Aalto University
School of Business

Abnormal returns in an efficient market?

Statistical and economic weak form efficiency of online sports betting in European soccer

Finance

Master's thesis

Tero Tiitu

2016

Department of Finance
Aalto University
School of Business

Aalto University, P.O. BOX 11000, 00076 AALTO
www.aalto.fi

Abstract of master's thesis

Aalto University
School of Business

Author Tero Tiitu

Title of thesis Abnormal returns in an efficient market? Statistical and economic weak form efficiency of online sports betting in European soccer

Degree Master of Science (M.Sc.)

Degree programme Finance

Thesis advisor(s) Vesa Puttonen

Year of approval 2016 **Number of pages** 122 **Language** English

Abstract

Online sports betting markets have expanded significantly during the recent decade. At the same time, due to their apparent similarities with traditional financial markets, the academic literature has perceived sports betting markets as a suitable empirical setting for different tests of efficiency.

Employing a data set more extensive than in any previous relevant study, this thesis investigates weak form efficiency of online sports betting in European soccer between seasons 2009 and 2014 from two perspectives: statistical and economic. The statistical tests examine whether subjective probability indicated by average market odds is an unbiased estimator of objective outcome probability. Economic tests, representing the stricter tests of efficiency, inspect whether any betting strategy yields positive returns, utilizing highest odds quoted in the market.

The statistical tests find clear evidence of a persistent favorite-longshot bias and hence of statistical weak form betting market inefficiency. On the aggregate level, both the linear and logit regression models reveal that as the subjective probability of an outcome increases, the objective probability increases more than implied by market efficiency. On the individual odds level, the technique that sorts odds into groups based on subjective probability discovers that the deviations between subjective and objective probability occur on both perimeters of the odds spectrum, not in the middle of it. Thus, in a statistical sense, the betting market appears to be weak form inefficient at high respective low probabilities, while being efficient in between.

The economic tests give a more ambiguous view on weak form efficiency. When considering all the matches in the sample, neither the tobit regression model nor any of the naïve strategies show chances for profitable betting, while some of the odds groups with highest subjective probability demonstrate moderately positive returns. When only including matches with a positive expected value and simulating the associated returns, the sophisticated strategy based on quasi-arbitrage yields no profits, but the strategy that takes into account the favorite-longshot bias generates consistent profits. Therefore, the betting market seems to be inefficient also in an economic sense.

The study concludes that the European online sports betting market in soccer is weak form inefficient. The results are consistent with the earlier literature in terms of statistical efficiency but not in terms of economic efficiency. Due to the lower bookmaker margins in the current market, it is shown that well known statistical biases now also lead to economic inefficiency. The study provides two explanations for the persistence of these inefficiencies: institutional arrangements and market immaturity. Since sports betting markets are not yet as sophisticated as many financial markets, they provide attractive investment opportunities for sharp bettors.

Keywords Efficient market hypothesis, sports betting markets, soccer, betting strategy, favorite-longshot bias

Aalto-yliopisto
Kauppakorkeakoulu

Aalto-yliopisto, PL 11000, 00076 AALTO
www.aalto.fi

Maisterintutkinnon tutkielman tiivistelmä

Tekijä Tero Tiitu

Työn nimi Abnormal returns in an efficient market? Statistical and economic weak form efficiency of online sports betting in European soccer

Tutkinto Kauppatieteiden maisteri (KTM)

Koulutusohjelma Rahoitus

Työn ohjaaja(t) Vesa Puttonen

Hyväksymisvuosi 2016 **Sivumäärä** 122 **Kieli** englanti

Tiivistelmä

Onlinevedonlyöntimarkkinat ovat laajentuneet huomattavasti viimeisen vuosikymmenen aikana. Koska vedonlyöntimarkkinoilla on runsaasti yhtäläisyyksiä rahoitusmarkkinoiden kanssa, vedonlyöntimarkkinat on myös tunnistettu kirjallisuudessa sopiviksi empiirisiksi puitteiksi tutkia markkinoiden tehokkuutta.

Tämä tutkielma perehtyy Euroopassa pelattavan jalkapallon onlinevedonlyöntimarkkinoiden heikon muodon tehokkuuteen kausina 2009–2014 käyttämällä aineistoa, joka on laajempi kuin missään aiemmassa vastaavassa tutkimuksessa. Tehokkuutta tutkitaan sekä tilastollisesta että taloudellisesta näkökulmasta. Tilastollinen näkökulma testaa, vastaako markkinoilla noteerattujen kertoimien keskiarvoon perustuva subjektiivinen todennäköisyys toteutuneisiin ottelutuloksiin pohjautuvaa objektiivista todennäköisyyttä. Taloudellinen näkökulma, joka edustaa tiukempaa tehokkuuden testiä, etsii voittoisia vedonlyöntistrategioita hyödyntämällä korkeimpia saatavilla olevia markkinakertoimia.

Tilastollisen tehokkuuden testit paljastavat pysyvän suosikki-altavastaaja-harhan ja täten selviä epätehokkuuden merkkejä. Kertoimien kokonaistarkastelussa sekä lineaariset että logit-regressiomallit osoittavat, että subjektiivisen todennäköisyyden noustessa objektiivinen todennäköisyys nousee enemmän kuin tehokkailla markkinoilla. Tarkastellessa erikorkuisia kertoimia erikseen havaitaan, että poikkeamat subjektiivisten ja objektiivisten todennäköisyyksien välillä esiintyvät kerroinalueen ääripäissä, eivät sen keskellä. Markkinat näyttäisivät siis olevan epätehokkaat korkeimpien ja matalampien todennäköisyyksien kohdalla ja tehokkaat näiden ääripäiden välillä.

Taloudellisen tehokkuuden testit antavat tehokkuudesta monitulkintaisemman kuvan. Mahdollisuuksia voittoisaan vedonlyöntiin ei löydy tarkastellessa kaikkia aineiston otteluita tobit-regressiomallilla ja naiiveilla strategioilla, mutta jotkin korkeimpien subjektiivisten todennäköisyyksien kerroinryhmät tuottavat pieniä voittoja. Kun hyödynnetään ainoastaan positiivisen odotusarvon otteluita ja simuloidaan näiden tuottoja, kvasi-arbitraasiin perustuva strategia on pääosin tappiollinen, mutta suosikki-altavastaaja-harhaa hyödyntävä sofistikoitunut strategia systemaattisesti voitollinen. Markkinat osoittautuvat näin ollen epätehokkaiksi myös taloudellisessa mielessä.

Tutkielman perusteella Euroopan onlinevedonlyöntimarkkinat jalkapallossa ovat heikossa muodossa epätehokkaat. Tulokset vastaavat aiempia tutkimuksia tilastollisen tehokkuuden osalta mutta eivät taloudellisen tehokkuuden osalta. Tutkimuksessa osoitetaan, kuinka aiemmin tunnistetut tilastolliset epätehokkuudet johtavat nykyisillä markkinoilla myös taloudellisiin epätehokkuuksiin aikaisempaa pienempien vedonvälittäjien marginaalien ansiosta. Tutkielma tarjoaa havaittujen epätehokkuuksien pysyvyydelle kaksi syytä: institutionaalisen asetelman ja markkinoiden kypsymättömyyden. Koska urheiluvedonlyöntimarkkinat eivät vielä ole yhtä kehittyneet kuin useat rahoitusmarkkinat, ne tarjoavat houkuttelevia sijoitusmahdollisuuksia järkiperäisille vedonlyöjille.

Avainsanat Tehokkaiden markkinoiden hypoteesi, urheiluvedonlyöntimarkkinat, jalkapallo, vedonlyöntistrategia, suosikki-altavastaaja-harha

“Chance always favors the prepared mind.”

– Louis Pasteur

Table of contents

List of tables

List of figures

1. Introduction

Online sports betting markets have experienced considerable growth during the last decade, not least due to widespread access to broadband, liberalized regulation of internet gambling, as well as the emergence of online betting channels and product innovations (see, e.g., Gainsbury, 2012; Williams et al., 2012). Anyone can these days bet on a huge variety of sporting events and different outcomes within them, both before and during the events. Sports betting is rapidly turning from being a monopoly-based governmental affair into a competitive global industry with an ever-increasing number of participants, both bookmakers and betting exchange platforms, providing attractive prices to consumers who place their bets rather online than over the counter in a high street shop. Despite these changes, the fundamental difference between sports betting and other forms of gambling still remains: as there are no pure objective probabilities in sports betting, it is possible to earn abnormal returns by deriving the outcome probabilities more sharply than the rest of the market.

At the same time, different betting markets have received lots of attention in the academic literature due to their similarities with financial markets (see, e.g., Pankoff, 1968; Thaler and Ziemba, 1988; Sauer, 1998; Avery and Chevalier, 1999; Durham et al., 2005). Most importantly, it is argued that betting markets are better suited to efficiency tests than traditional financial markets because they encompass several advantages as an empirical setting, while also having commonalities that provide insights into traditional markets. Because payoffs of bets are contingent on the occurrence of uncertain sporting events, sports betting markets generate information about the likelihood of these events in the same way that financial markets aggregate information about uncertain future asset payoffs. Moreover, researchers in finance have been interested in the behavioral motives that make people participate in gambling-related unprofitable investments at the same time that they buy insurance. As theories of decision making try to explain how individuals behave under risk and uncertainty, basing the analysis on probabilities and outcomes, decisions about betting can actually be seen as simplified versions of any decisions.

Similar to conventional financial markets, the efficient market hypothesis (EMH) with its three different forms applies also to sports betting markets. Given an information subset specified by the EMH, enquiries of market efficiency in the context of sports betting can be carried out through two different lenses: statistical and economic (see, e.g., Dowie, 1976; Gandar et al., 1988; Gray and Gray, 1997). The statistical lens inspects whether betting odds, which represent subjective probabilities assigned by market participants for different outcomes, are unbiased estimators of outcomes of sporting events. The economic lens examines the existence of profitable betting strategies. Because of the bookmaker margin, i.e. transaction costs included in betting odds, statistical efficiency does not necessarily translate into profits and is therefore mostly of academic interest. Economic efficiency, on the other hand, is full also of practical relevance and is often considered as the stricter and definite test of efficiency; the question whether a (statistical) bias is (economically) exploitable lies at the heart of the definition of betting market efficiency (Pope and Peel, 1989).

1.1. Purpose and contribution

The purpose of this study is to investigate weak form market efficiency of online sports betting in European soccer from both the statistical and economic perspective. In sports betting, weak form efficiency implies from the statistical perspective that betting odds are unbiased estimators of different outcomes when using only historical betting odds information. From the economic perspective, weak form efficiency means that a bettor cannot earn profits by using the same information. Using both the highest and average odds quoted in the online betting market for 95,789 soccer matches in 74 divisions in 45 European countries between 2009 and 2014, I examine the efficiency of the popular 1X2 betting with three possible match outcomes: a home win, a draw, and an away win.

The statistical efficiency part of the study performs three different tests. First, the relation between subjective (market consensus) and objective (actual historical) probabilities of home wins, draws, and away wins is modeled with linear regression. Second, the same relation is modeled with both binary and multinomial logit regression. Third, I will move from inspecting statistical efficiency on the aggregate level to investigating it on the individual odds level, grouping the average betting odds and testing the probabilities implied by these odds against the groups' empirical probabilities based on historical match results.

The economic efficiency part, employing the highest odds available in the market instead of the average odds used in the statistical part, is comprised of four sections. First, on the aggregate level, net returns of unit bets are regressed on the corresponding odds applying tobit regression, inspecting whether these returns are non-positive across all odds levels. Second, odds are grouped on the basis of odds level and rates of return of each group are tested against the profitability threshold. Third and fourth, profitability of some naïve and sophisticated value betting strategies is explored with the sample, respectively. The idea with the sophisticated strategies is that economic inefficiencies might stem from two alternative sources, either from statistical inefficiencies or from taking advantage of outlier odds even when the market is statistically efficient on aggregate.[1]

The study advances the literature that examines weak form betting market efficiency in three directions. First, to the best of my knowledge, the thesis uses a data set more comprehensive than in any relevant previous soccer study, both in terms of the number of matches considered and the number of bookmakers whose odds are taken into account when determining the average and highest odds for each outcome. When testing betting market efficiency, a large sample size is important for a number of reasons. If a betting market is only marginally inefficient, the sample size must be very large before the tests have reasonable power (see, e.g., Gandar et al., 1988; Golec and Tamarkin, 1991; Gray and Gray, 1997; Nyberg, 2014). A larger amount of data also allows us to perform more accurate tests on any given odds level. Even though the aggregate tests of efficiency might show a betting market to be inefficient, the deviations from efficiency often occur at some specific parts of the odds spectrum, which can be revealed only by the odds level analysis. Moreover, when building potentially profitable betting strategies, a greater number of bets with a positive expected value enables more precise specifications as well as more consistency in returns. Finally, when here covering an outstanding variety of European soccer divisions, we can draw conclusions that hold in the whole 1X2 betting market in European soccer.

As the second empirical contribution, the study comprises one of the first extensive examinations of 1X2 soccer betting during the online betting era, which is characterized by a greater number of bookmakers and other market participants, lower bookmaker margins, and intensive competition (Gainsbury, 2012; Williams et al., 2012). This contrasts starkly to the contexts in the majority of earlier studies of betting market efficiency, in which statistical biases, if found, have usually not resulted in economic inefficiency because of higher bookmaker margins (see, e.g., Pope and Peel, 1989; Cain et al., 2000; Deschamps and Gergaud, 2007). Third, as the study separates the statistical and economic tests of efficiency, it is able to search for value bets, i.e. sources of economic inefficiency, both from potential

[1] The latter scenario is actually behind the provocative title of this thesis. If the difference between the outlier odds and the (statistically efficient) market consensus is large enough, it might be possible to earn an abnormal (economic) return in an (statistically) efficient market.

statistical inefficiencies and quasi-arbitrage[2] opportunities. In this respect, the study differs from the majority of earlier literature that focus either on the statistical or economic dimension of betting market efficiency.

1.2. Motivation

The motivation for this study is also threefold. The first motive stems from the widely accepted fundamental view that sports betting markets resemble simple financial markets that provide ample opportunities for economic analysis not available in other financial markets (see, e.g., Pankoff, 1968; Dowie, 1976; Gandar et al., 1988; Thaler and Ziemba, 1988; Gray and Gray, 1997; Sauer, 1998; Avery and Chevalier, 1999; Durham et al., 2005; Stekler et al., 2010). Investments in sports betting markets have always a well defined start and end point to study as well as quick processing of returns, which allows for the examination of the efficient market hypothesis in an uncomplicated environment. Moreover, due to the large amount of data readily available, it is possible to obtain robust tests of various hypotheses. While a vast amount of studies explore efficiency for instance in stock markets—whether share prices represent unbiased views of true values of companies—the true value of a particular investment is never revealed because these markets are constantly forming expectations of future cash flows. Sports betting markets represent an appealing empirical laboratory to fill this gap and shed light on pure market efficiency.

The second source of motivation relates to the recent developments in online sports betting markets in relation to the results in the earlier betting market literature, which make these markets an even more relevant arena for testing market efficiency from both the statistical and economic perspective. While most of the academic research has focused on single betting markets and on a limited number of traditional bookmakers, it is interesting to perform similar studies in the online betting era, under very different circumstances than during the publication of the first betting market studies in the latter half of the 20th century (Vlastakis et al., 2009; Oikonomidis and Johnson, 2011, 204). Online betting comprises now a highly competitive market based on sophisticated technology and increasingly informed and demanding customers, having a substantial impact on society (Gainsbury, 2012, 1). Even more importantly, as bookmaker margins get smaller, systematic biases in odds, if they still

[2] Originally defined by Vaughan Williams (2001), quasi-arbitrage refers to situations in which a bettor can bet at odds better than implied by objective probabilities, assuming that the average market odds stand for a good approximation of the objective probabilities.

exist, might now be more easily economically exploited (Deschamps and Gergaud, 2007; Graham and Stott, 2008). In addition, at present bookmakers accept also single bets, i.e. bets regarding only one outcome, which gives a mathematical advantage to the bettor compared to combination bets that were the norm previously (Buchdal, 2003, 24; Goddard and Asimakopoulos, 2004; Forrest et al., 2005). All these changes make sports betting more attractive from the bettor point of view.

Regarding weak form efficiency in 1X2 soccer betting, a notable part of the previous literature rejects statistical efficiency in some form, often in the shape of a favorite-longshot bias, without being able to reject economic efficiency due to higher bookmaker margins (Pope and Peel, 1989; Paton and Vaughan Williams, 1998; Cain et al., 2000, 2003; Deschamps and Gergaud, 2007; Vlastakis et al., 2009). On the other hand, some more recent studies detect statistical inefficiency but do not investigate economic efficiency at all (Strumbelj and Sikonja, 2010; Koning, 2012; Nyberg, 2014). This study aims at filling these gaps in the literature. With respect to the studies that reject statistical but not economic efficiency, this study introduces fresh and extensive data in the online betting era, employing substantially more competitive odds. If statistical inefficiency still exists, I will examine whether this inefficiency could now also lead to economic inefficiency. If statistical inefficiency does not any longer exist, quasi-arbitrage might still result in economic inefficiency. With respect to the studies that reject statistical efficiency but do not investigate economic efficiency, this study extends the literature by performing also various tests of economic efficiency.

The third and final motive for the study concerns weak form betting market efficiency as the choice of study, instead of the semi-strong and strong forms in which the information subset under consideration is widened to publicly available information and all potential (even private) information, respectively. With connection to semi-strong form efficiency, various models for predicting soccer match outcomes have already existed for a long time both within the academia and in commercial markets, with an abundance of computer programs and expert knowledge (Rue and Salvesen, 2000). Earlier studies show that the forecasting performance of bookmakers has improved towards the online betting era with intensified competition and higher financial stakes, which might make the discovery of forecasts that would consistently beat the consensus estimates highly unlikely (see, e.g., Rue and Salvesen, 2000; Levitt, 2004; Forrest et al., 2005).

Strong form efficiency, on the other hand, has become a less interesting theme for research due to the unanimous evidence of the low capacity of insiders to provide any significant added value to outcome predictions, as well as of the absence of any considerable insider trading in the main betting markets (see, e.g., Pope and Peel, 1989; Avery and Chevalier, 1999; Forrest and Simmons, 2000; Cain et al., 2003; Song et al., 2007). Thus, it is more captivating to take the weak form efficiency point of view and explore on the one hand whether the market as a whole can build unbiased estimates of match outcomes and on the other hand whether some relatively simple betting strategies can lead to profits when using the highest odds available in the market.

In addition to the above sources of motivation, I find this study highly relevant also personally. As a student of finance, I have always been interested in financial markets in general and market efficiency in particular. On the other hand, as a professional athlete, I have been fascinated by various phenomena around sports, definitely not least by sports betting. Hence, this study gives me a brilliant opportunity to combine these interests around finance and sports. I would be more than glad to contribute both to the academia, by giving new insights into market efficiency, and to the practice of making money, by showing how betting market inefficiencies can be profitably exploited.

1.3. Main findings

The statistical tests find clear evidence of statistical weak form betting market inefficiency through a favorite-longshot bias (FLB), a systematic tendency of the betting public to overbet longshots and underbet favorites. The discovery of the FLB is in line with several studies in soccer (Pope and Peel, 1989; Paton and Vaughan Williams, 1998; Cain et al., 2000, 2003; Deschamps and Gergaud, 2007; Vlastakis et al., 2009; Koning, 2012; Direr, 2013; Nyberg, 2014) but also in other sports (see, e.g., Griffith, 1949; Dowie, 1976; Ali, 1977; Snyder, 1978; Asch et al., 1982, 1984; Vaughan Williams and Paton, 1997; Cain et al., 2003; Snowberg and Wolfers, 2010). On the aggregate level, both the linear and logit regression models reveal that subjective probability indicated by average market odds is not an unbiased estimator of objective outcome probability. As the subjective probability of an outcome increases, the objective probability increases more than implied by efficiency. On the individual odds level, the method that sorts odds into groups discovers that the deviations between subjective and objective probability occur on both perimeters of the odds spectrum, not in the middle of it. Thus, in a statistical sense, the betting market seems to be weak form inefficient at high respective low probabilities and efficient in between.

The economic tests give a more ambiguous view on weak form efficiency. When employing all the matches in the sample, neither the tobit regression model nor any of the naïve strategies reveal profitable betting opportunities, but some of the odds groups with highest subjective probability yield moderately positive returns. When only including bets with a positive expected value and simulating the associated returns, the sophisticated strategy based on quasi-arbitrage generates no profits, but the sophisticated strategy that takes into account the FLB generates consistent profits. For the whole sample period, a strategy specification that involves conservative staking and places bets only on strong favorites yields an annual return of 8%. Therefore, contrary to the majority of the literature that reveals statistical but not economic weak form inefficiency in 1X2 soccer betting (Pope and Peel, 1989; Paton and Vaughan Williams, 1998; Cain et al., 2000, 2003; Deschamps and Gergaud, 2007; Vlastakis et al., 2009), this study suggests that the current betting market is inefficient also in an economic sense. With relation to the more recent soccer literature that traces statistical inefficiency but does not investigate economic efficiency at all (Strumbelj and Sikonja, 2010; Koning, 2012; Nyberg, 2014), this study finds similar biases but also shows that these biases lead also to economic inefficiency.

Altogether, the study concludes that the European online sports betting market in soccer is weak form inefficient. The observed statistical biases appear to be persistent and similar to those discovered before the online betting era. Due to the lower bookmaker margins in the present-day online betting market, however, the distinction to the previous academic work is that these biases can now be economically exploited, providing attractive investment opportunities for sharp bettors. The study provides two intertwining streams of explanations for the peculiarity of why sharp bettors do not seem to exploit the irrationality similarly to what sharp investors and arbitrageurs would do in financial markets to eliminate mispricing. First, institutional arrangements might prevent sharp bettors from exploiting the value bets available in the market. Second, market immaturity highlights the presumption that a significant part of the bettor population may have less financial motives as well as smaller stakes than investors in financial markets, which might lead to suboptimal betting decisions that are taken into account in odds set by profit maximizing bookmakers. All things considered, the focal implication of this study for finance is to demonstrate how even a competitive market, characterized by an ocean of market participants, considerable volumes, and low transaction costs, can be weak form inefficient under some specific conditions.

