable odds. Additionally, combining a prediction model with market bias also seems to be profitable.

Goddard and Asimakopoulos (2004) and Vlastakis et al. (2009) find evidence that it is possible to make a profit using statistical models in the football betting markets. This thesis confirms these findings, and further contributes that this can also be observed in an unrestricted football betting market.

Additionally, it is observed that information is to some degree absorbed by the market, and that markets grow more efficient as match start approaches. The traded volume is also greater at closing than earlier points in time. This could indicate that the bookmaker intentionally sets a low volume and early odds in order to absorb information from informed actors or establish betting patterns.

5.3 Information overlap

An interesting observation is that a majority of odds posted for matches in the 2016 season have interim matches for the opposing teams, while this is not observed for vast majority of the matches in the seasons 2012–2015.

One possible explanation is that the odds recorded by `oddsportal.com` is not accurate enough, and that odds are actually posted earlier. This does not seem like a likely explanation, as odds recorded from the Pinnacle API are consistent with odds taken from `oddsportal.com` for the 2016 season. A more plausible explanation is that Pinnacle, and perhaps other bookmakers, are posting odds earlier to attract more bettors and gain a competitive advantage.

A different explanation is offered by the reduction in margin over time. This indicates that there is more uncertainty about the odds at earlier points in time. Pinnacle might offer odds at a lower margin well ahead of match start in order to absorb information at a lower exposure. This could however also be due to the uncertainty introduced by a number of unknown factors that are revealed closer to match start.

Since this overlap is mostly observed for the 2016 season, it is not given that this pattern will continue in the future.

5.4 Limitations and weaknesses

The regression model used can be improved upon. Country and distance effects were found to be significant predictors for match result, but were dropped due to model limitations. Potential interaction effects by country have not been evaluated either. The model is also only an approximation based on team statistics. Qualitative information such as injuries is not taken into account, and most news signals will be of this nature.

The low impact of the tested information signal in form of new results could be due to other factors being more important, and that the signal is mostly a valuable estimator in the context of the prediction model. This makes it difficult to measure how the market reacts to news, when it does not actually pick up the news signal fully or the signal is noisy.

Another weakness is that there is no measurement of the strength of the signal. New results that are mostly in line with what was already expected do not contribute much new information.

The model performs very well on home matches, especially when considering observed biases. This effect is however found after taking observed characteristics on the out-of sample data into account. These characteristics were anticipated, but further out-of sample data would be beneficial to eliminate any selection bias introduced.

6 Conclusion

This thesis has explored the efficiency of the football betting market by examining the pricing of a bookmaker that does not impose any limits on profits. Weak and semi-strong form inefficiency has been established, and profitable strategies have been developed.

6.1 Main findings

Evidence of weak form efficiency has been found through bias in pricing. Findings suggest that this effect is for some observed biases weakened by trading, and is less pronounced earlier to match start. While the observed bias is profitable, it can only be concluded that it is more favorable than alternative bets.

Strong form inefficiency has been established in the form of a prediction model providing finding outcomes where the market probabilities were underestimated. Additionally, it has been demonstrated that for a subset of matches the market underestimates the impact of new information in the form of match results.

The market absorbs information and is more efficient at match start than days before match start. Trading volume is also observed to be limited when information is at its most valuable. The change in odds is significantly correlated with match results for all three outcomes.

6.2 Recommendations for further research

The results in this thesis strengthens the view that there are common biases in the betting market, and a profitable strategy is demonstrated exploiting a home favorite and long-shot bias. These selection criteria based on bias would benefit from being tested on out-of sample data.

Closing odds are shown to be better predictors than opening odds, but there is significant market movement and at which point in the process information is picked up has not been looked into. More research into the price formation of the market, especially if incorporating stronger information signals, would be an interesting avenue of research.