1.4. Limitations

There are two main limitations to the study, relating to distribution of the bookmaker margin and prevalence of bookmaker stake limits, respectively. First, fully consistent with the practice of the discipline, the study assumes that the bookmaker margin is spread proportionally across each possible outcome of any match in the sample. This is a necessary condition for being able to calculate the subjective probabilities assigned by the market. If, however, the margin would not be dispersed in the hypothetical way, assessments of statistical betting market efficiency would become invalid because we could not determine the subjective probabilities accurately. In the context of this study, however, this limitation is less significant because bookmaker margins in the current online betting market are smaller than before, reducing their overall effect in probability calculations. Furthermore, the presumption that the subjective probabilities are correct is supported by the economic tests of efficiency that detect a similar bias (FLB) without needing to convert odds to probabilities in any way.

Second, when executing the economic tests of efficiency, it is assumed that the bettor is able to bet any given sum at the highest odds quoted in the market. As some bookmakers are known (and always allowed) to limit the stakes of sharp bettors, it might be that the given strategies could not be consistently implemented. In any case, this limitation is probably diminished due to three factors. On the one hand, the current online betting market includes various bookmakers who publically accept sharp bettors. On the other hand, sharp bettors can also use some specific techniques that make them look less sharp from the bookmaker point of view. Moreover, compared to arbitrage betting, the risk of becoming limited is less significant in value betting, which is the foundation of this study.

1.5. Structure of the thesis

The thesis is organized as follows. The subsequent two chapters provide the necessary introduction to online sports betting and testing efficiency within them, while Chapter 4 sieves through the relevant literature on weak form betting market efficiency. Chapters 5, 6, 7, and 8 include the hypotheses, data, methodology, and assumptions employed in the study, respectively. The results of the statistical and economic tests of efficiency are presented in Chapter 9 and discussed in Chapter 10. Finally, Chapter 11 concludes and Chapter 12 shortly highlights ideas on future research.

2. Fundamentals of online sports betting

Gambling refers to wagering a stake with monetary value in games in which the outcome is determined at least partly by chance. The main categories of gambling products are sports betting, casino games, poker, bingo, skill games, and lotteries. Sports betting, the focus of this study, can be described as speculative scenarios about sporting events regarding which participants exchange assets, known as bets or wagers, concerning a specific outcome (Alberola and Garcia-Fornes, 2013).[3] Sports betting distinguishes itself from other gambling products in one fundamental respect. Whereas in other gambling products probabilities of winning are known with certainty and the law of large numbers dictates profits for the organizers because odds are stacked in favor of them, market participants in sports betting can deduce only their subjective probabilities (Levitt, 2004). In other words, while other gambling products involve risk but no uncertainty, sports betting contains both risk and uncertainty (Figlewski, 1979). Consequently, in sports betting, bettors have a chance to be profitable (and organizers of betting unprofitable) in the long run if they are able to determine probabilities of sporting events more accurately than the other market participants.

2.1. Why do people participate in sports betting?

The question of why people gamble—irrespective of the negative expected return characteristic to gambling products—has been generally given two different explanations: the wealth maximizing explanation and the pleasure seeking explanation. The former views gamblers as economic agents who invest money in gambling purely aiming at maximizing wealth. Here, these agents are assumed to make decisions under risk and uncertainty within the expected utility theory (EUT) framework, originally proposed by Bernoulli (1738). After Neumann and Morgenstern (1944) developed the EUT, the wealth maximizing explanation for sports betting has dominated the research agenda. In horse racing, for example, betting behaviors are modelled and explained using the EUT already by Griffith (1949), McGlothlin (1956), Rosett (1965), Weitzman (1965), Harville (1973), Ali (1977; 1979), Snyder (1978),

[3] Here, sports betting is assumed to include horse racing as well as the rapidly expanding variety of other events on which bets are nowadays available. As an example of the blurring continuum between traditional gambling and speculation in financial markets, some well established betting companies have even financial indices betting (concerning stock markets, currencies, commodities, etc.) in their repertoire.

Losey and Talbott (1980), Hausch et al. (1981), and Asch et al. (1982, 1984). These studies consider bettors risk seekers with a convex utility function.

However, under the EUT, the assumption of homogeneity of bettors fails to account for the heterogeneity of motivations that might be present among them (see, e.g., Thaler and Ziemba, 1988). The latter explanation, which is often ignored in the economic literature, argues that betting is at least partly motivated by its pleasure of participation instead of being purely wealth oriented. Given that betting is a negative sum game, Busche and Hall (1988) ponder whether risk preferers anyway are the only people willing to bet and whether bettors are homogeneous with respect to risk attitudes. They suggest that betting could be considered as a hobby for a representative consumer, just like going to the opera or owning a boat. In this case, pleasure of participation would enter the utility function directly; even though it would still be possible that the utility function is convex in wealth, it is not any more a requirement for such rational bettors.

Conlisk (1993) also discovers that when resisting the hypothesis that betting confers utility from pleasure of participation, standard economic theory is left without a good explanation for betting. He argues that because risk aversion is at the heart of explanations for common economic behavior such as insurance purchase and portfolio diversification—whereas strict risk aversion rules out betting completely—it is difficult to explain why people simultaneously pay to decrease risk (insurance) and pay to increase risk (betting).[4] Conlisk claims that as economists do not model food preferences purely in terms of nutritional consequences for health, they should similarly not model betting solely in terms of consequences for wealth. Through appending a tiny utility of betting to the expected utility model for a risk averse individual, he is able to explain both small payoff gambles and large price lotteries, as well as many patterns of risk seeking in the experimental evidence that is puzzling from the standard theory point of view. Thus, when viewing betting as entertainment or consumption, its negative expected value can be seen as a fee for the pleasure it brings.[5]

[4] In their landmark paper, Friedman and Savage (1948) explain simultaneous purchase of insurance and lottery within the expected utility theory by supposing a special shape for an individual's utility-of-wealth function that is risk averting over high and low wealth ranges and risk seeking in between. Markowitz (1952) then modifies this theory by assuming that utility is a function of change in wealth rather than level of wealth and presumes a special shape for this function. The Markowitz hypothesis has also been basic to psychologists who study risky choice, like Kahneman and Tversky (1979). In any case, Conlisk (1993) emphasizes that all these approaches that view betting solely as a means to increase wealth include fundamental difficulties and can explain only very limited patterns of simultaneous purchase of insurance and bets.

[5] Some studies show how the demand for betting can be treated as an economic good, being subject to a variety determinants consistent with the standard consumer preference theory. Gruen (1976) and Ali and Thalheimer (1997), for example, show that betting in horse racing has a downward sloping demand curve and that these bettors behave in a rational, price elastic, and utility maximizing fashion, similar to other commodities.

Individual motivations underlying betting behavior have also been portrayed in other social sciences. Bruce and Johnson (1992) examine two groups of bettors: sophisticated and unsophisticated. They argue that sophisticated bettors place bets to increase their wealth but also because betting provides an intellectual challenge. This approach views betting as professional, following strict guidelines and systematic rules with the intention of maximizing financial gains. Unsophisticated bettors, in turn, strive for excitement and social interaction, sharing a common enthusiasm regarding the build-up to the betting events as well as the drama of the events themselves. As a result, these bettors enjoy the social aspects around the successes and losses of their bets, without putting too much emphasis on financial gains.

Combining the above two perspectives to discover why people participate in sports betting, we can say that incorporating the utility of gambling into the framework of expected utility theory is necessary, but that the magnitude of this utility might vary significantly between different individuals. For some, increasing wealth is the sole motivation for betting; for others, betting itself has its utility. In any case, this thesis views sports betting as a means to maximize wealth and explores the topic from the sophisticated bettor point of view. However, the study does not ignore the existence of unsophisticated bettors either. The proportion of "incompetent" betting volume generated by unsophisticated bettors has been shown to be significant and has a tendency to skew the odds, which, in turn, could possibly be exploited by the "competent money" of sophisticated bettors (Gandar et al., 1988; Paul and Weinbach, 2005b; Kotlyar and Smyrnova, 2012).

2.2. Characteristics of online sports betting markets

Online gambling refers to gambling that occurs for example via internet or mobile phone. Since its inception in 1995, online gambling has burgeoned; all forms of gambling are now widely available online 24 hours a day, seven days a week, from almost any location.[6] During the recent decade, the number of online gambling sites has fluctuated between 2,000 and 2,500. Online sports betting, in turn, has become the dominant sector in the online gambling industry, accounting for roughly a half of the industry revenue. It also represents the leading gambling product in terms of gross gambling yield[7]. The popularity of sports betting

[6] See Gainsbury (2012, 3) for an exhaustive list of the key trends that drive the online gambling industry.
[7] Gross gambling yield (GGY) equals the total amount gambled less prizes paid to customers. Without other income streams than gambling, GGY equals an operator's gross profit. Thus, GGY represents the real economic value of the gambling industry when compared to the rest of the economy. (Gainsbury, 2012, 7)

is mainly due to its obvious and close connections to sports as well as to the fact that betting on sports constantly challenges bettors’ opinions against odds. Moreover, sports betting is regarded as more acceptable by governments than many other forms of gambling, thereby being more permitted in certain markets. (Williams et al., 2012)

Along with the expansion and intensified competition in online betting markets, operators have been forced to offer their products in more attractive terms, corresponding to low margins and higher odds. Even though statistical efficiency of online sports betting markets is interesting to investigate also in the context of higher margins, examinations of economic efficiency become more alluring the lower the margins are. To generate profits by exploiting statistical biases, odds in the current online betting markets have to deviate less from true probabilities (Graham and Stott, 2008; Malarić et al., 2008). As margins decrease, it becomes more likely that some betting strategies generate positive returns, which contrasts strongly with margins that were greater than 10% until the early 2000s (see, e.g., Pope and Peel, 1989; Kuypers, 2000; Dixon and Pope, 2004; Deschamps and Gergaud, 2007).

2.2.1. Organization of markets

Sports betting markets are organized in three distinct ways. The traditional forms of betting existing both online and offline include the bookmaker system, which is the focus of this thesis, and the pari-mutuel system. Alongside these two traditional settings, exchange betting has emerged as a new market mechanism during the last decade. Next, I will give an overview on these systems.

The bookmaker system

The bookmaker system is the most popular form of organizing sports betting markets. In this system, fixed betting odds are unilaterally determined by bookmakers, i.e. gaming companies employing odds compilers who have special knowledge of specific sports to estimate true probabilities of events. Bookmakers usually publish their odds a few days before the start of an event. Bettors can then choose their bets at these odds, while bookmakers act as market makers and automatically take the opposite position, thereby carrying counterparty risk. Theoretically, bookmakers might accept unlimited betting volume at the odds they publish, but in practice some of them deploy limits for stakes and/or maximum winnings in their terms and conditions as well as in their regular customer analytics processes. In addition, even though bookmakers have the right to adjust their odds after a market has opened, they rarely adjust them significantly. Either way, the size of a bettor's claim is always tied to the initially chosen odds so the size of conceivable cash flows is known at the time of placing the bet. (Dixon and Pope, 2004; Forrest et al., 2005; Vlastakis et al., 2009; Franck et al., 2013)

Bookmakers' business model rests on managing betting probabilities, including a margin in their odds. Through slightly reducing the odds for each outcome in the book, the bookmaker will make a profit as long as the probabilities are managed correctly and/or the book is balanced. In this context, the purpose of academic studies has been to analyze the accuracy of these bookmaker probabilities and to explore profitable betting strategies, thereby drawing conclusions on market efficiency.[8] (Buchdal, 2003, 12)

The pari-mutuel system

In the pari-mutuel system, betting volumes on all possible outcomes of a sporting event are aggregated and then distributed to the winners according to their relative stakes. The odds fluctuate freely and the claim of an individual bettor is not fixed ex-ante but depends on all the incoming betting volumes until the market is closed before the start of the event. Thus, compared to the bookmaker system, conceivable cash flows of a bet are not yet determined at the time of placing the bet. In addition, as bettors are effectively competing against each other and the organizer of pari-mutuel betting takes a margin out of the win pool before delivering payouts to winners, the system is risk-free for the operator, whereas this is not necessarily the case in bookmaker betting. Even though the pari-mutuel system for betting is still common in horse racing, it is becoming a less important mechanism compared to the bookmaker and betting exchange systems. In any case, a significant part of earlier academic research on efficiency of betting markets has focused on pari-mutuel betting.[9] (Franck et al., 2013)

---

[8] Studies that investigate the bookmaker system include, for instance, Gandar et al. (1988), Pope and Peel (1989), Golec and Tamarkin (1991; 1995), Brown and Sauer (1993), Woodland and Woodland (1994; 2001; 2003), Gray and Gray (1997), Avery and Chevalier (1999), Cain et al. (2000), Kuypers (2000), Dixon and Pope (2004), Goddard and Asimakopoulos (2004), Schnytzer and Weinberg (2004; 2008), Durham et al. (2005), Forrest et al. (2005), Paul and Weinbach (2005a; 2005b), Deschamps and Gergaud (2007), Vlastakis et al. (2009), Koning (2012), Constantinou and Fenton (2013), Direr (2013), Krieger et al. (2013), and Nyberg (2014).

[9] These studies include, for example, Rosett (1965), Weitzman (1965), Ali (1977; 1979), Snyder (1978), Figlewski (1979), Losey and Talbott (1980), Hausch et al. (1981), Asch et al. (1982; 1984), Quandt (1986), Asch and Quandt (1987), Thaler and Ziemba (1988), Gabriel and Marsden (1990), Hausch and Ziemba (1990), Swidler and Shaw (1995), Terrell and Farmer (1996), Golec and Tamarkin (1998), and Edelman and O'Brian (2004).

The exchange betting system

Exchange betting has recently emerged as a novel betting market system. Originally inspired by electronic financial exchanges, developments in ICT, and the arrival of online betting, betting exchanges have marked a revolution in the industry. These exchanges are order-driven markets for fixed odds betting where odds are determined in a continuous double auction process that matches demand and supply, allowing bettors to bet directly with each other and thereby disintermediating bookmakers. Whereas in the two traditional betting market settings bettors can only buy bets (i.e. bet on a given outcome to occur), in the exchange setting they can also sell bets (i.e. bet against a given outcome to occur). Betting exchanges themselves do not carry any counterparty risk; they act only as middle men and charge a small commission. Exchange betting has been naturally considered as an attractive area of research during the recent years.[10] Out of all betting exchanges in the world, Betfair is by far the largest one, nowadays processing around seven million trades per day.[11] (Laffey, 2005; Franck et al., 2013; Croxson and Reade, 2014)

2.2.2. Other key characteristics

The other essential features of online sports betting markets concern types of bets offered, sports available for betting, timing of betting (in relation to the start of an event), and types of odds applied. First, in contemporary betting, there are many types of bets, both more and less exotic, most of them being suited to the full range of sports. In soccer, the most popular and common type is 1X2 betting, which will be investigated also in this study. In 1X2 betting, odds are determined for three conceivable outcomes: a home win (1), a draw (X), and an away win (2). Besides 1X2 betting, other popular types of soccer bets include handicap betting and total goals betting. A distinction should also be made between making a single bet and making a multiple bet. In the former, only the chosen outcome must take place for the bet to win; in the latter, all the selections must be successful for the bet to win. However, the larger the number of selections in a multiple bet is, the larger will also be the bookmaker's

[10] Empirical studies on exchange betting have been performed for example by Smith et al. (2006; 2009), Gil and Levitt (2007), Franck et al. (2010; 2013), as well as by Croxson and Reade (2014).
[11] Croxson and Reade (2014) highlight that this number of trades is greater than the number of daily trades on all the European stock exchanges combined. Moreover, they mention that during the last years the search term "betfair" has overtaken "FTSE" in popularity on Google.

expected profit margin. Thus, a sophisticated bettor would always prefer single bets, which is the focus of this study.[12] (Buchdal, 2003, 18–25; Gainsbury, 2012, 20)

Second, bets are nowadays offered for all kinds of sports, on several levels, and for both domestic and international events. This study focuses on soccer, which is the most popular sport in the world (Guttman, 1993, 129; Dunning, 1999, 103). Given the enormous popularity of the sport, combined with the increasing interest in online betting, it is no surprise that the online soccer betting market has also expanded. In most countries, soccer is by far the biggest sport in terms of turnover at online bookmakers (Finnigan and Nordsted, 2010; Oikonomidis and Johnson, 2011, 204) and constitutes the fastest growing gambling market (Constantinou and Fenton, 2013). The current soccer betting market is very large and competitive, companies operating with low margins, making it interesting to explore how odds-based information can be used to increase betting returns (Oikonomidis and Johnson, 2011, 208). Furthermore, European soccer constitutes one of the most liquid and focal betting markets (Vlastakis et al., 2009).

Third, online sports betting involves placing bets both before the start of an event (pre-match betting) and during the event (live betting). After publishing their pre-match odds, bookmakers have the option to change their odds if they need to react to new information concerning the event or to manage their projected liabilities around the event. This study focuses on pre-match betting, employing closing odds, i.e. odds quoted in the market just before the start of an event. Live betting has also grown in popularity during the recent years, now accounting for the major part of bookmakers' gross margins (Church-Sanders, 2011; Croxson and Reade, 2011). Compared to pre-match betting, live betting comprises a more fluctuating environment as different conditions in an ongoing event are changing constantly and sometimes even dramatically (Kotlyar and Smyrnova, 2012).

Fourth and finally, there is a difference between fixed odds and variable odds. Fixed odds imply that the size of the conceivable cash flows is determined by the odds on which the bettor places the bet. Thus, the size of the bettor's claim is tied to the initially chosen odds, not depending on subsequent price changes. Variable odds, in turn, refer to odds that are not fixed at the time of placing the bet. Instead, the size of the bettor's claim is tied to volumes placed on each outcome of the particular event at the time of closing the market. Odds provided by bookmakers and betting exchanges are fixed odds, whereas odds provided in

[12] Before the online betting era, single bets were usually not allowed at all, making bookmaking a lot more advantageous to bookmakers. Along with the emergence of online betting and intensified competition, these betting rules have been relaxed and the availability of single bets has become the norm. (Buchdal, 2003, 24–25; Goddard and Asimakopoulos, 2004; Forrest et al., 2005)

pari-mutuel betting markets are variable odds. As this study focuses on the bookmaker system, the odds discussed throughout the study are fixed odds.[13] (Franck et al., 2013)

2.3. Odds, probabilities, and the bookmaker margin

In the context of sports betting, prices of bets on different outcomes are called odds, which represents the amount that the bookmaker will pay out on a winning bet in relation to the bettor's stake (see, e.g., Buchdal, 2003). Odds, in turn, are inversely related to probabilities associated with particular outcomes; the higher the odds, the lower the implied probability of a certain outcome taking place and vice versa. Depending on the geographical area as well as the betting market in question, there are several conventions for quoting odds, out of which the most significant include decimal odds (that is commonly used in continental Europe), fractional odds, and moneyline odds. These conventions contain the same information and can be easily converted to each other. Decimal odds, which are used in this study, specify how many times the stake the payout will be if winning the bet. For example, if staking one hundred euros on a bet with the odds of 2.50, winning the bet would generate a payout of 250 euros and a net profit of 150 euros, whereas losing the bet means that the stake of one hundred euros would be lost.

In the context of decimal odds, probabilities implied by odds are calculated by taking the inverse of the odds. For any event, however, probabilities offered by a bookmaker for all possible outcomes never sum to one but exceed it, because of the bookmaker margin. This margin represents the bookmaker's brokerage fee and aims at guaranteeing a profit for the bookmaker regardless of the outcome, as long as the bets the bookmaker has received are distributed evenly across the outcomes. Conversely, from the bettor point of view, the bookmaker margin can be considered as a transaction cost for placing a bet.

Formally, when using decimal odds for a sporting event with n mutually exclusive outcomes, the offered probability of the j:th outcome is defined as

$$\theta_j = \frac{1}{\delta_j}, \tag{1}$$

[13] To be exact, the earlier literature sometimes gives another meaning for the term fixed odds. Pope and Peel (1989), for example, view fixed odds as odds that are set a few days before the start of an event, thereafter staying unchanged. However, as mentioned by Buchdal (2003, 15), this view is not relevant in the online betting era because odds provided by internet bookmakers and betting exchanges are often changing constantly to reflect new information as well as to manage projected liabilities.

where θ_j is the offered probability and δ_j is the odds quoted for the outcome. The bookmaker margin can then be obtained by adding up the offered probabilities of all possible outcomes so that

$$M = \sum_{j=1}^{n} \theta_j - 1, \tag{2}$$

where M is the bookmaker margin.

Whereas offered probability is used by the bookmaker to define the odds offered to bettors, subjective probability expresses the bookmaker's view on the true probability of an outcome occurring. To convert offered probabilities to subjective probabilities, offered probabilities must be normalized so that they sum to one, thereby eliminating the bookmaker margin. Thus, the subjective probability of the j:th outcome is defined as

$$\rho_j = \frac{\theta_j}{1+M}, \tag{3}$$

where ρ_j is the subjective probability.[14] Corresponding to the subjective probabilities, we can then calculate the scaled odds, i.e. fair odds that lead to a bookmaker margin of zero, by

$$\vartheta_j = (1+M)\delta_j, \tag{4}$$

where ϑ_j refers to the scaled odds for the outcome j.

While subjective probability can be viewed as the probability set by the market for an outcome to happen, objective probability, on the other hand, can be defined as the proportion of times an outcome takes place when the event is repeated a very large number of times. In the context of sports betting, however, objective probability is never revealed in practice, making the whole concept highly controversial. Neither classical nor statistical probability can be applied in sports betting because each sporting event is unique with its set of factors that are stochastic by their very nature. Consequently, sports betting odds represent only estimations of expected probabilities, which gives rise to the examinations of betting market efficiency in the first place. (Ali, 1977; Kotlyar and Smyrnova, 2012)

[14] This kind of standard normalization assumes that the bookmaker margin is distributed equally over the offered probabilities and the assumption is applied also in this study. Section 8.4 provides a further discussion.

2.4. Making money in online sports betting

Successful betting strategies in bookmaker markets contain two main elements: a method that gains an edge over bookmakers and money management that maximizes profits. Regarding the former, there are three methods that enable a long term mathematical advantage: value betting, sports arbitrage, and matched betting. Regarding the latter, staking strategies fall into four broad categories. This section presents these elements. As my focus is on legal ways of betting, I exclude match fixing[15] and exploitation of bookmakers' obvious errors in odds setting[16].

2.4.1. Value betting

Value betting is a sports betting strategy according to which a bettor seeks an edge through betting only on outcomes in which her notion of the objective probability exceeds the probability implied by the odds, with the bookmaker margin built in (see, e.g., Buchdal, 2003; Malarić et al., 2008; Kotlyar and Smyrnova, 2012). In this respect, it is equivalent to the search for underpriced securities in financial markets. Buchdal (2003, 42–53) views successful betting as a practice of understanding and managing probabilities and describes value betting as the only way to overcome bookmakers' odds, providing an accessible measure of a bettor's expectation to make a profit. If the edge exceeds one, a bet is a value bet and potentially profitable, according to the analysis that went into determining it in the first place. Put formally, a bet on a given outcome j is a value bet if

$$r_j = \frac{\pi_j}{\theta_j} > 1, \tag{5}$$

where r_j is the edge, π_j is the objective probability, and θ_j is the offered probability of the outcome.

[15] In match fixing, organizers of fixed matches (match fixers) provide an individual or a group of contestants in a sporting event with rewards for reducing or altering their effort contribution to generate a desired outcome that the organizers can bet on and thereby earn a profit (see, e.g., Preston and Szymanski, 2003; Caruso, 2009). Match fixing is illegal and considered as one of the most significant problems in present-day sports.

[16] Within the betting community, these obvious errors are known as palpable errors or palps. Placing bets at these odds may easily lead to bookmakers limiting or closing these customers' accounts.

In the literature, the value betting methodology is often employed when obtaining estimates of match outcomes with statistical models, thereafter screening potential value bets using Eq. (5) and replacing objective probability with the probability predicted by the given model.[17] Finally, when a value bet is found, a sharp bettor should always choose the bookmaker with the highest odds for the outcome to maximize value. This study rests also on value betting, assuming that value bets are now more prevalent both in terms of number and magnitude than before the online betting era.

2.4.2. Sports arbitrage

Even though each single bookmaker has a positive profit margin built into its odds, in some cases even negative profit margins might be available in the market on aggregate. Similar to finance, where arbitrage refers to the practice of making a risk-free profit by taking advantage of price differentials between different markets, sports arbitrage refers to carrying out sports betting so that a sure profit is guaranteed by combining odds offered by different bookmakers. Formally, a bet is an arbitrage bet if

$$\frac{1}{\sum_{j=1}^{n}\frac{1}{\delta_j}} - 1 > 0, \tag{6}$$

where n is the number of outcomes and δ_j is the odds quoted for the j:th outcome.

Occurrence of sports arbitrage opportunities depend on two factors: the divergence of odds between different bookmakers and the profit margins of these bookmakers. The larger the divergence of odds and the lower the profit margins, the more arbitrage opportunities there will be. Although arbitrage betting seems to provide risk-free profits at first sight, it includes some practical challenges that have to be carefully taken into account.[18, 19] (Constantinou and Fenton, 2013; Franck et al., 2013)

[17] These studies include, to mention a few, Kuypers (2000), Goddard and Asimakopoulos (2004), Gray et al. (2005), Paton and Vaughan Williams (2005), Deschamps and Gergaud (2007), Malarić et al. (2008), Milliner et al. (2009), Vlastakis et al. (2009), Franck et al. (2010), Hvattum and Arntzen (2010), Sessford and White (2010), and Constantinou et al. (2012; 2013).

[18] For a more detailed introduction to sports arbitrage, see, for instance, Buchdal (2003, 49–52), Laffey (2005), Vlastakis et al. (2009), Banks (2013), Constantinou and Fenton (2013), and Franck et al. (2013).

[19] Earlier empirical studies that deal with sports arbitrage include, for example, Pope and Peel (1989), Dixon and Pope (2004), Edelman and O'Brian (2004), Paton and Vaughan Williams (2005), Gil and Levitt (2007), Vlastakis et al. (2009), Constantinou and Fenton (2013), Constantinou et al. (2013), and Franck et al. (2013).

2.4.3. Matched betting

The third way of making money in online sports betting is about generating profits from bonuses and free bets that bookmakers offer to attract new and maintain existing customers, known as matched betting. It requires an account at a bookmaker that offers these incentives as well as another account at a betting exchange or at another bookmaker. By placing a bet on a particular outcome at the bookmaker offering the incentive, at the same time placing corresponding bets against these outcomes elsewhere, a profit will be generated regardless of the outcome. (Banks, 2013)

2.4.4. Principles of money management

Besides gaining a mathematical advantage, successful sports betting requires optimal staking; a bettor must know how much to invest in each bet with a positive expected value. Proper money management involves a predetermined bankroll set aside for betting as well as a staking plan, but it should also always maximize returns and reduce bankruptcy risk to acceptable levels. The problem is similar though more complex in stock markets; in this sense, a bettor can be viewed as an investor looking for excess risk adjusted return. (Buchdal, 2003, 96; Thorp, 2006, 386)

There are four categories of staking strategies: fixed staking, variable staking, progressive staking, and percentage staking. Fixed staking involves placing bets with the same stake regardless of odds. In variable staking, a bettor varies the size of her bets according to various criteria. For example, she may choose to standardize the amount she wins with every successful bet. Progressive staking refers to either increasing or decreasing the stake after each bet, depending on whether it was won or lost. Lastly, percentage staking standardizes the size of the stake as a percentage of the current betting bank at the time the bet is placed, instead of as a specific proportion of the initial bankroll. Kelly staking is the focal percentage staking plan in sports betting and will be applied also in this study, together with the basic fixed and variable staking plans. (Buchdal, 2003, 96–98)

2.4.5. Kelly staking

Kelly staking, or the Kelly criterion, is a capital growth model that maximizes the expected logarithm of wealth on a bet-by-bet basis, originally proposed by Kelly (1956) and later extended as well as thoroughly proved by Breiman (1961) and Algoet and Cover (1988). In the context of betting, Gramm and Ziemba (2008, 320–321) highlight the three most important properties of the model: it maximizes the asymptotic growth rate of capital, it asymptotically minimizes the expected time to reach a specific level of wealth, and it almost surely outperforms any other essentially different betting strategy in the long term.[20] Also in the context of financial markets, where investors must decide how much capital to invest after identifying a favorable investment opportunity, the methodology is equally important; the Kelly investment strategy has been applied in computations of optimal portfolio weights in multi asset as well as worldwide asset allocation problems (MacLean and Ziemba, 2006).[21]

Buchdal (2003, 155–156) argues that the Kelly staking provides an optimum risk-reward trade-off, allowing the betting bank to grow at the maximum rate for minimum risk, when repeated over a period of time. With Kelly staking, the bettor always bets percentages of the bankroll, so a growing bankroll means larger stakes and vice versa. For a particular size of the bankroll, stake size is dependent on both the odds and the edge of the given bet. Formally, the size of a stake according to the Kelly staking plan for placing a bet on the outcome j, which has been found to be a value bet, is defined as

$$k_j = \frac{r_j - 1}{\delta_j - 1}, \tag{7}$$

where k_j is the size of the Kelly stake as a decimal percentage of the bankroll, r_j is the edge determined by Eq. (5) and δ_j is the odds quoted for the outcome.

The main disadvantage of the Kelly criterion is that stakes suggested by the model may in some cases be very large (and intrinsically undiversified), making it risky in the short term, especially if estimations of the betting edge are somehow wrong. Most of the time bettors who apply the Kelly methodology will eventually increase their wealth, but it does not

[20] For further properties of Kelly staking, see, for example, MacLean et al. (1992; 2010) and Thorp (2006).

[21] More precisely, the Kelly methodology is known to economists and financial theorists as concepts like geometric mean maximizing portfolio strategy, maximizing logarithmic utility, the growth-optimal strategy, and the capital growth criterion (Thorp, 2006, 386).

exclude the possibility of losing one's wealth in the case of specific unsuccessful bets.[22] To avoid this pitfall, however, a bettor can apply fractional Kelly staking that compromises long term growth maximization but increases security. By first computing the Kelly optimal size of a bet and then placing the bet with only a fixed fraction of the computed amount, the bettor is able to combine the goals of capital growth and security as she moves the risk aversion away from zero to a higher level.[23] (MacLean et al., 1992; 2010)

Formally, the fractional Kelly staking plan for placing a value bet on the outcome j is given by

$$f_j = k_j * g, \tag{8}$$

where f_j is the size of the fractional Kelly stake as a decimal percentage of the bankroll, k_j is the size of the standard Kelly stake determined in Eq. (7), and g is the applied fraction so that $0 < g < 1$. One common choice is to apply the half Kelly staking plan, according to which $g = 0.5$, compromising the optimal growth rate by 25% (Thorp, 2006, 411–412).

[22] More precisely, these bets are subject to the Arrow-Pratt risk aversion—the reciprocal of current wealth—that is small compared to other commonly chosen utility functions (MacLean et al., 2010).
[23] In financial markets, many popular investors have applied Kelly or fractional Kelly strategies. In their analyses, MacLean et al. (2010) suggest that Warren Buffett and George Soros seem to act similarly to standard Kelly bettors, while John Maynard Keynes resembles an 80% fractional Kelly bettor.

3. Testing efficiency in sports betting markets

This chapter is concerned with testing efficiency in sports betting markets. The discussion centres par excellence around the efficient market hypothesis and its implications in sports betting. The chapter begins with a concise review of the theory of efficient markets as sketched for financial markets, then moving on to demonstrating the usefulness of sports betting markets as an arena for efficiency tests and discussing the similarities between these two markets. The latter part of the section gives the definitions for different forms of betting market efficiency as well as generally describes the methods for testing efficiency within these forms. The exact methodology applied in this study will be presented later in Chapter 7.

3.1. The efficient market hypothesis revisited

The theory of efficient markets is one of the most important paradigms in economics and finance, referring to the process of price formation in financial markets. While the emergence of the efficiency literature dates all the way back to the 16th century [see, e.g., Sewell (2011) for a chronological review of the notable literature], Bachelier (1900) provides the pioneering theoretical contribution and Cowles (1933) performs the first considerable empirical research. The modern literature around the topic begins with Samuelson’s (1965) proof of price changes that are not forecastable as long as prices fully incorporate the expectations and information of all market participants. However, the paradigm gained wide popularity after Fama (1970) presented the efficient market hypothesis (EMH), in his seminal review of theoretical and empirical studies done in capital markets. In 1978, Jensen (1978, 95) declared his view on the EMH so that “there is no other proposition in economics which has more solid empirical evidence supporting it”. Later, as originally done by Fama, most investigations on efficiency in the real world have focused on conventional financial markets in which large databases have been available.[24]

The efficient market hypothesis proposes that financial markets are informationally efficient, meaning that these markets are perfectly competitive and security prices reflect all

[24] Vaughan Williams (2005), among others, provides a comprehensive review of the academic literature investigating efficiency in financial markets. Osborne (2001) lists the wide variety of other markets in which the EMH has been tested.

available information.[25] If a discrepancy exists between the market price and the information-based fundamental value, informed investors would immediately exploit this deviation. As a result, under the EMH, market prices of financial assets are equivalent to their "right", fundamental values, i.e. the discounted values of the securities' future cash flows. Changes in market prices are random since they are influenced only by new random information, which in turn instantly adjusts the prices to reflect the new situation. Because the flow of information is unpredictable in nature, investors cannot make any estimations of the new information—otherwise predictions regarding new information would have already caused the market prices to change. (Fama, 1970; 1991; 1998; Malkiel, 2003)

Beyond the assumption of utility maximizing market participants, the efficient market hypothesis necessitates that agents have rational expectations on the population level. This implies that in the face of new relevant information agents are able to update their expectations correctly, though only on average. Consequently, the EMH claims that no one can consistently earn above average returns on a risk adjusted basis; neither technical analysis (the study of past prices to predict future prices) nor fundamental analysis (the exploration of financial information to detect misvalued assets) should generate returns greater than those achieved by holding a randomly selected portfolio of assets with comparable risk. Even though the EMH allows market pricing to be imperfect in the short term, it believes that the true values will win out in the long term. (Malkiel, 2003; Gray et al., 2005)

3.1.1. Three forms of market efficiency

Fama (1970) distinguishes three information subsets—weak, semi-strong, and strong—which have become widely accepted forms for appraising and testing market efficiency from different perspectives. Weak form tests determine whether past prices alone can predict future prices. Semi-strong tests use all publicly available information to predict prices. Strong form tests are concerned with demonstrating whether any special group is able to achieve a higher than average rate of return due to the group's monopoly over specific information. Thus, the basis for separation of the three information subsets lies in what is meant by "all available information". Each stronger form of efficiency incorporates all

[25] Literature recognizes two alternative interpretations of the statement that prices reflect all available information. In the strong version, information and trading costs, i.e. costs of getting prices to reflect information, are always zero (Grossman and Stiglitz, 1980). In the weaker and perhaps economically more sensible version, prices reflect information to the point in which the marginal benefits of acting on the information do not exceed the marginal costs (Jensen, 1978).

weaker forms of efficiency. Formally, a market is said to be efficient with respect to some information set if security prices would be unaffected by revealing that information to all participants (Malkiel, 1992).

Weak form efficiency claims that abnormal returns are impossible to obtain by analyzing prices from the past. Future price movements should be determined entirely by information not contained in the price series; asset prices must follow a random walk. If weak form efficiency holds, prices are composed of only three components: the last period's price, the expected return of the asset, and a random error term with an expected value of zero. The random error corresponds to new unexpected information released during the period observed. Weak form efficiency is usually examined through statistical tests of independence (such as autocorrelation tests and runs tests) and trading tests (such as filter rule tests). The majority of academic research supports weak form efficiency of capital markets.[26] (Fama, 1970; 1991; 1998; Malkiel, 2003)

Semi-strong form efficiency implies that all publicly available information, such as facts about firms' products, operations, patents, and balance sheets, are reflected in prices of relevant financial assets fully, immediately, and in an unbiased fashion. Thus, in addition to historical prices, it suggests that excess returns are impossible to generate by basing investment decisions on any publicly available information. Consequently, while weak form efficiency insists that technical analysis is useless in the search for abnormal returns, semi-strong form efficiency dismisses also all forms of fundamental analysis. The most common forms of semi-strong efficiency tests include event studies as well as different time series tests. Academic research shows that financial markets are typically efficient also in the semi-strong sense but this evidence is more blurred than that of weak form efficiency.[27] In any case, particularly event studies comprise the most obvious evidence of efficiency, which is supportive with only few exceptions. (Fama, 1970; 1991; 1998; Malkiel, 2003)

Strong form efficiency, the last information subset of the EMH, assumes that all available information, both public and private, is reflected in asset prices. In this case, excess returns cannot be achieved in the long run even if an investor holds insider information. This form of efficiency is naturally the most contested version of the EMH. In a theoretical sense, it is the most compelling form of the EMH, while in a practical sense it is the most difficult

[26] See Lim and Brooks (2011) for a systematic review of the weak form efficiency literature with an exclusive focus on stock markets.

[27] Studies inconsistent with the semi-strong form efficiency consider, for example, dividend payments (Charest, 1978; Ahrony and Swary, 1980; Asquith and Mullins, 1983) and earnings announcements (Ball and Brown, 1968; Bernard and Thomas, 1990).

form to confirm. Strong form efficiency is often viewed as a benchmark against which deviations from market efficiency in its strictest sense can be judged. Tests of strong form efficiency tend to focus on returns of groups of investors with excess information, including corporate insiders, institutional money managers, analysts, and exchange specialists. The majority of academic evidence refutes strong form efficiency of financial markets.[28] (Fama, 1970; 1991; 1998; Malkiel, 2003)

Altogether, during the last decades, empirical evidence of the different forms of the efficient market hypothesis have raised both support and opposition. After enjoying strong support earlier, the EMH has now come under relentless attack especially from the school of behavioral finance (Lim and Brooks, 2011). To offer reconciliation between the opposing views, Lo (2004) presents the Adaptive Market Hypothesis (AMH), an evolutionary alternative to market efficiency, under which the EMH can coexist along with behavioral finance in an intellectually consistent manner. According to the AMH, instead of being an all or none condition, market efficiency is a characteristic that varies continuously over time and across markets.

3.1.2. The joint hypothesis problem: the main obstacle to analyze market efficiency

As explained above, efficient markets are markets in which prices fully reflect available information within a given information subset (Fama, 1970). However, Fama (1991) later admits that market efficiency in itself is actually not testable because it must be tested jointly with some model of equilibrium, i.e. with an asset pricing model. Thus, even if we find anomalous evidence of the behavior on some asset returns, it will remain unclear whether these anomalies are really due to market inefficiency or, alternatively, an incorrect model of market equilibrium. This challenge, called the joint hypothesis problem (or the bad model problem), comprises the most significant obstacle to make inferences about market efficiency with the EMH, suggesting that regardless of the amount of evidence supporting the EMH it can never be perfectly validated.

This obvious drawback of the efficient market hypothesis has been thereafter widely recognized in the literature, allowing many other authors to conclude that the EMH per se is not a well defined and empirically refutable hypothesis. The common feature of this critique

[28] For example, corporate insiders have been shown to profit from trading on pertinent information (Jaffe, 1974; Seyhun, 1986). However, Rozeff and Zaman (1988) argue that even insiders' abnormal returns are economically insignificant and essentially diminish if transaction costs are higher.

concerns the fact that the focal phrase of the EMH, that prices fully reflect all available information, is a statement about two distinct aspects of prices: the information content and the price formation mechanism (see, e.g., Lo and MacKinlay, 1999; Lofthouse, 2001; Sewell, 2012). In short, despite the vast amount of studies applying the EMH, it does not seem to provide waterproof methodology for analyzing efficiency of financial markets.

3.2. Sports betting markets as a setting for efficiency tests

Due to the challenges in formulating direct tests of efficiency in financial markets, the finance literature has looked beyond conventional financial markets to find a setting that would be better suited for tangible studies of efficiency. Pankoff (1968) is the first to mention that sports betting markets permit an unusually direct test of market efficiency, whereas Thaler and Ziemba (1988) are among the first to propose that sports betting markets might be even better suited for efficiency tests than the stock market and many other financial markets. They emphasize that since a stock is basically infinitely lived, its value today depends both on the present value of future cash flows and on the price that will be paid for the security tomorrow. In contrast, Thaler and Ziemba argue that each asset in sports betting markets has a well defined termination point at which its value becomes certain.[29] Thaler and Ziemba also suggest that sports betting markets have a better chance of being efficient because their conditions, such as quick and repeated feedback, tend to facilitate learning.

Many other authors have later elaborated these views on the suitability of sports betting markets for tests of efficiency and choice under uncertainty, highlighting four distinct features (see, e.g., Gray and Gray, 1997; Sauer, 1998; Avery and Chevalier, 1999; Kuypers, 2000; Schnytzer and Weinberg, 2004, 71–72; Durham et al., 2005; Gray et al., 2005; Paul and Weinbach, 2005b; Croxson and Reade, 2014). First, in sports betting markets, market participants receive an objective signal about the fundamental value of an asset quickly, often even within a couple of hours. Efficiency tests do not have to focus on predictability of asset returns, which simplifies inferences about the learning behavior of investors. Second, the range of possible asset payoffs is simple and often known with certainty in advance. Third, tests of efficiency in sports betting markets reduce the scope of the pricing problem and

[29] To be exact, as shown by Figlewski (1978), even options and futures markets seem to have a similar objective end-of-horizon payoff at first sight. However, as Gandar et al. (1988) remind, the commodities underlying options and futures are anyway typically ongoing in nature and their prices at the option or future horizon could possess some irrationality that would therefore come over to the derivative securities as well.

remove the joint hypothesis problem. Fourth, sports betting markets seem to provide a useful middle ground between stock markets and experimental markets. Compared to the former, sports betting markets form an idealized laboratory setting; compared to the latter, they offer vast volumes of real money on the line over lengthy time series.

Besides praising the suitability of sports betting markets for efficiency tests, the literature also recognizes several similarities between sports betting markets and financial markets (see, e.g., Pankoff, 1968; Snyder, 1978; Asch et al., 1982; 1984; Gabriel and Marsden, 1990; Terrell and Farmer, 1996; Avery and Chevalier, 1999; Levitt, 2004; Durham et al., 2005). Sports betting constitutes an economic market in which people buy and sell assets at prices that reflect bettors' judgments about probabilities of different outcomes, offering an opportunity to study economic decision making under conditions of risk and uncertainty. These markets are characterized by a large number of participants with heterogeneous beliefs, ease of entry, and extensive market knowledge combined with rapid dissemination of information. Moreover, the markets are liquid, encompass large volumes, and involve arbitrageurs. Like trading in financial assets, sports betting is a zero-sum game with one trader on each side of any transaction, and also the possibility and profitability of insider information parallels closely that found in the stock market.[30]

3.3. Devising the efficient betting market hypothesis

Even though the framework of the efficient market hypothesis with its three different forms has been originally developed for financial markets, it applies in a natural way to betting markets where asset prices are replaced by betting odds. First proposed by Dowie (1976), investigations on betting market efficiency can be considered in three information subsets. In his approach, weak form efficiency equates available information with historical odds and returns, semi-strong form efficiency adds public information to the set of available information, and strong form efficiency allows for the existence of particular market participants who may possess monopolistic access to specific information.