References

- E. A. Baryla, Jr., R. A. Borghesi, W. H. Dare, and S. A. Dennis. Learning, Price Formation and the Early Season Bias in the NBA. *Finance Research Letters*, 4 (3):155–164, 2007.
- R. Borghesi. The home team weather advantage and biases in the NFL betting market. *Journal of Economics and Business*, 59:340–354, 2007.
- G. W. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78:1–3, 1950.
- A. Brown. Evidence of in-play insider trading on a UK betting exchange. *Applied Economics*, 44(9):1169–1175, 2012.
- M. Cain, D. Law, and D. Peel. The Favourite-Longshot Bias and Market Efficiency in UK Football Betting. *Scottish Journal of Political Economy*, 2000.
- M. Cain, D. Law, and D. Peel. The favourite-longshot bias, bookmaker margins and insider trading in a variety of betting markets. *Bulletin of Economic Research*, 55 (3):263–273, 2003.
- R. H. B. Christensen. ordinal: Regression Models for Ordinal Data, 2015. URL <https://cran.r-project.org/web/packages/ordinal/>. Last accessed 2017-01-08.
- S. R. Clarke and J. 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.
- A. C. Constantinou and N. E. Fenton. Determining the level of ability of football teams by dynamic ratings based on the relative discrepancies in scores between adversaries. *Journal of Quantitative Analysis in Sports*, 9(1):37–50, 2013a.
- A. C. Constantinou and N. E. Fenton. Profiting from arbitrage and odds biases of the European football gambling market. *Journal of Gambling Business and Economics*, 7(2):41–70, 2013b.

- D. Cortis. Expected Values and Variances in Bookmaker Payouts: A Theoretical Approach Towards Setting Limits on Odds. *The Journal of Prediction Markets*, 9(1):1–14, 2015.
- Crowdpark. Segmentation of online gambling market as of 2012, 2012. URL <https://www.statxista.com/statxistics/248655/segmentation-of-online-gambling-market/>. Last accessed 2017-01-08.
- K. Crosson and J. J. Reade. Information and efficiency: Goal arrival in soccer betting. *The Economic Journal*, 124(575):62–91, 2014.
- M. J. Dixon and S. 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.
- S. Dobson and J. Goddard. *The Economics of Football*. Cambridge University Pressm, 2001.
- A. Elo. *The rating of chess players, past and present*. Arco, 1978.
- E. F. Fama. Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, 25(2):383–417, 1970.
- E. F. Fama. Efficient Capital Markets: ii. *The Journal of Finance*, 46(5):1575–1617, 1991.
- D. Forrest and R. Simmons. Sentiment in the betting market on spanish football. *Applied Economics*, 40(1):119–126, 2008.
- D. Forrest, J. Goddard, and R. Simmons. Odds-setters as forecasters: The case of English football. *International Journal of Forecasting*, 21(3):551–564, 2005.
- E. Franck, E. Verbeek, and S. Nüesch. Prediction Accuracy of Different Market Structures – Bookmakers Versus a Betting Exchange. *International Journal of Forecasting*, 26(3):448–459, 2010.
- E. Franck, E. Verbeek, and S. Nüesch. Inter-market Arbitrage in Betting. *Economica*, 80(318):300–325, 2013.
- J. M. Gandar, W. H. Dare, C. R. Brown, and R. A. Zuber. Informed Traders and Price Variations in the Betting Market for Professional Basketball Games. *The Journal of Finance*, 53(1):385–401, 1998.

- J. M. Gandar, R. A. Zuber, and R. S. Johnson. A reexamination of the efficiency of the betting market on national hockey league games. *Journal of Sports Economics*, 5(2):152–168, 2004.
- J. Goddard. Regression models for forecasting goals and match results in association football. *International Journal of Forecasting*, 21(2):331–340, 2005.
- J. Goddard and I. Asimakopoulos. Modelling football match results and the efficiency of fixed-odds betting. *Journal of Forecasting*, 23(1):51–66, 2004.
- I. Graham and H. Stott. Prediction bookmaker odds and efficiency for UK football. *Applied Economics*, 40:99–109, 2008.
- L. M. Hvattum and H. Arntzen. Using ELO ratings for match result prediction in association football. *International Journal of Forecasting*, 26(3):460–470, 2010.
- James Stocks & Co and KPMG. Size of the online gambling market from 2009 to 2018, 2016. URL <https://www.statxista.com/statxistics/270728/market-volume-of-online-gaming-worldwide/>. Last accessed 2017-01-08.
- A. Joseph, N. Fenton, and M. Neil. Predicting football results using Bayesian nets and other machine learning techniques. *Knowledge-Based Systems*, 19(7):544–553, 2006.
- J. L. Kelly, Jr. A new interpretation of information rate. *The Bell System Technical Journal*, 35(4):917–926, 1956.
- R. H. Koning and B. van Velzen. Betting exchanges: the future of sports betting? *International Journal of Sport Finance*, 4(1):42–62, 2009.
- T. Kuypers. Information and efficiency: an empirical study of a fixed odds betting market. *Applied Economics*, 32(11), 2000.
- S. D. Levitt. Why are Gambling Markets Organized so Differently from Financial Markets? *Economic Journal*, 114(495):223–246, 2004.
- B. R. Marshall. How quickly is temporary market inefficiency removed? *The Quarterly Review of Economics and Finance*, 49:917–930, 2009.
- T. J. Moskowitz. Asset pricing and sports betting. *Chicago Booth Research Paper*, (15-26), 2015.