On the other hand, in the context of betting markets, efficiency has been viewed through two different lenses: statistical and economic (see, e.g., Snyder, 1978; Losey and Talbott, 1980; Asch et al., 1984; Gandar et al., 1988; Thaler and Ziemba, 1988; Woodland

[30] Furthermore, similarities between specific financial assets and bets have also been discussed in the literature. Ruhm (2003) shows how positions in financial options can be considered as simple bets, while Vecer et al. (2006) compare betting contracts with credit derivatives, viewing goals and red cards in soccer as credit events.

and Woodland, 1994; Gray and Gray, 1997; Sauer, 1998; Sessford and White, 2010). Statistical efficiency implies that betting odds (i.e. probabilities that are determined by the information subset at issue) are unbiased predictors of actual results. Economic efficiency implies that it is not possible to earn profits following any betting strategy. Thus, contrary to financial markets where market efficiency is commonly viewed only through the economic lens, betting markets provide an additional, statistical lens to shed light on efficiency.[31]

Due to the transaction cost incorporated in the bookmaker margin, statistical betting market inefficiency does not necessarily indicate economic inefficiency. While tests of statistical efficiency are primarily of academic interest, tests of economic efficiency represent the stricter and decisive test of efficiency, being full of practical interest (Asch et al., 1984; Gray and Gray, 1997; Schnytzer and Weinberg, 2004, 80). As suggested also by Fama (1991) in the context of financial markets, more important than finding market inefficiencies per se is to measure the extent of inefficiency. To generate profits, the larger the biases the greater can transaction costs be, and vice versa. In the current online betting markets characterized by low transaction costs, deviations from statistical efficiency do not have to be large to allow profitable betting (Deschamps and Gergaud, 2007; Graham and Stott, 2008).[32]

When combining the three information subsets provided by the efficient market hypothesis with the two lenses on betting market efficiency suggested by the earlier literature, I devise the "efficient betting market hypothesis" and give the following definitions for the three forms of betting market efficiency:

- *Statistical weak form betting market efficiency:* Betting odds capture all historical information on odds so that they are unbiased predictors of actual results.
- *Economic weak form betting market efficiency:* Using only historical information on odds, there is no consistently profitable betting strategy.

- *Statistical semi-strong form betting market efficiency:* Betting odds capture all publicly available information so that they are unbiased predictors of actual results.

[31] When using the concepts of statistical and economic efficiency in betting markets, it is important to bear in mind that in this context the word *statistical* refers to probabilities, not to research methodologies. Both statistical and economic efficiency are, naturally, investigated with the help of statistical techniques.

[32] Asch et al. (1984) draw an interesting analogy to empirical research in the stock market. While many departures from perfect efficiency have come out in different empirical tests, in general these departures are often not large enough to devise an active strategy that could consistently beat a passive buy-and-hold strategy.

- *Economic semi-strong form betting market efficiency:* Using all publicly available information, there is no consistently profitable betting strategy.

- *Statistical strong form betting market efficiency:* Betting odds capture all existent information so that they are unbiased predictors of actual results.
- *Economic strong form betting market efficiency:* Using all existent information, there is no consistently profitable betting strategy.

3.4. Methods for testing betting market efficiency

Empirical examination of betting market efficiency has many variations, depending on the form and lens through which efficiency is considered. Each form of efficiency can be examined with either pure statistical tests or direct economic tests; the former tests look at statistical properties of betting markets, while the latter tests attempt to detect unexploited profit opportunities (Gandar et al., 1988). This holds also in soccer; odds should constitute unbiased estimates of actual match outcomes and if there exist systematic biases in the market's ability to incorporate relevant information into odds, one might be able to formulate betting strategies that profitably exploit these biases (Gray et al., 2005). This section briefly describes the variety of methods available for statistical and economic tests of efficiency.

3.4.1. Statistical tests of efficiency

Statistical betting market efficiency requires that market probability distributions (i.e. odds), conditional on the relevant information subset, are equal to objective probability distributions, conditional on the same information set. The exact procedures depend on the institutional structure of the betting market associated with a given sport. Formally, statistical betting market efficiency implies that

$$E\left(\pi_j - \rho_j \middle| \emptyset \right) = 0, \tag{9}$$

where π_j is the objective probability of the outcome j, ρ_j is the subjective probability derived from odds for the outcome and $\emptyset$ is the subset of information at issue. As the betting market contains no systematic risk, we need to consider only first moments of conditional probability distributions. Moreover, a zero expectation of the betting market's conditional forecast error comprises a necessary condition for efficiency. (Gandar et al., 1988; Stekler et al., 2010)

Statistical tests of betting market efficiency can be divided into regression-based tests, grouping tests, and other tests. Regression-based tests use various types of regression analysis, both linear and nonlinear, for modeling the relationship between odds and outcomes. Standard linear regression has been widely employed but it has been considered weak, revealing only whether a betting market is efficient in an aggregate sense (see, e.g., Zuber et al., 1985; Gandar et al., 1988; Sauer et al., 1988; Golec and Tamarkin, 1991; Gray and Gray, 1997; Kuypers, 2000; Paul et al., 2004; Sinkey and Logan, 2014). To reduce aggregation, however, it can be applied to each level of odds individually (Woodland and Woodland, 1994). Besides these standard models, the relationship between subjective and objective probability has been investigated using linear probability models (LPMs), first introduced by Pope and Peel (1989). However, LPMs are subject to two major econometric issues: heteroscedasticity and allowance of unreal probabilities. To overcome these issues, probit and logit models, both binary and multinomial, have also been implemented (see, e.g., Forrest et al., 2005; Franck et al., 2010; Koning, 2012; Nyberg, 2014).

Grouping tests, originally developed for horse racing by Ali (1977), apply two different grouping methods to compare subjective probabilities with objective probabilities. The first method categorizes contestants into groups based on their favorite rank in a given race, while the second method carries out the categorization based on the level of odds of the contestants. Thereafter, in both cases, the groups' average subjective probabilities are calculated and compared to the estimators for the groups' objective probabilities. Even though the grouping tests are applied to a large extent only in pari-mutuel horse and greyhound racing, they contain specific features that are applicable also in other contexts.

The other tests include two approaches to evaluate the effectiveness of betting market odds as forecasts: the Brier score and the rank probability score (RPS), both of which are proper score functions. These tests are descriptive measures on prediction accuracy over time and/or between different bookmakers for mutually exclusive and collectively exhaustive outcomes (Franck et al., 2010; Hvattum and Arntzen, 2010). Developed by Brier (1950), the Brier score treats each of the outcomes in a betting event as a separate binary event and each market odds for these outcomes as an individual forecast. Each score is obtained by calculating the squared difference between the subjective probability and the realized outcome (zero if the outcome did not take place and one if it did) in a given match. Thus, the lower the score, the closer the odds are to the true probabilities. Formally, for a set of forecast-event pairs, the average Brier score is defined as

$$B = \frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{O}(\rho_{ij} - \varphi_{ij})^2, \tag{10}$$

where B refers to the Brier score, N is the number of forecast-event pairs, O is the number of possible outcomes in an event, ρ_{ij} is the subjective probability of the j:th outcome to occur in the i:th event, and φ_{ij} is a binary indicator for the different outcomes so that $\varphi_{ij} = 1$ if the outcome j took place in the i:th event and $\varphi_{ij} = 0$ otherwise.

First introduced by Epstein (1969), the rank probability score is another descriptive measure on odds accuracy, evaluating the forecasts for different outcomes as a multi-category forecast. In the context of soccer, Constantinou and Fenton (2012) argue that Brier scores do not address the ordinality of match outcomes when scoring the forecasts and suggest that the RPS is a more appropriate measure. It is calculated as the sum of differences between the cumulative forecast probability and the cumulative outcome probability. The major difference to Brier scores is that the score given to a forecast with the RPS is weighted with the number of possible outcomes. Hence, the average RPS is formally defined as

$$RPS = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{O-1}\sum_{j=1}^{O}\left(\sum_{k=1}^{j}\rho_{ijk} - \sum_{k=1}^{j}\varphi_{ijk}\right)^2, \tag{11}$$

where RPS refers to the rank probability score. Similar to Brier scores, a lower average RPS indicates higher forecast accuracy and vice versa. Brier scores are thereby actually a special case of the RPS when the number of categories is two.

3.4.2. Economic tests of efficiency

Economic tests of efficiency explore whether any betting strategy yields consistent profits. At the same time, these tests are always incomplete per se because the number of different betting rules and strategies is in practice unlimited. Formally, economic betting market efficiency means that

$$E(s|\,\emptyset\,) \leq 0, \tag{12}$$

where *s* stands for the applied betting strategy. (Tryfos et al., 1984; Gandar et al., 1988)

Similar to financial markets, different betting strategies can be broadly divided into two categories: fundamental and technical strategies (see, e.g., Hausch et al., 1981; Thaler and Ziemba, 1988). Fundamental strategies are commonly based on publicly available information according to which a bettor attempts to detect bets that have greater probabilities of winning than implied by odds quoted in the market. Technical strategies use only current betting data and require less information; here, a bettor attempts to find inefficiencies in the market and to place bets that have a positive expected value. Most of the earlier academic research has concentrated on technical strategies, and as a weak form examination of efficiency the economic tests applied in this study are also based on technical strategies.

There are various methods to test the significance of returns obtained with different betting strategies (see, e.g., Tryfos et al., 1984; Gandar et al., 1988; Paton and Vaughan Williams, 1998; 2005; Sessford and White, 2010). For example, z-values of returns can simply be evaluated against the null hypotheses of randomness and/or unprofitability. Moreover, after calculating the actual return to a unit stake on each outcome, net returns on unit bets can be regressed on odds or, inversely, on implied probability.

4. Earlier studies on weak form betting market efficiency

This chapter focuses on the literature on weak form betting market efficiency. While much of the earlier literature concentrates on racetrack betting, team sports have begun to attract attention more recently. In sports whose betting markets are centred around handicaps, researchers have mostly focused on potential biases in the favorite-longshot and home-away dimensions. On the other hand, in markets where bets are offered on different outcomes at odds varying according to the relative strengths of the participating teams, as is the case in soccer, the literature has concentrated on biases in the short odds versus long odds dimension.

Regarding statistical tests, Pankoff (1968) performs the first ever test of betting market efficiency by regressing match outcomes in NFL on bookmakers' spreads, later followed by Zuber et al. (1985), Gandar et al. (1988), and Dare and MacDonald (1996), all of whom conclude that the market seems to be efficient on aggregate, analogously to the stock market. The first economic tests, also in NFL, are provided by Vergin and Scriabin (1978), Tryfos et al. (1984), Zuber et al. (1985), and Sauer et al. (1988), all of whom are able to demonstrate some profitable betting strategies, though often with much less consistency when applied with other samples. In the following, I will present the later literature most relevant from this study point of view, divided in three parts: evidence of the favorite-longshot bias, results in soccer, and findings in sports arbitrage.

4.1. The favorite-longshot bias

The most prominent bias reported in the betting market literature is unquestionably the favorite-longshot bias (FLB), first discovered by Griffith (1949). The FLB implies that placing bets on favorites yields a higher return than placing bets on longshots; in other words, the betting public has been observed to have a systematic tendency to overbet longshots and underbet favorites. Asch and Quandt (1987, 290) describe it as "the most widely established empirical regularity" in betting. This phenomenon has been identified in a variety of sports.[33] The FLB is most extensively investigated in horse racing, in which it is detected for example by Dowie (1976), Ali (1977), Snyder (1978), Hausch et al. (1981), Asch et al. (1982, 1984),

[33] Rubinstein (1985) detects a similar bias also in the equity options market as he finds that shorter maturity options (i.e. longshots) are overpriced. Moreover, Hodges et al. (2003) demonstrate how the FLB exists when investors pay more for S&P 500 and FTSE 100 put options than they are subsequently worth, while the degree of overpaying increases monotonically as the probability of finishing in the money decreases.

Vaughan Williams and Paton (1997), Jullien and Salanié (2000), and Snowberg and Wolfers (2010). Most often the FLB is found by comparing objective and subjective probabilities and/or economic returns between different odds classes or favorite positions.

Thaler and Ziemba (1988) review the earlier literature on betting markets in horse racing and highlight the consistency of the favorite-longshot bias. At the same time, however, they propose that the betting market is surprisingly efficient in an economic sense; the FLB does not appear to translate into profitable betting strategies. More generally, Thaler and Ziemba argue that modelling betting behavior is complicated because it seems to depend on numerous factors that are existent also in other investment behavior. Professional portfolio managers, for example, might be more concerned with beating some index than with maximizing returns, which might lead to betting on longshots at the end of the day if behind.

In soccer, evidence of the favorite-longshot bias is provided by Pope and Peel (1989), Paton and Vaughan Williams (1998), Cain et al. (2000, 2003), Deschamps and Gergaud (2007), Vlastakis et al. (2009), Koning (2012), Direr (2013), and Nyberg (2014). Besides soccer, Cain et al. (2003) depict that the FLB appears in many sports, including boxing, cricket, greyhound and horse racing, as well as snooker and tennis, even though it remains unclear whether the bias can be viewed as a monotonic relationship throughout the probability spectrum. In most cases, however, regardless of the sport at issue, the studies propose that it is not possible to exploit the bias to devise a profitable betting strategy, despite the variation in rates of return between different bets along the favorite-longshot continuum.

The economics and psychology literature provides several explanations for the favorite-longshot bias without a clear consensus on its causes. In any case, when recharacterizing the FLB as a human tendency to overestimate the probability of unlikely events, Coleman (2004) accentuates the need for a robust explanation for the phenomenon because it might be pervasive throughout finance, management, and society, often causing the return from investment to fall as the probability of loss increases.[34] The different explanations for the FLB can be broadly divided into two categories: demand side explanations and supply side explanations. In the former, the FLB is caused by irrational behavior of bettors; as they place more money on longer than shorter odds, they progressively reduce the return of longshots. In the latter, the FLB is driven by biased odds setting behavior of bookmakers. Next, I will shortly present the main ideas behind these explanations.

[34] As real life examples of this tendency, Coleman (2004) mentions how the return is often lower than expected in investments in developing markets, mergers and acquisitions, mineral exploration, and innovative business models, essentially covering a wide spectrum of institutions from marriage to crime.

4.1.1. Demand side explanations for the FLB

The demand side reasons for the favorite-longshot bias can be further classified into two distinct explanations: bettors either have risk loving utility functions and/or they misperceive probabilities associated with different outcomes. The risk loving class presumes that bettors give up some return in exchange for the additional risk of low probability bets. The view that bettors are purely risk lovers is supported by Weitzman (1965), Ali (1977), Quandt (1986), and Jullien and Salanié (2000), for example. In addition, some other theories are observationally equivalent to the risk loving explanation; these include the existence of bragging rights resulting from eventually winning a bet with long odds (Snyder, 1978; Thaler and Ziemba, 1988), bettors' preference for skewness instead of variance (Golec and Tamarkin, 1998), as well as utility conferred by purchasing a bet with long odds (Conlisk, 1993). Because globally risk loving bettors should not buy insurance—even though they do it in most cases—it is often suggested that bettors are actually locally risk loving but globally risk averse (Rosett, 1965; Quandt, 1986; Golec and Tamarkin, 1995).

In the misperceptions class it is assumed that cognitive errors and misperceptions of probabilities play a bigger role in the FLB than risk love. Many laboratory studies have shown how people overestimate small probabilities and underestimate large probabilities; people are systematically poor at discerning between small and tiny probabilities, while also having a strong preference for certainty over extremely likely outcomes (Snowberg and Wolfers, 2010). These findings, in turn, constitute an important foundation for the prospect theory proposed by Kahneman and Tversky (1979). Moreover, Golec and Tamarkin (1995) examine whether bettors' risk preferences or overconfidence better explains the preference for low probability bets and find that overconfidence fits the data more closely.[35]

4.1.2. Supply side explanations for the FLB

The supply side factors behind the favorite-longshot bias rest on the fact that the phenomenon is more pronounced in bookmaker markets (in which firms accept bets from the public and might therefore carry a risk if their books are not balanced) than in pari-mutuel

[35] In addition to the above explanations, it has been suggested that transaction and information costs might cause the FLB at least partly (Terrell and Farmer, 1996; Paton and Vaughan Williams, 1998; Smith et al., 2006). The reasoning here is that becoming an informed bettor is hindered because of these costs, which slows down the removal of market mispricing and maintains the bias. Here, the extent of the FLB depends on the incidence of these costs on informed bettors; the greater the transaction and information costs, the greater the bias.

markets (in which payoffs represent proportional shares of the total money bet on all outcomes and market makers carry no risk) (Coleman, 2004). At the same time, however, the supply side cannot provide a complete explanation as the FLB is slightly present also in pari-mutuel betting markets. In his discussion on bookmakers' strategies, Henery (1985) concludes that because betting on favorites is most competitive, bookmakers quote higher odds on favorites to attract bettors, while they compensate for the implied reduction in margin by providing less advantageous odds on longshots. Shin (1991), on the other hand, argues that due to the existence of insiders it is optimal from the bookmakers' point of view to employ a square root rule where the ratio of quoted odds is equal to the square root of the ratio of objective probabilities. As a result, betting odds would underestimate the chances of favorites and overestimate the chances of longshots.

4.1.3. Exceptions to the FLB

Despite the wide evidence of the favorite-longshot bias, some authors have also provided exceptions, either in the form of no bias at all or in the form of a reverse favorite-longshot bias (RFLB). Busche and Hall (1988), Busche (1994), and Swidler and Shaw (1995), for example, find no bias in racetracks, whereas Schnytzer and Weinberg (2008) observe no bias in Australian football. More recently, after the emergence of betting exchanges, Smith et al. (2006; 2009) show how the FLB is demonstrably lower in betting exchange markets than in traditional betting markets, most likely due to lower transaction costs. They also suggest that the level of the FLB is inversely related to the amount of information available to bettors.

The reverse favorite-longshot bias means that placing bets on favorites yields a lower return than placing bets on longshots; as subjective probability of an outcome increases, objective probability increases less than implied market efficiency and vice versa. Signs of a RFLB have been detected in American football (Golec and Tamarkin, 1991; Borghesi, 2012), basketball (Paul and Weinbach, 2005a), and ice hockey (Woodland and Woodland, 2001; Gandar et al., 2004; Paul and Weinbach, 2012), occasionally allowing also profitable betting strategies. Furthermore, in baseball, Woodland and Woodland (1994; 2003) find individual betting lines efficient but observe a RFLB when considering all lines simultaneously, whereas Gandar et al. (2002) find only a subtle bias after making a correction to the Woodland-Woodland methodology.

4.2. Findings in soccer

During the recent decades scholars have begun to examine weak form efficiency also in soccer, initiated by Pope and Peel (1989) who analyze odds quoted by four national high street bookmakers in the UK. With a linear probability model, they regress match outcomes on bookmakers' subjective probabilities using weighted least squares estimation. Their evidence shows not only that there are differences in pricing of certain types of bets between the bookmakers, but also that there are systematic differences in the apparent odds setting processes of the firms. In this respect, Pope and Peel argue that pooling of information contained in odds leads to more efficient forecasts. At the same time, however, none of these superior forecasts results in profits after considering the bookmaker margin. The authors end up concluding that the betting market appears to meet the most important (economic) criterion for market efficiency, even though all the odds, especially those for the draw outcome, do not seem to meet the axioms of rational expectations.

Paton and Vaughan Williams (1998) employ tobit regression with a small sample of 1X2 Premier League odds and detect a favorite-longshot bias, explaining the phenomenon with transaction costs. Also in English soccer, Cain et al. (2000) report evidence of a FLB, similar to that found in horse racing, both for match results and scores when comparing the estimated fair odds with a bookmaker's actual odds. Their calculations of probability for different categories of bets suggest that betting on strong favorites may offer limited profits. Also in a later study, Cain et al. (2003) find a FLB for match results, but they underline the need for inspecting more closely whether the FLB takes place only at the high and low probability ends of the odds spectrum or monotonically throughout the spectrum. Kuypers (2000), on the contrary, performs weak form tests of efficiency in four English divisions, using odds of only one bookmaker. He detects neither differences between subjective and objective probabilities nor strategies that would lead to positive returns. Moreover, Kuypers' results give no support to the Pope and Peel (1989) assertion on biased odds for draws.

In an essential paper, Paton and Vaughan Williams (2005) hypothesize that the existence of outlier odds may provide otherwise uninformed bettors with forecasting information that can be used to build improved betting strategies. Using data from the UK bookings points spread betting market, the authors define quasi-arbitrages (or quarbs) as situations in which the average or mid-point of all the quoted spreads lies outside the top or bottom end of the spread quoted by at least one bookmaker. When detecting a quarb, they analyze whether it is the average market position or the outlying position which provides most information and whether it is possible to implement a profitable betting strategy based on this information. More specifically, the logic here is that in the absence of other information, the mid-point of all spreads should provide an obvious point estimate of the value of the asset; then, as long as this value is greater (lesser) than the price at which the bettor buys (sells), one can expect positive returns. The results clearly show that the market mid-point price comprises a better forecast of betting asset values than the outlier odds. Even more importantly, when taking advantage of the market mid-point as a predictor, Paton and Vaughan Williams demonstrate positive and superior returns both in-sample and out-of-sample, even after controlling for differential risk.

Deschamps and Gergaud (2007) analyze also weak form efficiency in English soccer, observing a favorite-longshot bias both in the odds for home wins and away wins, and a reverse bias in the draw odds. Despite this evidence, they find no betting strategy with a positive return. Anyway, Deschamps and Gergaud confirm that the strategy of choosing systematically the highest available odds between different bookmakers improves the return significantly. In addition, consistent with Paton and Vaughan Williams (2005), using a more sophisticated strategy according to which the variance of odds between bookmakers acts as a signal that the bookmaker with the highest odds for an outcome might underestimate the probability of this outcome, generates an even higher return.

After taking into consideration the two major changes introduced by online sports betting, low transaction costs and the greater number of bookmakers from which to choose the highest odds, Malarić et al. (2008) test weak form efficiency in ten European soccer leagues during three seasons. Their simple strategy based on value betting yields a return of 7.9% when placing only bets with an expected return greater than 5%. Malarić et al. perform their study with a limited number of leagues and with not more than fifteen bookmakers, but they suggest that the model could be easily expanded to include both more leagues and bookmakers to increase the number of profitable bets per season.

Vlastakis et al. (2009) undertake a comprehensive empirical study in European soccer, which according to them constitutes one of the most liquid and important betting markets, focusing on the predictability of match outcomes based on information contained in odds. When evaluating the average returns of a series of different betting rules, the authors perceive a favorite-longshot bias but also evidence of overestimation of home ground advantage, together leading to a new reverse home-underdog bias, i.e. an away-favorite bias. Vlastakis et al. also employ two regression techniques, a Poisson count for forecasting scores and a multinomial logit model for forecasting outcome probabilities, separately for each bookmaker. They then combine the predictions from individual bookmakers with two different encompassing techniques; their idea with this is that if bookmakers use different methodologies when quoting their odds, aggregating this information might lead to superior forecasting performance. The authors suggest that formal econometric models result in more accurate forecasts which can then be utilized to build profitable betting strategies, indicating weak form inefficiency. In particular, they highlight that encompassing goal forecasts seem to provide superior results; at their best, regression-based encompassing of goal forecasts in the away-favorite subsample yields an average return of 13.3% per bet. However, when employing 1X2 odds only, they cannot reject economic efficiency.

Strumbelj and Sikonja (2010) use both Brier scores and ranked probability scores to evaluate bookmaker odds as probabilistic forecasts of match outcomes in six major European soccer leagues. They discover that odds of specific bookmakers are better forecasts than those of others but also that the effectiveness of odds has increased over time, even though the effectiveness varies between different leagues. Without investigating any betting strategies themselves, Strumbelj and Sikonja highlight that in the future it would be interesting to study whether more effective bookmaker odds and/or consensus among bookmakers would provide additional information to build a profitable strategy.

More recently, Direr (2013) investigates weak form efficiency of European soccer betting with odds from twelve bookmakers in 21 divisions over eleven years, totally including 79,446 matches and around 1,800,000 odds, being the largest data set thus far in the relevant literature. His method of systematically picking out odds inferior to a threshold—simply betting on favorites whose probability of winning exceeds 90%—generates a return of 4.5% with best odds and 2.8% with average odds available in the market, being robust to out-of-sample tests. In other words, only with information contained in odds and without relying on complex econometric models, Direr observes a favorite-longshot bias similar to other betting markets and, consequently, provides evidence of deviations from weak form efficiency.

Koning (2012) examines the efficiency of soccer betting odds in ten highest level European leagues during eight seasons, employing a binary logit model. He finds a persistent favorite-longshot bias, which holds across different countries and applies both to home and away wins. Moreover, in a related study, Nyberg (2014) employs a binary logit model to make a comparison to previous models but also introduces a multinomial logit model to test the three match outcomes simultaneously. His sample consists of the top four divisions in England between the years 2000 and 2013. Overall, with both models, Nyberg shows that statistically there are no large deviations from efficiency, but that the Premier League betting market is weak form inefficient, also featuring a FLB, which leads to the rejection of the null hypothesis of efficiency also for the whole sample.

4.3. Evidence of sports arbitrage

Even though the search for arbitrage opportunities would be a natural starting point to discover weak form efficiency in betting markets, it has received relatively little attention in the literature, at least partly due to the fact that betting markets were previously segmented across countries with neither any intensive competition between bookmakers nor proper possibilities to carry out arbitrage analysis (Vlastakis et al., 2009). In any case, arbitrage opportunities have been found in racetrack betting for example by Asch and Quandt (1987) among different exotic bets, by Hausch and Ziemba (1990) in cross-track betting, and by Edelman and O’Brian (2004) with their specific model involving several bet types. However, Edelman and O’Brian list also some practical obstacles (such as time lags and limited liquidity) which make arbitrage betting not completely risk-free and easily exploitable.

In soccer, Pope and Peel (1989) list a small number of conceivable arbitrage bets between odds of four UK bookmakers, whereas Dixon and Pope (2004) find no arbitrage situations between three and Forrest et al. (2005) between five UK bookmakers. Using data on the 2002 World Cup, Gil and Levitt (2007) detect only a few arbitrage opportunities, which also disappear rather quickly. Vlastakis et al. (2009), in turn, study the odds quoted by five major online bookmakers and one regular bookmaker for soccer matches in 26 different countries and events during 2002–2004, when online sports betting was in its early phase. When considering all the bookmakers, they find 63 arbitrage situations, representing 0.5% of all matches, with an average return of 21.8% (that in my view represents an impossibly high return to be exploited in practice). However, when taking into account the online bookmakers only, Vlastakis et al. identify only ten arbitrage opportunities.

More recently, even if Koning (2012) finds only a minimal number of arbitrage bets, Constantinou et al. (2013) identify arbitrage opportunities in not less than 70 matches out of the 380 matches played during the English Premier League season 2011–2012, guaranteeing an average profit of 0.57% per such match. Constantinou and Fenton (2013) also find a notable number of arbitrage bets in 14 European soccer leagues. In addition, Franck et al. (2013) provide the first paper analyzing inter-market arbitrage, i.e. the possibility to combine bets between bookmaker and betting exchange markets. They find 102 intra-market arbitrage opportunities among bookmakers but a whole of 2,287 inter-market arbitrage opportunities (in 19.2% of matches), yielding an average return of 1.4%. The authors note that these inter-market opportunities arise mostly from too generous bookmaker odds that a bettor can then sell at a higher price (i.e. at lower odds) at the betting exchange.

4.4. Summary of the earlier literature

Overall, the earlier findings on betting market efficiency in its different forms provide no waterproof evidence neither on efficiency nor inefficiency. As expected, statistical efficiency has been rejected more often than the decisive test of economic efficiency. The situation in which a betting market is found to be statistically efficient but economically inefficient can be explicated with Jensen’s (1978) weaker formulation on financial market efficiency, according to which prices reflect information to the point at which the marginal benefits of acting on the information do not exceed the marginal costs. Because bookmaker margins, which can be viewed as bookmakers’ safety buffer against biased subjective probabilities, have been large before the online betting era, marginal benefits (placing value bets that exploit biases) have not exceeded marginal costs (bookmaker margins) and the statistical inefficiencies have thereby often had only scant practical relevance. As highlighted throughout this study, this setting of large bookmaker margins does not exist anymore.

In terms of weak form efficiency in 1X2 soccer betting, a notable part of the previous literature has rejected statistical inefficiency in some form or another, often in the shape of a favorite-longshot bias, without being able to reject economic efficiency (Pope and Peel, 1989; Paton and Vaughan Williams, 1998; Cain et al., 2000, 2003; Deschamps and Gergaud, 2007; Vlastakis et al., 2009). On the other hand, some more recent studies detect statistical inefficiency but do not investigate economic efficiency at all (Strumbelj and Sikonja, 2010; Koning, 2012; Nyberg, 2014). This study aims at filling these gaps in the literature. On the one hand, the study inspects whether the statistical inefficiencies—if there still are some—could now lead also to economic inefficiencies, thanks to the more competitive odds in the present-day betting market. On the other hand, the study builds on the recently observed statistical inefficiencies by performing also economic tests of efficiency.

Regarding the choice of methodology of this study, the previous results lead to four important conclusions. First, to obtain more powerful results, the academic evidence clearly demonstrates that tests of statistical efficiency must be carried out by analyzing odds on a level-by-level basis, along with aggregate tests. Second, pooling of odds information from diverse bookmakers has been shown to lead to more efficient forecasts, making it an applicable procedure to reach a market consensus estimate. Third, when utilizing these consensus estimates, economic profits have been generated by various value betting strategies in general and quasi-arbitrage strategies in particular. Thus, when performing economic tests of efficiency, it is reasonable to apply a variety of betting strategies. Fourth, alongside value betting strategies, the research has proven the (fairly self-evident) proposition that the exploitation of highest odds instead of average market odds or odds of a specific bookmaker only yields higher returns. Even when comparing bookmakers and betting exchanges, it is usually the former who give rise to economic inefficiencies due to their too generous odds.

5. Hypotheses

This chapter presents the two hypotheses of this study, founded on the efficient market hypothesis in general and its applications in the context of sports betting markets in particular. Within the EMH, the focus is on weak form efficiency, assuming that future prices cannot be predicted by analyzing prices from the past. As described in Section 3.3, in sports betting this presumption translates into a condition according to which all information concerning upcoming sports outcomes is already included in the record of betting odds from the past.

As demonstrated above, weak form efficiency of betting markets can be further separated into statistical and economic weak form efficiency, the latter being considered as the decisive test of efficiency. Statistical efficiency implies that subjective probabilities denoted by betting odds are unbiased estimators of match outcomes. Due to bookmaker margins (that represent transaction costs in betting markets), conceivable statistical inefficiencies might not lead to economic profits and hence have only limited practical relevance, but the results are important from the academic viewpoint. Moreover, information on statistical efficiency and/or inefficiency can be used to build certain betting strategies to gain more insight into economic efficiency. Therefore, with connection to statistical weak form efficiency, we get the following hypothesis:

Hypothesis I: The European online soccer betting market is statistically weak form efficient, meaning that 1X2 betting odds are unbiased predictors of match outcomes.

Economic weak form betting market efficiency implies that there is no betting strategy that generates consistent positive returns. As explained in Section 2.2, bookmaker margins have decreased in the online betting era, which means that smaller biases in odds might these days lead to economic profits. On the other hand, as odds quoted for each and every outcome vary between different bookmakers, it might be possible to build a profitable betting strategy even under unbiased market odds by exploiting the highest odds available in the market. Thus, the hypothesis with respect to economic weak form efficiency can be written as:

Hypothesis II: The European online soccer betting market is economically weak form efficient, meaning that no betting strategy yields consistent positive returns.

6. Data

The data in the study consist of full time results as well as both average and highest 1X2 odds quoted for each outcome in the online betting market for 95,789 soccer matches in 74 divisions in 45 European countries during the seasons 2009–2014. As there are six separate odds for each match (three average and three highest), the data include in total 574,734 odds. The data are obtained from oddsportal.com, a leading sports betting odds comparison site that compiles odds quoted by dozens of online bookmakers. Collecting and processing the primary data stood for a notable part of the study's total workload.

Table 1 describes the match outcomes, odds, and probabilities in the data, while Table A.1 in Appendix A lists the number of matches per season in each country and division considered in the study. Table 1 shows that the proportions of home wins, draws, and away wins have been relatively stable during the sample period. Moreover, even though the frequency of the different outcomes is fairly close to their subjective probability, the table indicates on aggregate that home wins have occurred somewhat more often than expected, whereas the opposite is the case for draws and away wins. In addition, Table 1 highlights that the standard deviation for the subjective probability of draws is notably lower than that for home wins and away wins. Moreover, as illustrated in Table 2, the data include the number of bookmakers whose odds were available and thereby taken into account when determining the average and highest odds for each match. Table 2 proves that the number of available bookmakers has increased each year during the sample period.

For the purposes of this study, some revisions have been made to the data obtained from the odds comparison site. First, matches that were either cancelled or awarded have been removed from the data. Second, matches for which no 1X2 odds were quoted have also been eliminated from the data; however, this was the case only in a few dozen matches in the less significant divisions during the earliest seasons considered in the study. Third, to improve robustness as well as practical applicability of the results, all matches with highest odds that could have resulted in arbitrage bets of at least 5% have been taken out from the data. Even though such arbitrage bets appear occasionally in the current online betting market, odds for these considerable arbitrage bets might also include palps, i.e. obvious errors made by bookmakers who reserve the right to cancel bets made at these odds. Thus, it is the norm in the value betting community to avoid substantial outlier odds. Overall, slightly more than three thousand matches have been removed from the primary data and the results of a little

Table 1

Match outcomes, odds, and probabilities

This table presents summary statistics for match outcomes, odds, and subjective probabilities for each season in the data. The first two columns give the number of given outcomes (n) and their frequency during each season (later on referred to as objective probability). The third and fourth columns list the mean and median odds, employing the average odds quoted in the market. The last two columns portray the mean ($\overline{\rho}$), as defined in Eq. (3), and standard deviation (σ_{ρ}) of subjective probability.

	n	Frequency	Mean odds	Median odds	$\overline{\rho}$	σ_{ρ}
2009–10						
Home wins	8,824	0.460	2.41	2.05	0.452	0.156
Draws	4,958	0.258	3.57	3.31	0.264	0.042
Away wins	5,401	0.282	4.24	3.41	0.284	0.138
Total	19,183	1	3.41	3.15	0.333	0.149
2010–11						
Home wins	8,945	0.462	2.42	2.07	0.451	0.156
Draws	4,983	0.257	3.60	3.32	0.263	0.043
Away wins	5,441	0.281	4.22	3.39	0.286	0.139
Total	19,369	1	3.41	3.14	0.333	0.149
2011–12						
Home wins	8,742	0.452	2.45	2.09	0.449	0.157
Draws	5,013	0.259	3.62	3.32	0.263	0.042
Away wins	5,606	0.290	4.23	3.33	0.289	0.140
Total	19,361	1	3.43	3.13	0.333	0.149
2012–13						
Home wins	8,524	0.450	2.45	2.13	0.442	0.151
Draws	4,913	0.259	3.59	3.33	0.264	0.040
Away wins	5,502	0.291	4.02	3.25	0.294	0.135
Total	18,939	1	3.36	3.13	0.333	0.142
2013–14						
Home wins	8,526	0.450	2.55	2.13	0.442	0.161
Draws	4,851	0.256	3.70	3.36	0.260	0.044
Away wins	5,560	0.294	4.14	3.26	0.299	0.146
Total	18,937	1	3.46	3.16	0.333	0.150
Total						
Home wins	43,561	0.455	2.45	2.09	0.447	0.156
Draws	24,718	0.258	3.62	3.33	0.263	0.042
Away wins	27,510	0.287	4.17	3.32	0.290	0.150
Total	95,789	1	3.41	3.14	0.333	0.148

Table 2

Availability of bookmaker odds

Shown below are summary statistics for the number of bookmakers whose odds were available in the online betting market for each match in the data. In other words, these figures describe how many different bookmakers have been involved when determining the average and highest odds in the sample.

	2009–10	2010–11	2011–12	2012–13	2013–14	Total
Mean	33	38	43	50	56	44
Median	36	43	47	54	59	45
Standard deviation	10.3	10.9	10.4	11.4	11.1	13.5

over one hundred matches that were decided after either overtime or penalties have been manually altered to reflect their full time result.

The data set in this study is very extensive. To the best of my knowledge, in comparison with the earlier literature on betting markets in soccer, the number of matches considered is larger than in any previous study. Moreover, the data set is outstanding also in terms of the number of bookmakers taken into consideration when determining the average and highest odds quoted in the market. While the absolute majority of previous studies employ odds quoted only by a few bookmakers, the average and highest odds employed here represent the view of the whole market instead of only some specific bookmakers. Table 3 depicts the number of matches and bookmakers included in the data sets of some earlier investigations of betting market efficiency in soccer.

Table 3

Coverage of earlier betting market studies in soccer

This table portrays the number of matches and bookmakers (B's) considered in some earlier studies on betting market efficiency in soccer. Compared to all the studies mentioned below, the current study (95,789 matches and a median of 45 bookmakers per match) dominates in both standards.

Study	Matches	B's	Study	Matches	B's
Direr, 2013	79,446	12	Dixon & Pope, 2004	6,629	1–3
Hvattum & Arntzen, 2010	30,524	various	Constantinou et al., 2012	6,244	28–40
Nyberg, 2014	26,463	5–10	Kuypers, 2000	3,382	1
Koning, 2012	25,744	up to 12	Goddard & Asimakopoulos, 2004	3,139	1
Demir et al., 2012	12,880	various	Cain et al., 2000	2,855	1
Vlastakis et al., 2009	12,841	6	Cain et al., 2003	2,855	1
Malarić et al., 2008	12,128	15	Forrest & Simmons, 2000	1,694	3**
Franck et al., 2013	11,933	11*	Forrest & Simmons, 2008	1,510	7
Graham & Stott, 2008	11,000	1	Pope & Peel, 1989	1,291	10***
Strumbelj & Sikonja, 2010	10,699	10	Sessford & White, 2010	787	1
Forrest et al., 2005	9,727	5	Paton & Vaughan Williams, 2005	447	up to 5
Deschamps & Gergaud, 2007	8,377	6	Constantinou et al., 2013	380	26–49
Dixon & Coles, 1997	6,629	1–3	Milliner et al., 2009	194	1

Notes: * Includes one betting exchange. ** Represents tipsters, not bookmakers. *** Includes six tipsters.

The odds applied in the study are closing odds, i.e. odds quoted in the market just before the start of an event. Intuitively, closing odds should be the most efficient odds amongst the pre-match odds at different points in time. For example, when compared to opening odds that are primarily based on statistical analysis of the participating teams' performance and other relevant information at the time of publishing the odds for the match, closing odds can be assumed to best reflect all news, statistics, as well as the bettor sentiment regarding the match. Moreover, while bookmakers usually accept lower stakes at opening odds, which can also change quickly after they have been published without a chance for the bettor to place her bet at the initial odds in practice, bookmakers accept much higher stakes at closing odds that also tend to remain stable at the end of the pre-match betting period.

The data provide interesting insights also into bookmaker margins and arbitrage opportunities in the online betting market in soccer. Table 4 portrays the evolution of the bookmaker margin during the sample period separately for the average and highest odds. As foreseen earlier in the study, the margin has steadily decreased each year in both cases, resulting in, ceteris paribus, better chances for profitable betting. In practice, the use of highest odds in the current online betting market seems to imply zero transaction costs.

Consistent with this observation, when considering the highest odds quoted in the market for each match, sports arbitrage opportunities have increased each year both in number and size, as shown in Table 5. While for the whole sample period arbitrage bets appeared in roughly 42% of the matches, during the season 2013–2014 there was such an opportunity available in more than 50% of the matches. This finding gives initial support for rejecting Hypothesis II and contrast with the earlier studies that find either no or a minimal number of arbitrage bets (Dixon and Pope, 2004; Forrest et al., 2005; Gil and Levitt, 2007; Vlastakis et al., 2009; Koning, 2012) and a moderate number of these bets (Pope and Peel, 1989; Constantinou and Fenton, 2013; Constantinou et al., 2013; Franck et al., 2013). The abundance of arbitrage opportunities in the sample can be explained to a large extent by the low margins of the current online betting market and by the outstanding number of bookmakers considered in the study. However, as there in practice are some challenges in capitalizing on arbitrage opportunities, the existence of arbitrage bets per se is not a sufficient condition for the overall conclusion that the market is economically weak form inefficient.

Table 4

Evolution of the bookmaker margin

This table gives summary statistics for the bookmaker margin, as defined in Eq. (2), for both the average and highest odds in the data.

	2009–10	**2010–11**	**2011–12**	**2012–13**	**2013–14**	**Total**
Average odds						
Mean	0.098	0.097	0.095	0.092	0.090	0.095
Median	0.096	0.095	0.093	0.091	0.088	0.093
Standard deviation	0.015	0.016	0.015	0.015	0.015	0.015
Highest odds						
Mean	0.014	0.011	0.007	0.002	–0.002	0.006
Median	0.008	0.007	0.005	0.000	–0.002	0.004
Standard deviation	0.032	0.028	0.024	0.022	0.020	0.026

Table 5

Prevalence of arbitrage opportunities

Listed below are summary statistics for the prevalence of sports arbitrage opportunities in the data, as defined in Eq. (6), employing the highest odds quoted for each match. The first two rows represent the number of such matches (n) and their frequency during the given period, while the remaining rows describe the scale of these bets, i.e. the return that could have been earned by placing bets in correct proportions on each outcome. As explained in this chapter, bets with an arbitrage return of over five percent have been removed from the data set.

	2009–10	**2010–11**	**2011–12**	**2012–13**	**2013–14**	**Total**
n	6,456	6,697	7,574	9,230	10,442	40,399
Frequency	0.337	0.346	0.391	0.487	0.551	0.422
Mean	0.014	0.014	0.015	0.015	0.016	0.015
Median	0.011	0.011	0.011	0.012	0.014	0.012
Standard deviation	0.012	0.012	0.012	0.012	0.012	0.012

Lastly, we will take a look at the descriptive measures on the prediction accuracy of the odds in the sample. Table 6 gives the Brier scores and the rank probability scores for the sample, using the average odds. Both scores clearly exhibit that the effectiveness of the odds has remained very stable during the sample period and has especially not improved. Moreover, the scores in Table 6 indicate that the prediction accuracy of home wins is somewhat weaker than that of draws and away wins. These observations, combined with the above finding that the bookmaker margin has decreased constantly at the same time, make up a fascinating setting for the efficiency tests carried out in the study.

Table 6

Brier scores and rank probability scores

Employing the average odds, this table lists the Brier scores for each match outcome during the sample period, as defined in Eq. (10), as well as the rank probability scores (RPS) for each season, as defined in Eq. (11). For comparability, both scores have been scaled so that they vary between zero and one, zero indicating the best achievable forecast and one denoting the worst possible forecast. Coin tossing would generate a score of 0.5.

	2009–10	**2010–11**	**2011–12**	**2012–13**	**2013–14**	**Total**
Brier scores						
Home wins	0.220	0.220	0.218	0.221	0.217	0.219
Draws	0.189	0.188	0.189	0.190	0.187	0.189
Away wins	0.179	0.179	0.182	0.185	0.183	0.182
Total	0.196	0.196	0.196	0.199	0.196	0.196
RPS						
Total	0.199	0.199	0.200	0.203	0.200	0.200

7. Methodology

This chapter presents the methodology applied in the study, divided into statistical and economic parts. The statistical part performs two regression-based tests, employing linear and logit (both binary and multinomial) regression, as well as compares subjective and objective probabilities by grouping odds in two different ways. The economic part regresses the net returns of unit bets on odds applying tobit estimation, tests the rates of return of each odds group against the profitability threshold, and explores the profitability of both naïve and simulated sophisticated strategies. Altogether, this methodology toolbox is composed of the techniques most widely used in the betting market literature. On many occasions, I will use various specifications to generate robust results.

7.1. Statistical tests of weak form efficiency

In 1X2 betting, three outcomes are possible. Compared to those forms of betting that have only two conceivable outcomes, the analysis of statistical efficiency in 1X2 betting becomes more complicated because of dependence issues. As more than one outcome might have similar or almost similar odds, methods that group these bets based on the level of odds might place more than one outcome of an event into the same odds group, leading to within-group correlation as only one of the outcomes can take place. To overcome these complications, this study performs all the statistical tests of efficiency, except the one that applies multinomial logit, in three parts, i.e. separately for home wins, draws, and away wins. This way there will be only two possible outcomes in each case, for example a home win and a complement of the home win, removing the potential problem of dependency.

7.1.1. Linear regression: subjective vs. objective probability of odds groups

I will start exploring the relation between subjective and objective probability on the aggregate level with linear regression using odds groups, estimating the simple model

$$\overline{\pi}_h = \alpha + \beta \times \overline{\rho}_h + e_h, \tag{13}$$

where $\overline{\pi}_h$ is the objective probability of the odds group h, $\overline{\rho}_h$ is the subjective probability of this odds group, and e_h is the error term. Here, weak form betting market efficiency implies that $\alpha = 0$ and $\beta = 1$; if this holds, subjective probability is an unbiased estimator of objective probability (see, e.g., Zuber et al., 1985; Gandar et al., 1988; Sauer et al., 1988; Golec and Tamarkin, 1991; Gray and Gray, 1997). If $\beta > 0$, the betting market predicts the outcome in the correct direction, while the opposite holds if $\beta < 0$.