- R. J. Paul and A. P. Weinbach. The Determinants of Betting Volume for Sports in North America: Evidence of Sports Betting as Consumption in the NBA and NHL. *International Journal of Sport Finance*, 5:128–140, 2010.
- H. Rue and Ø. Salvesen. Prediction and retrospective analysis of soccer matches in a league. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 49(3):399–418, 2000.
- H. S. Shin. Optimal betting odds against insider traders. *The Economic Journal*, 101:1179–1185, 1991.
- W. W. Snyder. Horse Racing: Testing the Efficient Markets Model. *The Journal of Finance*, 33(4):1109–1118, 1978.
- R. H. Thaler and W. T. Ziemba. Parimutuel Betting Markets: Racetracks and Lotteries. *Journal of Economic Perspectives*, 2(2):161–174, 1988.
- UCLA Institute for Digital Research And Education. R data analysis examples: Ordinal logistic regression, 2016. URL <http://www.atxs.ucxla.edu/stat/r/dae/ologit.htm>.
- M. Verbeek. *A Guide to Modern Econometrics*. John Wiley & Sons, Ltd, 2 edition, 2004.
- N. Vlastakis, G. Dotsis, and R. N. Markellos. Nonlinear Modeling of European Football Scores Using Support Vector Machines. *Applied Economics*, 40(1):111–118, 2008.
- N. Vlastakis, G. Dotsis, and R. Markellos. How Efficient is the European Football Betting Market? Evidence from Arbitrage and Trading Strategies. *Journal of Forecasting*, 28:426–444, 2009.
- L. V. Williams. Information efficiency in betting markets: a survey. *Bulletin of Economic Research*, 51(1):1–39, 1999.
- L. V. Williams, editor. *Information Efficiency in Financial and Betting Markets*. Cambridge University Press, 2005.

A Appendix

A.1 Pi-rating calibration

Table A.1: Mean squared error for pi-ratings

$\lambda - \gamma$	0.25	0.30	0.35	0.40	0.45	0.50	0.55	0.60	0.65	0.70	0.75
0.025	2.4521	2.4502	2.4487	2.4477	2.4473	2.4475	2.4485	2.4506	2.4539	2.4590	2.4664
0.030	2.4456	2.4438	2.4424	2.4414	2.4407	2.4406	2.4411	2.4425	2.4450	2.4491	2.4555
0.035	2.4438	2.4422	2.4408	2.4397	2.4390	2.4386	2.4387	2.4396	2.4414	2.4446	2.4499
0.040	2.4448	2.4433	2.4420	2.4409	2.4401	2.4396	2.4394	2.4399	2.4411	2.4436	2.4479
0.045	2.4475	2.4461	2.4449	2.4439	2.4430	2.4424	2.4421	2.4423	2.4431	2.4450	2.4485
0.050	2.4514	2.4501	2.4490	2.4480	2.4472	2.4466	2.4462	2.4462	2.4467	2.4481	2.4509
0.055	2.4562	2.4550	2.4540	2.4530	2.4522	2.4516	2.4512	2.4511	2.4514	2.4525	2.4547
0.060	2.4617	2.4606	2.4596	2.4587	2.4580	2.4574	2.4570	2.4568	2.4570	2.4578	2.4597
0.065	2.4678	2.4667	2.4658	2.4650	2.4643	2.4637	2.4633	2.4632	2.4633	2.4639	2.4654
0.070	2.4744	2.4734	2.4725	2.4717	2.4710	2.4705	2.4702	2.4700	2.4701	2.4706	2.4719
0.075	2.4814	2.4804	2.4796	2.4789	2.4782	2.4778	2.4774	2.4773	2.4774	2.4778	2.4789

A.2 Regression model 1 – 2011

A.2.1 Initial specification

```
link   threshold nobs logLik   AIC       niter max.grad cond.H
probit flexible  7758 -7731.61 15491.23 6(0)   3.29e-12 1.2e+05
```

Coefficients:

	Estimate	Std. Error	z value	Pr(> z)	
pi_diff	0.618390	0.024573	25.166	< 2e-16	***
ap_diff	0.121287	0.023995	5.055	4.31e-07	***
dap_diff	-0.003365	0.036430	-0.092	0.9264	
cup_diff	-0.036185	0.028327	-1.277	0.2015	
promoted_diff	0.017355	0.033155	0.523	0.6007	
weather_bad	0.006858	0.043683	0.157	0.8753	
distance	0.028610	0.015115	1.893	0.0584	.
derby	0.829187	0.323171	2.566	0.0103	*
countryESP	-0.013360	0.044212	-0.302	0.7625	
countryFRA	-0.095717	0.045196	-2.118	0.0342	*
countryGER	-0.090710	0.041812	-2.169	0.0300	*
countryITA	-0.039710	0.044501	-0.892	0.3722	

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Threshold coefficients:

	Estimate	Std. Error	z value
0 0.5	-0.1303	0.1791	-0.728
0.5 1	0.6259	0.1792	3.493

A.2.2 First stepwise model

```
link   threshold nobs logLik   AIC       niter max.grad cond.H
probit flexible  7758 -7732.59 15485.18 6(0)   2.78e-12 1.2e+05
```

Coefficients:

	Estimate	Std. Error	z value	Pr(> z)	
pi_diff	0.61978	0.02175	28.494	< 2e-16	***
ap_diff	0.12313	0.02289	5.378	7.52e-08	***
distance	0.02875	0.01511	1.902	0.05715	.
derby	0.83512	0.32304	2.585	0.00973	**
countryESP	-0.01389	0.04420	-0.314	0.75336	
countryFRA	-0.09556	0.04491	-2.128	0.03336	*
countryGER	-0.09124	0.04181	-2.182	0.02908	*
countryITA	-0.04046	0.04448	-0.910	0.36302	

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Threshold coefficients:

	Estimate	Std. Error	z value
0 0.5	-0.1291	0.1790	-0.721
0.5 1	0.6270	0.1791	3.501

A.2.2.1 Likelihood-ratio test

	Df	logLik	AIC	LRT	Pr(>Chi)
<none>		-7732.6	15485		
pi_diff	1	-7732.0	15486	1.1126	0.291517
ap_diff	1	-7732.5	15487	0.1561	0.692800
distance	1	-7732.4	15487	0.2882	0.591381
derby	1	-7728.8	15480	7.4927	0.006195 **
country	4	-7723.9	15476	17.3088	0.001683 **

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

A.2.3 Final model

link	threshold	nobs	logLik	AIC	niter	max.grad	cond.H
probit	flexible	7758	-7739.67	15487.34	6(0)	3.40e-13	6.8e+00

Coefficients:

	Estimate	Std. Error	z value	Pr(> z)
pi_diff	0.61946	0.02173	28.511	< 2e-16 ***
ap_diff	0.12252	0.02289	5.353	8.64e-08 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Threshold coefficients:

	Estimate	Std. Error	z value
0 0.5	-0.44273	0.01691	-26.18
0.5 1	0.31245	0.01664	18.77

A.2.3.1 Likelihood-ratio test

	Df	logLik	AIC	LRT	Pr(>Chi)
<none>		-7739.7	15487		
pi_diff	1	-7739.2	15488	1.01278	0.3142
ap_diff	1	-7739.6	15489	0.13222	0.7161

A.3 Regression model 2 – 2011

A.3.1 Initial specification

```
link   threshold nobs logLik   AIC       niter max.grad cond.H
probit flexible  7740 -7721.63 15461.25 6(0)   1.74e-12 3.4e+04
```

Coefficients:

	Estimate	Std. Error	z value	Pr(> z)	
pi_diff_1	0.616745	0.024562	25.110	< 2e-16	***
ap_diff	0.122630	0.024013	5.107	3.27e-07	***
dap_diff	-0.005428	0.036500	-0.149	0.882	
cup_diff	-0.037008	0.028325	-1.307	0.191	
promoted_diff	0.015037	0.033294	0.452	0.652	
weather_bad	-0.005922	0.043182	-0.137	0.891	
distance	0.002536	0.011008	0.230	0.818	

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Threshold coefficients:

	Estimate	Std. Error	z value
0 0.5	-0.4124	0.1380	-2.989
0.5 1	0.3427	0.1380	2.484

A.3.2 Final model

```
link   threshold nobs logLik   AIC       niter max.grad cond.H
probit flexible  7740 -7722.63 15453.26 6(0)   5.43e-13 6.8e+00
```

Coefficients:

	Estimate	Std. Error	z value	Pr(> z)	
pi_diff_1	0.61921	0.02177	28.438	< 2e-16	***
ap_diff	0.12399	0.02290	5.414	6.16e-08	***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Threshold coefficients:

	Estimate	Std. Error	z value
0 0.5	-0.44237	0.01694	-26.12
0.5 1	0.31259	0.01667	18.75

A.3.2.1 Likelihood-ratio test

```
link   threshold nobs logLik   AIC       niter max.grad cond.H
probit flexible  7740 -7722.63 15453.26 6(0)   5.43e-13 6.8e+00
```

Coefficients:

```
          Estimate Std. Error z value Pr(>|z|)
pi_diff_1  0.61921    0.02177  28.438 < 2e-16 ***
ap_diff    0.12399    0.02290   5.414 6.16e-08 ***
```

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Threshold coefficients:

```
          Estimate Std. Error z value
0|0.5 -0.44237    0.01694 -26.12
0.5|1  0.31259    0.01667  18.75
```

A.4 Regression model evolution

Table A.2: Model coefficient values based on training data from different seasons

Seasons	Model 1				Model 2			
	0-0.5	0.5-1	pi_diff	ap_diff	0-0.5	0.5-1	pi_diff	ap_diff
2007–2011	-0.4480	0.3131	0.6259	0.1112	-0.4472	0.3143	0.6266	0.1122
2007–2012	-0.4435	0.3178	0.6251	0.1061	-0.4426	0.3188	0.6250	0.1073
2007–2013	-0.4313	0.3294	0.6349	0.1031	-0.4304	0.3304	0.6353	0.1042
2014–2014	-0.4172	0.3346	0.6451	0.1021	-0.4161	0.3359	0.6461	0.1029
2015–2014	-0.4133	0.3382	0.6424	0.0918	-0.4127	0.3389	0.6426	0.0930

A.5 Opening and closing odds

Table A.3: Mean opening and closing probabilities formatted as odds

	Home	Draw	Away
Open	2.148	3.822	3.381
Close	2.146	3.841	3.395

Table A.4: Mean opening and closing probabilities for home wins grouped by opening odds

	$p \leq 0.25$	$0.25 < p \leq 0.75$	$0.75 < p$
Open	5.957	2.095	1.216
Close	5.793	2.093	1.214

Table A.5: Odds distribution for information based home win strategy

	Mean	Count	Sum
$p < 0.25$	0.242	340	82.28
$0.25 \leq p < 0.75$	0.035	1,511	53.11
$p \geq 0.75$	0.165	26	4.30

A.6 Price change/match result correlation

Pearson's product-moment correlation

```
data: pa_home.diff and res_home
t = 6.864, df = 9097, p-value = 7.135e-12
alternative hypothesis: true correlation is not equal to 0
95 percent confidence interval:
 0.05130832 0.09219193
sample estimates:
      cor
0.07178028
```

```
-----
data: pa_draw.diff and res_draw
t = 4.6402, df = 9097, p-value = 3.529e-06
alternative hypothesis: true correlation is not equal to 0
95 percent confidence interval:
 0.02807361 0.06907188
sample estimates:
      cor
0.04859321
```

```
-----
data: pa_away.diff and res_away
t = 6.9843, df = 9097, p-value = 3.064e-12
alternative hypothesis: true correlation is not equal to 0
95 percent confidence interval:
 0.05256288 0.09343904
sample estimates:
      cor
0.07303163
```

A.7 Odds change before and after new results

Paired t-test

```
data: pa_home.pre and pa_home.post
t = -6.981, df = 639, p-value = 7.362e-12
alternative hypothesis: true difference in means is not equal to 0
95 percent confidence interval:
 -0.005214183 -0.002924771
sample estimates:
mean of the differences
      -0.004069477
```

```
-----
data: pa_home.post and pa_home.close
t = -0.58404, df = 639, p-value = 0.5594
alternative hypothesis: true difference in means is not equal to 0
95 percent confidence interval:
 -0.003421487  0.001852810
sample estimates:
mean of the differences
      -0.0007843385
```

```
-----
data: p_home.pre and p_home.post
t = -5.2227, df = 639, p-value = 2.389e-07
alternative hypothesis: true difference in means is not equal to 0
95 percent confidence interval:
 -0.004352775 -0.001973976
sample estimates:
mean of the differences
      -0.003163376
```

```
-----
data: p_home.post and p_home.close
t = 1.4966, df = 639, p-value = 0.135
alternative hypothesis: true difference in means is not equal to 0
95 percent confidence interval:
 -0.000642912  0.004763343
sample estimates:
mean of the differences
      0.002060215
```