To carry out the linear regression (as well as the grouping-based methods introduced later in this chapter), odds must be grouped according to a criterion so that nearby odds, i.e. all odds falling within specific intervals, are assigned into a same group. The literature recognizes two such criteria, originally developed by Ali (1977) and Snyder (1978). The first alternative is to determine the intervals so that each group would include approximately the same number of odds, as later done by Asch and Quandt (1987) and Busche and Hall (1988), for example. The second choice is to define the intervals so that each probability interval is of similar size, as later conducted in soccer by Kuypers (2000), for instance.

Beyond these two alternatives there are no widely accepted guidelines for defining the intervals. Snyder (1978), for example, uses only his own discretion and convenience. In any case, grouping may introduce measurement error and the groupings should be chosen so that they maximize between group (and minimize within-group) variation (Busche and Hall, 1988; Vaughan Williams and Paton, 1998). The larger the group size, the more observations there are in each group, generating more reliable test statistics; on the other hand, a larger group size means also a fewer number of groups and therefore more aggregation, i.e. higher variation of odds within the groups. Thus, to obtain robust results, it is advisable to apply multiple intervals and check whether different choices give similar results.

I will perform the linear regression with three different groupings, taking advantage of the considerable size of the data set. First, I estimate the model with one thousand equally sized groups. Second and third, I repeat the estimation with one hundred equally sized groups and fifty groups with an equal probability interval. To be exact, since we are here more interested in the subjective probabilities implied by the odds than in the odds themselves, I carry out all the groupings based on subjective probability, as defined in Eq. (3).[36]

[36] I see this as an important procedural choice, considering that each match has its unique bookmaker margin. For example, if we have two matches in which the average odds for a home win are 1.38 and 1.42, these odds would most likely be (incorrectly) classified in different groups. If, however, we take into account that the bookmaker margins for these matches are 0.08 and 0.05, respectively, the subjective probability of a home win in both matches is actually 0.671 and the matches would be (correctly) classified in the same group.

7.1.2. Logit regressions: outcome-specific binary and all-inclusive multinomial

After employing standard linear regression and getting the first impression of weak form betting market efficiency on the aggregate level and concerning each outcome separately, I will investigate the issue also with logit regression, a nonlinear regression model specifically designed for binary dependent variables. More specifically, I will apply both binary logit regression to test the efficiency of each outcome separately and multinomial logit regression to test the efficiency of all the three outcomes simultaneously.

Generally, nonlinear binary response models use formulations that force the predicted values to be between zero and one. For a binary dependent variable Y, these models estimate the probability that $Y = 1$. There are two such models: probit regression that uses the standard normal cumulative probability distribution function (c.d.f.) and logit regression that applies the logistic c.d.f. Even though the logistic c.d.f. has thinner tails than the normal c.d.f., these models often produce very similar results. Both probit and logit regression have been employed in the betting market literature. Different versions of probit regression have been utilized by Kuypers (2000), Boulier and Stekler (2003), Goddard and Asimakopoulos (2004), Forrest et al. (2005), Graham and Stott (2008), and Franck et al. (2010; 2013), for example. Logit regression has been used by Pope and Peel (1989), Forrest and Simmons (2000), Vlastakis et al. (2009), Hvattum and Arntzen (2010), Koning (2012), and Nyberg (2014), for instance.

Following Koning (2012), the binary logit model, which tests the efficiency of each outcome separately, can be defined as follows. Since the outcome of a soccer match i is either a home win (HW), a draw (D), or an away win (AW), these events can be indicated by dummy variables $Y_i^{HW} = 1$, $Y_i^{D} = 1$, and $Y_i^{AW} = 1$, while the variables are zero if the outcome does not occur. Thus, when the match is finished, we have three dummy variables that sum up to one. Given that odds reflect all information, the expected payout (including the stake) on a unit bet on a home win in the match i is then

$$\vartheta_i^{HW} \times \Pr(Y_i^{HW} = 1) + 0 \times \Pr(Y_i^{HW} = 0) = 1, \tag{14}$$

where ϑ_i^{HW} is the scaled odds for a home win in the match, defined in Eq. (4). In other words, informationally efficient odds for a home win imply that $\Pr(Y_i^{HW} = 1) = 1/\vartheta_i^{HW}$. Because

$$\frac{1}{\vartheta_i^{HW}} = \frac{1}{1+\vartheta_i^{HW}-1} = \frac{1}{1+\exp(\log(\vartheta_i^{HW}-1))} = \frac{1}{1+\exp(-0-(-1)\log(\vartheta_i^{HW}-1))} = \frac{1}{1+\exp((-)(0-\log(\vartheta_i^{HW}-1)))},$$

weak form betting market efficiency can be investigated by estimating the logit model

$$\Pr(Y_i^{HW} = 1) = \frac{1}{1+\exp((-)(\alpha+\beta\log(\vartheta_i^{HW}-1)))}, \tag{15}$$

testing the joint hypothesis that $\alpha = 0$ and $\beta = -1$. Similar procedure naturally applies to draws and away wins. Besides performing this analysis for each outcome, I will estimate the above model also for each season separately to see if there are changes in the degree of efficiency during the sample period. Moreover, similar to Nyberg (2014), to combine the information of the independent logit models and their test statistics, I consider a simple and intuitively plausible procedure in which the total value of the log-likelihood function is the sum of the independent logit models. In that case, a joint test of the hypothesis that $\alpha = 0$ and $\beta = -1$ can be employed in the conventional way for all the outcomes simultaneously.

After employing binary logit regression for each outcome separately, I will also use multinomial logit regression to model the three possible soccer match outcomes in tandem within one model. Here, efficiency testing can be based on only one test statistic instead of the three obtained above. In fact, multinomial logit regression is a generalization of the binary logit model and, following Greene (2012, 803–805) and Nyberg (2014), can be defined as follows.

If we now denote the outcomes of a match i as Y_i so that $Y_i^{HW} = 1$, $Y_i^D = 0$ and $Y_i^{AW} = -1$, we can have three binary indicator variables so that $Y_{i,1} = 1$ and $Y_{i,0} = Y_{i,-1} = 0$ if the result of the match i was a home win. Similarly, we have $Y_{i,0} = 1$ and $Y_{i,1} = Y_{i,-1} = 0$ if the result was a draw, as well as $Y_{i,-1} = 1$ and $Y_{i,1} = Y_{i,0} = 0$ if the result was an away win. Assume also that the subjective probabilities assigned by the market for the match are contained in the vector

$$\rho_i = [\rho_i^{HW} \quad \rho_i^D \quad \rho_i^{AW}]. \tag{16}$$

The multinomial logit model is then specified when we determine the conditional probabilities $\rho_{i,j}$ $(j = -1, 0, 1)$ of the match outcomes $Y_i = j$, conditional on the relevant predictive information, which under weak form betting market efficiency should be completely included in Eq. (16). The model can be written applying the log-odds ratios so that

$$\log\left(\frac{\pi_{i,j}}{\pi_{i,0}}\right) = \tau_{i,j}, \tag{17}$$

where the linear functions $\tau_{i,j}$, $j = -1, 1$, should be determined to complete the model. Using the subjective probabilities in Eq. (16), the log-odds ratios reduce to

$$\log\left(\frac{\rho_{i,j}}{\rho_{i,0}}\right) = \gamma_{i,j}. \tag{18}$$

Thus, as we want to investigate the predictive power of the subjective probabilities, the linear functions given in Eq. (17) can be specified as

$$\tau_{i,j} = \alpha_j + \gamma_{i,j}\beta_j, \qquad j = -1, 1, \tag{19}$$

where the parameters α_j and β_j are outcome-specific. Similar to Vlastakis et al. (2009) and Nyberg (2014), for example, I will use draws as the benchmark category in Eq. (17), which indicates that the linear function in Eq. (19) is not determined for that outcome. For robustness, however, I will repeat the analysis using also home wins and away wins as the benchmark category, respectively. Finally, when solving the conditional probabilities $\pi_{i,j}$ from Eq. (17), we get

$$\begin{aligned}
\pi_{i,1} &= \Pr(Y_i = 1|\rho_i) = \frac{\exp(\tau_{i,1})}{1+\exp(\tau_{i,1})+\exp(\tau_{i,-1})} \\
\pi_{i,0} &= \Pr(Y_i = 0|\rho_i) = \frac{1}{1+\exp(\tau_{i,1})+\exp(\tau_{i,-1})} \\
\pi_{i,-1} &= \Pr(Y_i = -1|\rho_i) = \frac{\exp(\tau_{i,-1})}{1+\exp(\tau_{i,1})+\exp(\tau_{i,-1})},
\end{aligned} \tag{20}$$

where $\sum_{j=-1}^{1} \pi_{i,j} = 1$. This equation shows that the linear functions in Eq. (19) completely determine the conditional probabilities of the different soccer match outcomes. The parameters of this multinomial logit model can be estimated by maximum likelihood, which also facilitates the use of the conventional likelihood-based test statistics when investigating market efficiency.

Statistical weak form betting market efficiency can be then tested with a restricted multinomial logit model in which $\tau_{i,1} = \gamma_{i,1}$ and $\tau_{i,-1} = \gamma_{i,-1}$. Therefore, we get the null hypothesis

$$H_0 : \alpha_j = 0, \beta_j = 1, j = -1, 1. \tag{21}$$

If the null hypothesis cannot be rejected, subjective probabilities are statistically unbiased and informationally efficient predictors of match outcomes. If, on the other hand, the null hypothesis is rejected, match results deviate systematically from subjective probabilities and it would be possible to get more accurate probability forecasts using the unrestricted model.

7.1.3. Grouping based on subjective probability

As the last test of statistical efficiency, I will employ the widely used method that sorts odds into groups based on subjective probability and compares the groups' subjective probabilities with their objective probabilities. This turns the focus from aggregate level to odds level, allowing us to go through the spectrum of odds on a level-by-level basis and clarify whether there are deviations between subjective and objective probabilities on some specific levels of odds. After Ali's (1977) groundwork in developing a grouping method for horse racing based on the favorite rank, Snyder (1978) is the first to provide a test that groups the observations (odds) based on the level of odds. The method proceeds in the following way.

First, similar to standard linear regression presented in Section 7.1.1, the odds are classified into groups so that nearby odds are assigned into a same group. Second, the groups' average subjective probabilities are calculated by taking the average of the subjective probabilities implied by all the odds included in each group. The estimators for the objective probabilities are calculated as a ratio of the number of winning bets in each group to the total number of bets in the same group. Formally, the average subjective probability and the estimator for the objective probability are given as

$$\overline{\rho}_h = \frac{\sum_{i=1}^{n} \rho_{hi}}{n_h} \quad \text{and} \tag{22}$$

$$\overline{\pi}_h = \frac{\sum_{i=1}^{n} Y_{hi}}{n_h}, \tag{23}$$

where $\overline{\rho}_h$ is the subjective probability of the group h, ρ_{hi} is the subjective probability of the bet in the group h and match i, n_h is the number of matches in the group h, $\overline{\pi}_h$ is the estimator for the objective probability of the group h, and Y_{hi} is a variable for the bet in the group h and match i so that $Y_{hi} = 1$ if the bet wins in the i:th match and zero otherwise. As the matches can be considered as independent binomial trials, Y_{hi} follows a binomial distribution so that $E(Y_{hi}) = \overline{\pi}_h$ and $Var(Y_{hi}) = \frac{\overline{\pi}_h(1-\overline{\pi}_h)}{n_h}$.

Finally, after determining the subjective and objective probabilities I will calculate whether these probabilities are equal within each group h, applying the null hypothesis

$$H_0 : \overline{\rho}_h - \overline{\pi}_h = 0 \tag{24}$$

for all h = 1, 2, …, H, where H refers to the number of groups. This analysis reveals whether the betting market is able to predict the winners and whether the calibration of odds varies between different levels of odds. Given the size of the sample, the binomial distribution can be approximated with the normal distribution and, due to the Central Limit Theorem, the test statistic for the group h can be formally defined as

$$z_h = \frac{\overline{\rho}_h - \overline{\pi}_h}{\sqrt{\frac{\overline{\pi}_h(1-\overline{\pi}_h)}{n_h}}} \sim N(0,1). \tag{25}$$

In this study, I will conduct the grouping analysis with four different categorizations. First, I sort the odds into fifty equally sized groups, thereafter repeating the procedure with fifty groups with an equal probability interval. Third and fourth, I carry out the analysis with twenty equally sized groups as well as twenty groups with an equal probability interval, respectively. In soccer, this method has been employed with varying categorizations by Pope and Peel (1989), Kuypers (2000), and Koning (2012), for instance.

7.2. Economic tests of weak form efficiency

This section presents the four economic tests of efficiency performed in the study. Instead of average market odds, the economic tests employ highest odds available in the market and explores whether historical information on odds can be utilized to make a profit. Similar to the tests of statistical efficiency, also the economic tests begin with an aggregate test, employing tobit regression that examines the whole spectrum of odds simultaneously. The second test makes a shift in focus to odds level, inspecting whether profits can be made on some specific levels of odds. The third part tests the profitability of some naïve betting strategies, while the fourth and final part encompasses tests of some specific sophisticated strategies that apply value betting and simulation.

7.2.1. Tobit regression: returns of unit bets on odds

As the first test of economic efficiency, I will apply tobit regression to investigate whether the returns of unit bets are systematically non-positive across all odds levels, originally proposed by Paton and Vaughan Williams (1998) and Vaughan Williams and Paton (1998), and later applied for example by Sessford and White (2010). Since this method aggregates returns and odds, we do not need to make any groupings and therefore avoid the potential challenges of doing (sometimes artificial) classifications. Separately for each match i and outcome j in the sample (due to the correlation issues discussed earlier), I will first calculate the net returns of a unit bet so that $R_{i,j} = \delta_{i,j}^{B} - 1$ if the bet won and $R_{i,j} = -1$ if the bet did not win. We can then perform the regression

$$R_{i,j} = \alpha + \beta \times \delta_{i,j}^{B} + e_{i,j}, \tag{26}$$

where $R_{i,j}$ is the realized net return of a unit bet on the outcome j in the match i, $\delta_{i,j}^{B}$ is the best odds for the given outcome, and $e_{i,j}$ is the error term. Because the dependent variables are now left-censored at –1, OLS regression would lead to biased estimates and tobit estimation is more appropriate.

In an economically weak form efficient market, returns should be zero or negative throughout the spectrum of odds. Thus, given that the coefficient β does not exceed zero, we get the null hypothesis

$$H_0 : \alpha \leq 0. \tag{27}$$

On the other hand, tobit regression also reveals whether the returns vary between different odds levels. If $\beta = 0$, the level of odds does not have any effect on returns. However, if $\beta < 0$, returns are lower at higher odds, indicating a favorite-longshot bias; if $\beta > 0$, returns are higher at higher odds, which would point to a reverse favorite-longshot bias.

7.2.2. Returns of odds groups

After inspecting economic weak form betting market efficiency on the aggregate level, I will move on to exploring it on the individual odds level by sorting the odds into groups in the same way that in the statistical efficiency part and perform a simple test of the returns of unit bets in each group. In soccer, this procedure has been carried out by Pope and Peel (1989), Cain et al. (2000; 2003), Kuypers (2000), Dixon and Pope (2004), Deschamps and Gergaud (2007), Malarić et al. (2008), and Direr (2013). Similar to the statistical test of grouping based on subjective probability, I will use four different groupings: fifty and twenty equally sized groups as well as fifty and twenty groups with an equal probability interval. In each case, after assigning the odds into the groups, I will use the returns of unit bets as obtained in the previous section and employ the standard one-sided t-test to examine whether the return of any group is positive. Thus, for each group h, we get the null hypothesis

$$H_0 : \overline{R}_h \leq 0, \tag{28}$$

where $\overline{R}_h$ is the average net return of unit bets in the group h.

To examine the profitability of various betting strategies, different versions of the standard t-tests as well as z-tests have been the norm in the literature (see, e.g., Snyder, 1978; Tryfos et al., 1984; Gandar et al., 1988; Paton and Vaughan Williams, 2005; Demir et al., 2012). However, there is no unanimous view concerning whether one should apply one-sided or two-sided tests. Woodland and Woodland (2003), for example, discuss that when inspecting the equality of observed outcomes with conditions for efficiency, a two-sided test is appropriate. Meanwhile, if the objective is to test betting market returns directly, as is the case both in this and their own study, Woodland and Woodland recommend a one-sided test.

7.2.3. Returns of naïve betting strategies

The last two economic tests of efficiency compose different betting strategies and examine whether they result in consistent profits. These strategies are divided into two categories: naïve and sophisticated. In the former, I will calculate the returns obtained with five different naïve strategies, applying two simple staking plans. The naïve strategies consist of betting on every home win, draw, and away win, respectively, as well as of betting on the favorite (the outcome with the lowest odds in each match) and on the longshot (the outcome with the highest odds). In other words, these strategies place one bet on every match in the sample following a specific straightforward rule, without making any other distinction between different matches and the odds relating to them (see, e.g., Kuypers, 2000; Forrest et al., 2005; Deschamps and Gergaud, 2007; Vlastakis et al., 2009; Franck et al., 2013).

The staking strategies applied with the naïve strategies are unit bets and unit wins, introduced in Section 2.4.4. After calculating the net return of each bet, I will perform the simple one-sided t-test to find out whether the returns are systematically zero or negative. The null hypothesis for all the naïve (as well as sophisticated) strategies will therefore be

$$H_0 : \overline{R}_s \leq 0, \tag{29}$$

where $\overline{R}_s$ is the average return of the given strategy s. Moreover, to see whether there are systematic differences between the returns generated by the two staking strategies, I will test the equality of the returns of each naïve strategy, using the Wilcoxon signed-rank test[37], as done in the context of soccer betting for example by Demir et al. (2012).

[37] Developed by Wilcoxon (1945), the Wilcoxon signed-rank test is a nonparametric statistical hypothesis test that can be used for example to compare two related samples.

7.2.4. Returns of sophisticated betting strategies

The economic tests—as well as the whole study—culminate in the exploration of the returns earned by the sophisticated strategies. Here, I will utilize the two cornerstones of profitable betting presented in Section 2.4: value betting and Kelly staking. Regarding value betting, I will seek an edge [cases in which r_j in Eq. (5) exceeds one] through two strategies that rest on quasi-arbitrage and the binary logit model constructed in the statistical efficiency part of this study, respectively. The idea with applying both quasi-arbitrage and the logit model is that bets with a positive expected value might stem from two distinct sources: either from taking advantage of outlier odds in a statistically efficient market or from exploitation of the observed statistical inefficiencies, in which case our logit model would be useful to determine the true probabilities. Overall, different variants of value betting have been extensively covered in the betting market literature (see, e.g., Dixon and Coles, 1997; Kuypers, 2000; Dixon and Pope, 2004; Vlastakis et al., 2009; Franck et al., 2010; Hvattum and Arntzen, 2010; Constantinou et al., 2013; Direr, 2013).

First, introduced by Vaughan Williams (2001), quasi-arbitrage refers to a setting in which the highest odds quoted in the betting market lies so far outside the average market odds that, taking into account the bookmaker margin, the bettor can place a bet with a positive expected value. Therefore, in the quasi-arbitrage strategy, I will place bets only on outcomes and matches in which such quasi-arbitrage opportunity is available, similar to Paton and Vaughan Williams (2005) and Deschamps and Gergaud (2007). The numerator in Eq. (5) is now represented by the subjective probability of the average market odds, while the offered probability of the highest odds stands for the denominator. I will calculate the annual returns of this strategy with six different thresholds for r_j: 1.00, 1.05, 1.10, 1.20, 1.30, and 1.50; the higher the threshold, the higher the deviation between the average and highest market odds.

Second, I will seek an edge by utilizing the coefficients for the binary logit model obtained in the statistical tests of efficiency. In soccer, similar techniques have been employed by Pope and Peel (1989) and Vlastakis et al. (2009). Here, the procedure is otherwise similar to the quasi-arbitrage strategy, but the numerator in Eq. (5) is now represented by the probability estimate generated by the binary logit model, based on the logit coefficients obtained for the whole sample period.

In theory, the returns of this logit-based strategy might be subject to in-sample bias because the logit coefficients applied in the strategy are founded on a sample that is similar in terms of actual match results to the sample that is then used to simulate the strategies.[38] In practice, however, the risk of bias is here considerably diminished due to two factors. On the one hand, the used logit coefficients represent those of the whole sample period, not of single seasons. On the other hand, even though the sample is similar in terms of match results, the strategy anyway employs a different set of odds, i.e. highest odds instead of average market odds that were utilized when building the logit model. In any case, for robustness, I will perform also out-of-sample tests, employing the seasons 2009–2012 as the in-sample period, according to which the logit model is determined, then operating the seasons 2012–2014 as the out-of-sample period.

In connection with both the sophisticated strategies, I will follow the Kelly staking plan to determine the stakes. As discussed in Section 2.4.5, in betting practice it is advisable to apply a fractional Kelly plan that increases security at the expense of somewhat compromising long term growth maximization. Thus, besides employing the full Kelly plan [in which $g = 1$, as defined in Eq. (8)], I will operate three fractional Kelly plans in which $g = 1/2$, $g = 1/4$, and $g = 1/20$.

Moreover, since the size of a Kelly stake is dependent on the current size of the bettor's bankroll, inherent in Kelly staking is the fact that outcomes of consecutive matches have their influence on forthcoming stakes. At the same time, soccer matches can be considered independent of each other per se. Therefore, to make statistical inferences on both the sophisticated strategies, I will use simulation in which I let the match order to differ randomly within each season in the sample, thereby obtaining a distribution of the returns separately on each above threshold for r_j as well as on each above level of Kelly staking. For each season, I will generate a random match order, calculate the returns of the two sophisticated strategies with the above specifications and repeat the procedure one thousand times, eventually performing the simple t-test to find out whether the returns are not positive, operating the null hypothesis defined in Eq. (29). In the betting market literature, different forms of simulation have been employed by Constantinou et al. (2012), Demir et al. (2012), and Nyberg (2014), for instance.

---

[38] Thus, in terms of a potential in-sample bias, the logit-based sophisticated strategy differs from the strategy that is based on quasi-arbitrage. Because in the quasi-arbitrage strategy subjective probability is based only on the average market odds quoted for the given match, that strategy contains definitely no risk of any in-sample bias.

8. Assumptions

This chapter presents the key assumptions for carrying out the tests of efficiency. The first two assumptions interconnect with both the statistical and economic tests of efficiency. The third and fourth assumptions relate to the statistical tests of efficiency only, while the fifth, sixth, and seventh assumptions are linked merely to the economic tests of efficiency.

8.1. Determination of market consensus

When determining the subjective probabilities assigned by the betting market, the study assumes that market consensus is represented by the average odds quoted in the market. This is a painless and natural assumption and common in the betting market literature (see, e.g., Buchdal, 2003, 43–48; Levitt, 2004; Paton and Vaughan Williams, 2005; Deschamps and Gergaud, 2007; Hvattum and Arntzen, 2010; Koning, 2012). If the betting market processes all known information in the most efficient manner, one might expect that average odds provide the best indicator of the actual outcome.[39] On the other hand, if a specific bookmaker possesses privileged information or at least has superior expertise in processing public information, outlier odds might provide a more accurate predictor of the outcome.

8.2. Bettor rationality

Betting strategies employed in the study assume that bettors are risk neutral and rational wealth maximizers with a linear utility of wealth. This assumption translates into two separate attributes. First, risk neutrality denotes that no bet imposes additional costs or benefits due to its probability of winning, i.e. depending on the riskiness of the bet. Second, as odds offered by different bookmakers on same outcomes vary, rational bettors always place their bets at the bookmaker offering the most favorable odds. By combining bets at different bookmakers and disregarding other factors such as habits, personal favorites, and levels of maximum payouts, they are able to reduce their exposure to risk and/or increase their expected return from betting strategies (Pope and Peel, 1989).

---

[39] Also in other contexts, aggregation of information has been shown to provide valuable information in predicting future outcomes. Regarding future GNP growth, for instance, Clemen and Winkler (1986) and Fomby and Samanta (1991) observe that the consensus estimate is a better predictor than any one single estimate.

8.3. Bookmaker book balancing

Following the convention of the literature that examines betting market efficiency under the traditional models of sportsbook behavior, this study assumes that bookmakers set market clearing prices (odds) by balancing their books (see, e.g., Pankoff, 1968; Zuber et al., 1985; Sauer et al., 1988; Dare and MacDonald, 1996; Gray and Gray, 1997; Gandar et al., 2002). In other words, bookmakers are assumed to play the role of market makers who match buyers and sellers without taking positions by themselves. If odds are set this way, betting action will be split between the different outcomes, bookmakers earning risk-free returns as soon as a balance is achieved. When odds are determined as a result of bettor actions, sports betting markets become a natural place to test the efficient market hypothesis with its different variations (Paul and Weinbach, 2012).

This view on bookmaker behavior has been dominant in the literature, but it does not necessarily hold. The traditional model has been challenged by assuming that bookmakers set odds to maximize profits rather than to balance their books and might therefore take positions with respect to match outcomes, quoting inefficient odds to exploit bettor biases (see, e.g., Treynor, 1987; Avery and Chevalier, 1999; Millman, 2001; Levitt, 2004; Surowiecki, 2004). If this would be the case, odds would not represent bettors' views on outcome probabilities, questioning the methodology applied in this study.

Even though studies that have recently tested this contradictory view do not discard it, they demonstrate that the strategic odds setting by bookmakers is substantially more uncommon and less profitable than proposed earlier (Paul and Weinbach, 2009; 2012; Krieger et al., 2013). Instead, it has been suggested that bookmakers might set their odds as a forecast of actual match outcomes, which would be a profit maximizing strategy in the long run even in the presence of imbalanced betting volumes (Paul and Weinbach, 2012; Franck et al., 2013).[40] If this would be the case, odds would not represent bettors' views on outcome probabilities but would instead reflect bookmakers' views on these probabilities. In any case, from this study point of view, the methodology for determining subjective probabilities is justifiable as long as bookmakers set their odds either as forecasts or to balance their books, the only difference then relating to the actual source of these probabilities.

[40] Paul and Weinbach (2012) also add that by setting odds as a forecast of actual game outcomes bookmakers would still earn their commission on losing bets in the long run, without the transaction costs that are necessary to attempt to balance the book or to price to exploit bettor biases on an event-by-event basis. Furthermore, the authors note that this strategy lessens the incentive for informed bettors to enter the market, eventually preventing these bettors to take away additional profits from bookmakers in the long run.

8.4. Distribution of the bookmaker margin

In addition to assuming that bookmakers balance their books, this study assumes that the bookmaker margin is spread proportionally across each possible outcome of any event. Together with the above assumption about bookmakers balancing their books, this is a necessary condition for calculating subjective probabilities using Eq. (3), i.e. determining these probabilities by dividing odds quoted in the market by the margin, and is widely used in the literature.[41] If the books of a bookmaker are not in balance and/or if its margin is spread unproportionally, questions of market efficiency would become more difficult to address because the subjective probabilities could not be exposed (Woodland and Woodland, 1994).

8.5. Search and transaction costs

The thesis also assumes that search and transaction costs, naturally excluding the built-in costs in bookmaker margins, are zero. This is a relatively undemanding condition and widely applied in the literature. Direr (2013) notes that there are nowadays several websites offering odds comparisons between bookmakers for free and argues that shopping around for highest possible odds does not entail significant search costs. When it comes to transaction costs, depositing and withdrawing funds to and from accounts at different bookmakers, which constitutes an unavoidable task when using multiple bookmakers, comprises another potential source of transaction costs besides bookmaker margins. In this respect, there are nowadays money transfer sites that offer zero costs for transferring money to and from bookmakers. Moreover, as bettors almost always receive a welcome bonus when joining a new bookmaker, overall transaction costs may well turn out to be even negative in practice.

[41] Studies in which bookmaker margins are assumed to spread this way include, among others, Kuypers (2000), Dixon and Pope (2004), Goddard and Asimakopoulos (2004), Forrest et al. (2005), Deschamps and Gergaud (2007), Graham and Stott (2008), Malarić et al. (2008), Milliner et al. (2009), Vlastakis et al. (2009), Demir et al. (2012), Koning (2012), and Nyberg (2014).

8.6. Stakes and bookmaker limits

When carrying out the economic tests of efficiency, the thesis assumes that one might bet any given sum of money at any given odds in the data. It is therefore assumed that bookmaker stake limits are absent. This condition might be restrictive in some occasions, possibly harming the implementation of certain betting strategies. This is because bookmakers currently have the technical means to identify skilled bettors and arbitrageurs and always reserve the right to limit the stakes of individual customers or to close the accounts of these customers completely (Franck et al., 2013). As the strategies necessitate placing bets on the highest odds quoted in the market as well as using exact stakes, for example betting 2,147 euros instead of betting exactly two thousand euros, such bets might raise the attention of bookmakers who counteract their sharpest clients and lead to stake limits.

8.7. Bettor taxation

When testing economic efficiency, the study assumes that online bettors do not have to pay any taxes on their betting activities in general and winnings in particular. This assumption is due to two facts that together in practice lay the foundation for tax-free betting in a significant part of Europe, for instance. First, winnings from online betting are currently not taxed at all in many European countries or taxed only for winnings received from bookmakers that are based outside the EU; overall, the current outlook of bettor taxation is manifold and subject to constant change. Second, even if living in a country in which non-EU online betting winnings are taxed, a bettor can always reduce or completely eliminate her tax burden by transferring funds into some of her EU-bookmaker accounts using matched betting introduced in Section 2.4.3. In this case, if necessary, she would withdraw funds only from her EU-bookmakers and never withdraw any net winnings from non-EU bookmakers.[42]

[42] As such, the intuition behind this procedure is the same than in tax planning activities of multinational corporations who aim at earning accounting profits in countries with lower taxation.

9. Results

This chapter presents the results of the statistical and economic tests of weak form betting market efficiency performed in the study, as introduced in Chapter 7. A discussion on the results follows then in Chapter 10.

9.1. Statistical tests of weak form efficiency

In the following, I will display the results of the three statistical tests of efficiency as described in Section 7.1, applying average market odds. The first two tests reveal whether the 1X2 betting market in European soccer is statistically weak form efficient on the aggregate level, whereas the third test exposes the degree of efficiency on the individual odds level.

9.1.1. Linear regression: subjective vs. objective probability of odds groups

The comparison between the subjective and objective probability of odds groups casts doubt on statistical weak form betting market efficiency on the aggregate level. As shown in Table 7, standard linear regression using one thousand equally sized groups rejects market efficiency for all outcomes with high values of the F-statistic. Moreover, the coefficient β is significantly larger than one, which means that when subjective probability increases, objective probability increases more than implied by market efficiency and vice versa. In other words, average market odds seem to be too high for favorites and too low for longshots. This indicates a favorite-longshot bias, the most common departure from market efficiency in the earlier literature highlighted in Section 4.1. These results contrast sharply with Kuypers (2000) who neither is able to reject weak form efficiency nor detects any systematic bias, though with a sample roughly thirty times smaller than that of this study, thereby providing much less power.

Another interesting aspect of the linear regression results in Table 7 is the notably different results for draws when compared to those for home wins and away wins. The estimated coefficients and the value for R-squared reveal that odds for draws appear to be considerably more biased than those for the other outcomes and the linear regression model fits the data for draws much worse. This feature has been noticed earlier by Pope and Peel

Table 7

OLS estimates and efficiency testing results of 1,000 equally sized groups

This table presents the OLS estimates and the efficiency testing results of the standard linear regression model, as defined in Eq. (13), for 1,000 equally sized odds groups. The regressions have been run separately for home wins, draws, and away wins. The groups are sorted by subjective probability, as defined in Eq. (3), so that the group size is 95 for 211 groups and 96 for 789 groups. Robust standard errors of the estimated coefficients are reported in parentheses. In the case of the coefficient β, the null hypothesis is that $\beta = 1$. The sixth and seventh rows give the F-statistics for the joint test that $\alpha = 0$ and $\beta = 1$, which denotes the test of weak form betting market efficiency. *** indicates significance at the 1% level.

	Home wins	**Draws**	**Away wins**
α	–0.0338***	–0.0724***	–0.0302***
	(0.0041)	(0.0075)	(0.0028)
β	1.0930***	1.2582***	1.0934***
	(0.0087)	(0.0292)	(0.0097)
R^2	0.929	0.572	0.923
F	68.61***	51.60***	57.74***
(p-value)	(0.000)	(0.000)	(0.000)
n	1,000	1,000	1,000

(1989), Dixon and Pope (2004), as well as Koning (2012), all of whom highlight bookmakers' potential inability to predict draws and propose that bookmakers might systematically underestimate the variance of the probabilities associated with this outcome.

The repetition of the procedure with one hundred equally sized groups and fifty groups with an equal probability interval produces similar results. The only major differences are that in the former the model fits the data better also for draws and that in the latter market efficiency for draws cannot be rejected, even though the coefficients are still parallel to those in the other groupings. For convenience, the results of these two additional estimations are given in Tables B.1 and B.2 in Appendix B. As all the three groupings generate similar outcomes, the results can be considered robust and there is no need for repeating the analysis with other groupings.

9.1.2. Logit regressions: outcome-specific binary and all-inclusive multinomial

The results of the binary logit regression give further support to the rejection of statistical weak form betting market efficiency. As illustrated in Table 8, the joint hypothesis of market efficiency is rejected for each outcome and season, as well as for the whole sample period and all the outcomes simultaneously. Thus, also here subjective probabilities of soccer matches seem to deviate systematically from their objective probabilities. Moreover, when

Table 8

ML estimates and efficiency testing results of the binary logit model

This table presents the maximum likelihood estimates and the efficiency testing results of the binary logit model, as defined in Eq. (15). Robust standard errors of the estimated coefficients are reported in parentheses. In the case of the coefficient β, the null hypothesis is that $\beta = -1$. Wald's χ^2 refers to the Wald's chi-square test statistic for the estimated model, while $\chi_1^2(\alpha = 0; \beta = -1)$ stands for the test statistic for the Wald's test of coefficient restrictions, the null hypothesis of weak form efficiency being that $\alpha = 0$ and $\beta = -1$. Pseudo-R^2 refers to McFadden's pseudo R-squared. $\chi_6^2(\alpha = 0; \beta = -1)$ stands for the joint test that $\alpha = 0$ and $\beta = -1$ for home wins, draws, and away wins simultaneously. *** and ** indicate significance at the 1% and 5% levels, respectively.

	2009–10	**2010–11**	**2011–12**	**2012–13**	**2013–14**	**Total**
Home wins						
α	0.0595***	0.0712***	0.0385**	0.0569***	0.0642***	0.0583***
	(0.0160)	(0.0160)	(0.0161)	(0.0164)	(0.0164)	(0.0072)
β	–1.1257***	–1.1213***	–1.1382***	–1.0953***	–1.1201***	–1.1204***
	(0.0250)	(0.0250)	(0.0251)	(0.0260)	(0.0251)	(0.0113)
Wald's χ^2	2,020.85***	2,012.32***	2,061.05***	1,778.26***	1,983.57***	9,856.36***
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\chi_1^2(\alpha = 0; \beta = -1)$	31.12***	33.73***	31.18***	19.18***	29.55***	140.25***
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
pseudo-R^2	0.0912	0.0907	0.0949	0.0824	0.0971	0.0913
n	19,183	19,369	19,361	18,939	18,937	95,789
Draws						
α	0.2688***	0.2680***	0.2540***	0.3874***	0.4446***	0.3212***
	(0.0830)	(0.0818)	(0.0825)	(0.0924)	(0.0870)	(0.0380)
β	–1.2968***	–1.2953***	–1.2697***	–1.4055***	–1.4522***	–1.3409***
	(0.0810)	(0.0795)	(0.0800)	(0.0905)	(0.0839)	(0.0369)
Wald's χ^2	256.17***	265.74***	252.02***	241.27***	299.80***	1,320.11***
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\chi_1^2(\alpha = 0; \beta = -1)$	16.48***	17.01***	12.63***	21.52***	29.94***	94.66***
(p-value)	(0.000)	(0.000)	(0.002)	(0.000)	(0.000)	(0.000)
pseudo-R^2	0.0130	0.0139	0.0128	0.0136	0.0176	0.0141
n	19,183	19,369	19,361	18,939	18,937	95,789
Away wins						
α	0.0950***	0.0828***	0.1174***	0.0711***	0.0560**	0.0840***
	(0.0273)	(0.0273)	(0.0270)	(0.0272)	(0.0258)	(0.0120)
β	–1.1312***	–1.1345***	–1.1384***	–1.1158***	–1.1075***	–1.1251***
	(0.0262)	(0.0264)	(0.0260)	(0.0272)	(0.0254)	(0.0117)
Wald's χ^2	1,859.21***	1,846.37***	1,920.90***	1,687.19***	1,908.39***	9,2221.54***
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\chi_1^2(\alpha = 0; \beta = -1)$	25.51***	28.24***	28.50***	19.51***	20.36***	118.35***
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
pseudo-R^2	0.0976	0.0978	0.0990	0.0882	0.0994	0.0965
n	19,183	19,369	19,361	18,939	18,937	95,789
All outcomes						
$\chi_6^2(\alpha = 0; \beta = -1)$	73.11***	78.98***	72.31***	60.21***	79.85***	353.26***
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

inspecting the results of each season separately, we see no clear changes in the degree of inefficiency in the course of the sample period (excluding odds for draws that appear to have become even more inefficient during the two last seasons), which indicates that the betting market has not become more efficient during the more recent years. These results are strikingly consistent both with Koning (2012) who rejects weak form betting market efficiency while not observing any differences in the degree of efficiency between countries or seasons and with Nyberg (2014) who rejects efficiency for his whole sample as well as for the English Premier League separately. At the same time, however, the results contrast with Pope and Peel (1989) who find that odds are set in a weak form efficient manner for home wins and away wins but are anyway similar in the sense that odds are more devoid of information content for draws.

We also see that the coefficient β is significantly smaller than minus one in all cases. Within the binary logit model, this means that when subjective probability decreases, objective probability falls more than implied by market efficiency and vice versa. Therefore, well in line with Koning (2012) and Nyberg (2014), a favorite-longshot bias is found also with the binary logit regression; favorites win more often and longshots less frequently than indicated by the average betting market odds. Altogether, these results are identical to those of the linear regression presented above.

Portrayed in Table 9, statistical weak form betting market efficiency is rejected also when employing multinomial logit regression, both for the whole sample and for each season separately. In fact, as highlighted by the test statistics for the joint hypothesis of market efficiency, the multinomial logit model suggests that the betting market has become more inefficient during the latter part of the sample period. Based on the above results of binary logit regression, this change in the overall degree of inefficiency is most likely due to the even more biased draw odds during the seasons 2012–2014. In line with the rejection of weak form efficiency, the estimated coefficients show deviations from efficiency, consistent with Nyberg (2014). The coefficients β_1 and β_{-1} are significantly greater than one, which here denotes that as subjective probability of a home win or an away win rises, objective probability of the given outcome climbs more than implied by efficiency and vice versa. Thus, this finding is explicitly in line with the FLB found in the previous tests of statistical efficiency.

For robustness, instead of draws, I repeated the analysis using home wins and away wins as the benchmark category, respectively. The results of these alternative specifications, depicted in Tables B.3 and B.4 in Appendix B, are intuitively the same as those above in the sense that the market seems to be weak form inefficient when modeling all the outcomes in tandem, while the change in the degree of efficiency is induced by the draw odds.[43]

[43] In addition, to further test the validity of the results, I also performed a similar test applying the multinomial probit model that assumes the standard normal c.d.f. The results of this method are, as expected, uniform with those of the multinomial logit model and available upon request.

Table 9

ML estimates and efficiency testing results of the multinomial logit model with draws as the benchmark category

This table presents the maximum likelihood estimates and the efficiency testing results of the multinomial logit model, as defined in Eq. (20). Draws are used as the benchmark category. Robust standard errors of the estimated coefficients are reported in parentheses. In the case of the coefficients β_1 and β_{-1}, the null hypothesis is that $\beta_1 = \beta_{-1} = 1$. Wald's χ^2_4 refers to the Wald's chi-square test statistic for the estimated model, while $\chi^2_4(\alpha = 0; \beta = 1)$ stands for the test statistic for the Wald's test of coefficient restrictions, the null hypothesis of weak form efficiency being that $\alpha_j = 0$ and $\beta_j = 1$ for $j = -1, 1$. Pseudo-R^2 refers to McFadden's pseudo R-squared. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	2009–10	**2010–11**	**2011–12**	**2012–13**	**2013–14**	**Total**
α_1	–0.0510	–0.0499	–0.1175**	–0.1596***	–0.1841***	–0.1019***
	(0.0403)	(0.0407)	(0.0472)	(0.0542)	(0.0524)	(0.0204)
β_1	1.1780**	1.1924**	1.2765***	1.4171***	1.4490***	1.2832***
	(0.0852)	(0.0838)	(0.0998)	(0.1137)	(0.1057)	(0.0425)
α_{-1}	–0.0569	–0.0805	–0.0557	–0.2057***	–0.1705***	–0.1032***
	(0.0510)	(0.0501)	(0.0545)	(0.0591)	(0.0551)	(0.0239)
β_{-1}	1.2471**	1.2838***	1.2580**	1.5600***	1.4008***	1.3262***
	(0.1153)	(0.1099)	(0.1229)	(0.1307)	(0.1138)	(0.0523)
Wald's χ^2_4	2,498.37***	2,498.38***	2,574.00***	2,263.92***	2,538.24***	12,360.37***
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\chi^2_4(\alpha = 0; \beta = 1)$	9.10*	11.70**	12.77**	23.16***	21.86***	66.87***
(p-value)	(0.059)	(0.020)	(0.012)	(0.000)	(0.000)	(0.000)
pseudo-R^2	0.0745	0.0747	0.0765	0.0681	0.0786	0.0745
n	19,183	19,369	19,361	18,939	18,937	95,789

9.1.3. Grouping based on subjective probability

Thus far we have found powerful evidence of statistical weak form betting market inefficiency on the aggregate level, considering the whole spectrum of odds simultaneously. Consequently, when later conducting different betting strategies, it might be possible to obtain more accurate probability forecasts by utilizing the above regression results, especially those generated by logit regression that is a binary response model. The third and final part of the statistical tests of efficiency shifts from aggregate level to odds level, investigating whether the above evidence of a favorite-longshot bias translates into systematic deviations between subjective and objective probabilities on some individual odds levels.

Table 10 presents the results of the grouping analysis based on odds level for fifty groups with equal size, performed separately for each possible outcome. The results reveal statistically significant deviations from efficiency at both ends of the odds spectrum. At the one end with highest probability, i.e. for groups with lowest group number, objective probability tends to be systematically higher than subjective probability especially for home and away wins, as indicated by the negative z-values. In other words, average betting odds

Table 10

Subjective vs. objective probability: z-tests of 50 equally sized groups

This table presents the results of the z-tests that compare subjective probabilities ($\overline{\rho}_h$) of 50 odds groups with their objective probabilities ($\overline{\pi}_h$), as defined in Eq. (22) and Eq. (23). The groups are sorted by subjective probability so that the group size is 1,915 for 49 groups and 1,954 for the remaining group. The lower the group number, the higher the subjective probability and vice versa. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Home wins			***Draws***			***Away wins***		
h	$\overline{\rho}_h$	$\overline{\pi}_h$	**z-value**	$\overline{\rho}_h$	$\overline{\pi}_h$	**z-value**	$\overline{\rho}_h$	$\overline{\pi}_h$	**z-value**
1	0.824	0.897	–10.4***	0.334	0.349	–1.4	0.733	0.789	–5.9***
2	0.763	0.807	–4.8***	0.312	0.331	–1.8*	0.622	0.655	–3.1***
3	0.726	0.762	–3.7***	0.307	0.341	–3.2***	0.564	0.579	–1.4
4	0.695	0.723	–2.7***	0.303	0.304	–0.1	0.521	0.536	–1.3
5	0.671	0.696	–2.4**	0.301	0.305	–0.4	0.487	0.508	–1.8*
6	0.649	0.671	–2.0**	0.298	0.335	–3.4***	0.461	0.486	–2.2**
7	0.630	0.651	–1.9*	0.296	0.285	1.1	0.440	0.462	–2.0**
8	0.612	0.617	–0.4	0.294	0.283	1.1	0.421	0.422	–0.2
9	0.597	0.622	–2.3**	0.293	0.304	–1.1	0.405	0.409	–0.3
10	0.583	0.608	–2.2**	0.291	0.304	–1.2	0.392	0.399	–0.6
11	0.570	0.570	0.0	0.290	0.288	0.1	0.381	0.377	0.4
12	0.558	0.579	–1.9*	0.289	0.315	–2.5**	0.370	0.373	–0.3
13	0.546	0.554	–0.6	0.287	0.292	–0.4	0.360	0.351	0.8
14	0.536	0.551	–1.4	0.286	0.269	1.7*	0.351	0.360	–0.8
15	0.525	0.554	–2.5**	0.285	0.297	–1.1	0.343	0.359	–1.5
16	0.515	0.532	–1.5	0.284	0.290	–0.6	0.335	0.340	–0.5
17	0.505	0.512	–0.6	0.284	0.263	2.1**	0.328	0.339	–1.0
18	0.496	0.496	0.0	0.283	0.286	–0.3	0.322	0.304	1.7*
19	0.487	0.502	–1.3	0.282	0.285	–0.3	0.315	0.312	0.3
20	0.478	0.470	0.7	0.281	0.285	–0.4	0.309	0.297	1.2
21	0.471	0.477	–0.6	0.280	0.281	–0.1	0.303	0.289	1.4
22	0.463	0.465	–0.1	0.279	0.273	0.6	0.297	0.288	0.8
23	0.456	0.473	–1.5	0.278	0.261	1.7*	0.291	0.268	2.2**
24	0.448	0.460	–1.0	0.277	0.291	–1.4	0.285	0.281	0.4
25	0.441	0.447	–0.5	0.276	0.278	–0.3	0.278	0.255	2.3**
26	0.434	0.430	0.3	0.274	0.280	–0.6	0.272	0.267	0.5
27	0.426	0.449	–2.0**	0.273	0.274	–0.1	0.266	0.251	1.5
28	0.419	0.440	–1.8*	0.272	0.265	0.6	0.260	0.257	0.3
29	0.412	0.419	–0.6	0.270	0.254	1.7*	0.254	0.239	1.6
30	0.406	0.425	–1.7*	0.269	0.264	0.4	0.248	0.227	2.2**
31	0.400	0.389	1.0	0.267	0.264	0.3	0.242	0.237	0.5
32	0.394	0.419	–2.2**	0.265	0.232	3.4***	0.235	0.213	2.4**
33	0.387	0.373	1.3	0.263	0.245	1.9*	0.229	0.223	0.6
34	0.380	0.375	0.4	0.261	0.241	2.0**	0.222	0.197	2.7***
35	0.373	0.385	–1.0	0.259	0.248	1.1	0.214	0.192	2.5**
36	0.365	0.358	0.7	0.256	0.242	1.5	0.207	0.210	–0.3
37	0.357	0.354	0.3	0.253	0.257	–0.4	0.200	0.207	–0.8
38	0.348	0.353	–0.4	0.250	0.240	1.1	0.192	0.194	–0.2
39	0.339	0.348	–0.8	0.246	0.228	1.9*	0.185	0.173	1.3
40	0.329	0.305	2.3**	0.242	0.223	2.0**	0.176	0.151	3.1***
41	0.319	0.314	0.4	0.237	0.225	1.3	0.168	0.169	–0.1
42	0.307	0.280	2.6***	0.232	0.226	0.6	0.159	0.152	0.9
43	0.294	0.292	0.1	0.225	0.216	1.0	0.150	0.142	1.0
44	0.278	0.260	1.9*	0.218	0.212	0.6	0.141	0.137	0.5
45	0.260	0.260	0.1	0.209	0.185	2.7***	0.130	0.124	0.8
46	0.240	0.221	2.1**	0.199	0.183	1.8*	0.119	0.106	1.9*
47	0.215	0.209	0.6	0.187	0.153	4.1***	0.108	0.094	2.1**
48	0.185	0.178	0.8	0.171	0.146	3.1***	0.095	0.079	2.5**
49	0.148	0.143	0.7	0.151	0.133	2.4**	0.079	0.057	4.2***
50	0.094	0.073	3.6***	0.116	0.074	7.0***	0.056	0.027	7.8***

quoted in the market for strong favorites seem to be higher than what weak form market efficiency would imply. This anomaly is most extensive for home favorites. Therefore, from a betting strategy point of view, this result suggests that a (weak form) sharp bettor should exploit especially the excessively generous odds offered for favorites. At the same time, the opposite takes place at the other end of the odds spectrum. For groups with lowest probability, i.e. with highest group number, subjective probability appears to be systematically higher than objective probability, particularly for draws and away wins.

Due to the statistically significant and positive z-values, average market odds for the more obvious longshots seem to be lower than suggested by efficiency. Altogether, while the market shows to be efficient in between, the deviations found at both perimeters reflect a favorite-longshot bias observed already in the first two tests of statistical efficiency, but also provide insights on the shape of this bias. Most importantly, the relation between subjective and objective probability turns out not to be linear throughout the odds spectrum; instead, the deviations are borne particularly when the probabilities in question are high respective low. In other words, when comparing subjective probabilities with objective probabilities on a level-by-level basis, the results show that the betting market is statistically efficient in the middle part of the odds spectrum, whereas it is inefficient near the both ends of the spectrum.

Shown in Table 11, the results of fifty groups with an equal probability interval are on the whole similar to those of fifty groups with equal group size, still demonstrating a FLB as well as aggregate efficiency in the middle part of the odds spectrum. However, we can distinguish three differences. First, the deviations in probabilities assigned by the market for home and away favorites are now even more extensive, objective probability being higher than subjective probability for several groups. To be exact, we see that the average market odds seem to be biased in a statistically significant sense approximately down to the probability of 0.60. Thus, the betting market for home and away wins appears to be inefficient roughly up to the odds of 1.60. Second, an opposite deviation is in this case visible also for home longshots, in addition to the deviations for draw and away longshots found previously. Third, due to the very limited number of matches with high probabilities of a draw, the specification does not provide information on the equality of subjective and objective probability for this end of the draw odds spectrum.

Put together, the differences between the results of the two classifications point out that the average odds are upward (downward) biased up to (down to) a certain level, but that the number of matches in which this low (high) odds are available is somewhat limited. In European top soccer, namely, subjective probabilities of the majority of matches are relatively

Table 11

Subjective vs. objective probability: z-tests of 50 groups with equal interval

This table presents the results of the z-tests that compare subjective probabilities ($\overline{\rho}_h$) of 50 odds groups with their objective probabilities ($\overline{\pi}_h$), as defined in Eq. (22) and Eq. (23). The groups have been formed so that they have an equal probability interval. The lower the group number, the higher the subjective probability and vice versa. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Home wins			***Draws***			***Away wins***		
h	$\overline{\rho}_h$	$\overline{\pi}_h$	**z-value**	$\overline{\rho}_h$	$\overline{\pi}_h$	**z-value**	$\overline{\rho}_h$	$\overline{\pi}_h$	**z-value**
1	0.908	0.923	–0.2	0.74	1.00	n/a	0.90	0.86	0.3
2	0.892	0.973	–4.2***	0.73	1.00	n/a	0.88	1.00	n/a
3	0.874	0.969	–6.2***	0.71	1.00	n/a	0.86	0.96	–2.6***
4	0.857	0.915	–3.0***	n/a	n/a	n/a	0.85	0.96	–3.6***
5	0.839	0.935	–5.9***	n/a	n/a	n/a	0.83	0.94	–3.4***
6	0.821	0.906	–5.8***	0.67	1.00	n/a	0.81	0.87	–1.9*
7	0.803	0.869	–4.6***	n/a	n/a	n/a	0.79	0.85	–1.4
8	0.785	0.838	–3.7***	0.64	0.50	0.4	0.78	0.82	–1.3
9	0.768	0.817	–3.5***	0.63	0.00	n/a	0.76	0.80	–1.4
10	0.750	0.783	–2.5**	0.61	1.00	n/a	0.74	0.77	–1.0
11	0.732	0.761	–2.1**	0.60	0.33	1.0	0.72	0.80	–2.7***
12	0.714	0.754	–3.1***	0.59	0.00	n/a	0.70	0.78	–2.7***
13	0.696	0.725	–2.3**	0.58	0.00	n/a	0.69	0.74	–2.0**
14	0.678	0.702	–2.0**	0.56	0.50	0.2	0.67	0.71	–1.8*
15	0.660	0.688	–2.4**	0.54	1.00	n/a	0.65	0.72	–2.8***
16	0.642	0.667	–2.2**	0.53	0.33	0.7	0.63	0.66	–1.4
17	0.625	0.629	–0.4	0.52	0.67	–1.0	0.62	0.68	–3.2***
18	0.606	0.628	–2.1**	0.51	1.00	n/a	0.60	0.59	0.4
19	0.589	0.607	–1.9*	0.49	0.75	–1.2	0.58	0.57	0.3
20	0.571	0.577	–0.6	0.48	0.13	3.0***	0.56	0.58	–1.2
21	0.553	0.570	–1.9*	0.46	0.50	–0.2	0.55	0.58	–1.7*
22	0.535	0.553	–2.1**	0.45	0.63	–1.0	0.53	0.53	–0.4
23	0.517	0.535	–2.1**	0.43	0.25	1.2	0.51	0.53	–1.6
24	0.499	0.505	–0.7	0.42	0.63	–1.2	0.49	0.51	–1.5
25	0.481	0.484	–0.4	0.41	0.38	0.2	0.47	0.49	–1.0
26	0.463	0.472	–1.2	0.39	0.29	0.6	0.46	0.48	–1.9*
27	0.446	0.453	–1.1	0.38	0.33	0.4	0.44	0.46	–1.9*
28	0.428	0.443	–2.1**	0.37	0.42	–0.7	0.42	0.42	–0.2
29	0.410	0.420	–1.5	0.35	0.33	0.3	0.40	0.41	–0.3
30	0.392	0.393	–0.1	0.34	0.43	–1.5	0.38	0.38	0.0
31	0.374	0.378	–0.5	0.32	0.33	–0.4	0.37	0.37	–0.2
32	0.356	0.354	0.2	0.31	0.33	–1.9*	0.35	0.36	–1.1
33	0.338	0.336	0.3	0.30	0.31	–2.5**	0.33	0.34	–0.8
34	0.321	0.308	1.5	0.29	0.29	–1.0	0.31	0.30	1.9*
35	0.303	0.282	2.3**	0.27	0.27	1.3	0.30	0.28	2.3**
36	0.285	0.277	0.8	0.26	0.25	3.4***	0.28	0.27	1.6
37	0.267	0.255	1.2	0.25	0.23	2.9***	0.26	0.25	2.1**
38	0.249	0.247	0.2	0.23	0.22	1.9*	0.24	0.23	2.6***
39	0.231	0.203	2.6**	0.22	0.21	0.5	0.22	0.21	3.2***
40	0.213	0.206	0.7	0.20	0.18	3.4***	0.21	0.21	0.2
41	0.195	0.203	–0.6	0.19	0.15	4.5***	0.19	0.19	0.7
42	0.177	0.163	1.2	0.18	0.16	1.7*	0.17	0.16	2.4**
43	0.159	0.162	–0.2	0.16	0.14	2.0**	0.15	0.15	0.8
44	0.141	0.126	1.3	0.15	0.13	2.1**	0.14	0.13	1.1
45	0.124	0.102	1.9*	0.14	0.10	3.5***	0.12	0.11	1.9*
46	0.106	0.107	–0.1	0.12	0.08	3.6***	0.10	0.08	3.0***
47	0.088	0.066	2.0**	0.11	0.05	4.4***	0.08	0.06	3.5***
48	0.071	0.036	3.2***	0.09	0.05	3.2***	0.07	0.04	5.4***
49	0.053	0.014	4.0***	0.08	0.03	3.7***	0.05	0.01	9.5***
50	0.038	0.026	0.5	0.07	0.04	0.9	0.03	0.01	2.0**

even, hence not including any clear favorite and/or longshot. These findings both align and contradict with earlier results in soccer. Pope and Peel (1989) observe some evidence of a FLB in the tails of odds distributions for draws and away wins at the same time that Koning (2012) shows that more likely events tend to happen slightly more frequently than implied by odds in the case of home wins and away wins, holding across different countries. According to Kuypers (2000), in contrast, subjective probability appears to be a good estimator of objective probability except around the probability of 0.5 in his small sample.

The FLB is illustrated also in Fig. 1, which provides nonparametric lowess regressions of objective probability on subjective probability for home wins and away wins using one thousand odds groups with equal size. As explicated above, the results for draws are more erratic and therefore not presented in the figure. If subjective probability would be an unbiased estimator of objective probability, the regressions should lie on the grey 45-degree lines that in Fig. 1 symbolize statistical weak form betting market efficiency. However, we clearly see that events whose subjective probability is high tend to occur more frequently than suggested by efficiency, the deviation getting larger as subjective probability increases. The opposite holds for events with very low probability, even though the magnitude of the bias is smaller at this extreme of the odds spectrum. Otherwise the regressions for the two outcomes behave very similarly and there appears to be no bias in the middle of the odds spectrum.

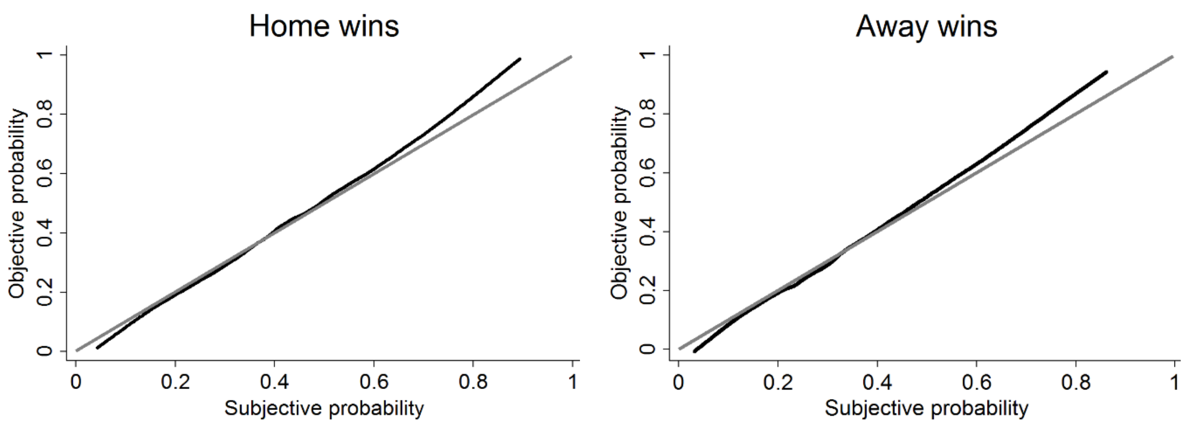

Fig. 1. Lowess regressions of objective probability on subjective probability. The regressions (black line) have been run separately for home wins and away wins using 1,000 odds groups with equal size. Statistical weak form betting market efficiency (grey line) implies that subjective probability equals objective probability throughout the odds spectrum.

For robustness, in addition to the two grouping specifications presented above, I performed a similar comparison also for twenty groups with equal size as well as for twenty groups with equal probability interval. The results of these specifications, which imply more within-group variation, provide no additional information to those portrayed above and are given in Tables B.5 and B.6 in Appendix B.

In summary, all the three tests of statistical efficiency give unanimous evidence of statistical weak form betting market inefficiency and therefore strong support to the rejection of Hypothesis I. In European soccer, 1X2 betting odds appear not to be unbiased estimators of match outcomes. Moreover, since all the tests detect a favorite-longshot bias, we also have stark evidence of a FLB consistent with the literature presented in Section 4.1. Next, I will turn the focus from probabilities assigned by average market odds to highest odds quoted in the market and go through the results of the economic tests of efficiency too see whether the betting market is economically weak form efficient, albeit not being statistically efficient.

9.2. Economic tests of weak form efficiency

This section displays the results of the four economic tests of efficiency introduced in Section 7.2, employing highest odds quoted in the market. Similar to the statistical tests, the economic tests begin with an aggregate test that exposes whether the betting market is economically weak form efficient on the aggregate level. The latter tests reveal the degree of efficiency on the individual odds level as well as explore the profitability of both naïve and sophisticated betting strategies.

9.2.1. Tobit regression: returns of unit bets on odds

Tobit regression does not reveal any signs of economic inefficiency and hence of any profitable betting strategies but, as a by-product, provides extra support to the existence of a favorite-longshot bias. The result is therefore consistent with Paton and Vaughan Williams (1998). The coefficients α in Table 12 are systematically below zero, which means that we cannot reject the null hypothesis of economic efficiency. In other words, returns of unit bets are negative throughout the odds spectrum, even when using the highest available odds, and the betting market seems to be economically weak form efficient on this aggregate level.

Table 12

ML estimates and efficiency testing results of the tobit model

This table presents the maximum likelihood estimates of the tobit model, as defined in Eq. (26). Robust standard errors of the estimated coefficients are reported in parentheses. In the case of the coefficient α, the null hypothesis is that $\alpha \leq 0$. Pseudo-R^2 refers to McFadden's pseudo R-squared. *** indicates significance at the 1% level.

	2009–10	**2010–11**	**2011–12**	**2012–13**	**2013–14**	**Total**
Home wins						
α	–0.3624	–0.3576	–0.4269	–0.6139	–1.1599	–0.6597
	(0.0707)	(0.0720)	(0.0679)	(0.0815)	(0.3529)	(0.1419)
β	–0.2767***	–0.2719***	–0.2713***	–0.2057***	–0.0554	–0.1855***
	(0.0256)	(0.0258)	(0.0239)	(0.0274)	(0.0896)	(0.0451)
F	116.91***	111.10***	128.41***	56.41***	0.38	16.88***
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.537)	(0.000)
pseudo-R^2	0.0083	0.0084	0.0087	0.0049	0.0006	0.0046
n	19,183	19,369	19,361	18,939	18,937	95,789
Draws						
α	–1.9686	–2.0352	–2.2376	–2.2318	–2.4486	–2.2112
	(0.2421)	(0.2390)	(0.2463)	(0.2840)	(0.2937)	(0.1220)
β	–0.5271***	–0.5179***	–0.4600***	–0.4593***	–0.4316***	–0.4719***
	(0.0586)	(0.0569)	(0.0578)	(0.0677)	(0.0660)	(0.0288)
F	80.80***	82.82***	63.35***	45.96***	42.71***	268.62***
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
pseudo-R^2	0.0029	0.0032	0.0029	0.0025	0.0032	0.0029
n	19,183	19,369	19,361	18,939	18,937	95,789
Away wins						
α	–2.2908	–2.4177	–2.1887	–2.0721	–2.0922	–2.2124
	(0.1688)	(0.1913)	(0.1741)	(0.1356)	(0.1532)	(0.0747)
β	–0.3754***	–0.3524***	–0.3547***	–0.3656***	–0.3490***	–0.3591***
	(0.0269)	(0.0304)	(0.0276)	(0.0231)	(0.0250)	(0.0121)
F	194.70***	134.24***	164.93***	251.07***	195.36***	881.90***
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
pseudo-R^2	0.0120	0.0111	0.0129	0.0120	0.0134	0.0123
n	19,183	19,369	19,361	18,939	18,937	95,789

However, while all negative, the differences in the degree of negative value for the coefficient α are worth noting; Table 12 shows that the returns are least negative for home wins and most negative for away wins. Thus, betting on home wins appears to be more favorable, although still unprofitable, than betting on draws or away wins.

At the same time, the coefficients β are statistically significantly negative for all outcomes during every season separately as well as during the whole sample period, except for home wins during the season 2013–2014. This gives an additional indication on the FLB earlier discovered in each statistical test of efficiency. An increase in odds (here represented by highest instead of average odds) decreases the expected return of unit bets and vice versa. Moreover, as demonstrated by the relatively stable values of the coefficient β for each outcome, there is no change in the degree of the bias over the course of the sample period.

9.2.2. Returns of odds groups

The analysis on the individual odds level does not expose considerable signs of economic weak form inefficiency. Tables 13 and 14 depict the returns of unit bets for fifty equally sized groups and fifty groups with an equal probability interval, respectively. Overall, home wins and away wins generate a marginally negative return, while placing unit bets on draws performs much worse. For all outcomes, the odds groups with highest subjective probability generate positive returns, whereas the groups with lowest probability perform worst, especially in the case of draws and away wins. Otherwise, along the odds spectrum outside the perimeters, the returns show no clear pattern and give no hints on potentially profitable betting strategies.

Table 13

Returns of unit bets of 50 equally sized groups

This table presents the returns of unit bets of 50 equally sized odds groups based on subjective probability, employing the one-sided t-test defined in Eq. (28). The lower the group number, the higher the subjective probability and vice versa. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Home wins				*Draws*				*Away wins*			
h	$\overline{R}_h$	h	$\overline{R}_h$	h	$\overline{R}_h$	h	$\overline{R}_h$	h	$\overline{R}_h$	h	$\overline{R}_h$
1	0.030***	26	–0.024	1	0.009	26	0.014	1	0.026**	26	0.014
2	0.006	27	0.039*	2	0.030	27	–0.001	2	0.015	27	–0.026
3	0.002	28	0.036*	3	0.082***	28	–0.030	3	–0.002	28	0.024
4	–0.003	29	0.004	4	–0.024	29	–0.065	4	0.003	29	–0.027
5	–0.005	30	0.029	5	–0.013	30	–0.021	5	0.022	30	–0.050
6	–0.006	31	–0.039	6	0.092***	31	–0.016	6	0.038*	31	0.023
7	–0.004	32	0.050**	7	–0.062	32	–0.125	7	0.037*	32	–0.050
8	–0.028	33	–0.048	8	–0.064	33	–0.071	8	–0.008	33	0.019
9	0.007	34	–0.024	9	0.012	34	–0.077	9	–0.001	34	–0.064
10	0.010	35	0.023	10	0.017	35	–0.040	10	0.007	35	–0.057
11	–0.031	36	–0.028	11	–0.030	36	–0.049	11	–0.020	36	0.071*
12	0.007	37	–0.011	12	0.068**	37	0.022	12	0.002	37	0.102**
13	–0.015	38	0.018	13	–0.008	38	–0.033	13	–0.030	38	0.075*
14	–0.000	39	0.034	14	–0.081	39	–0.058	14	0.025	39	0.003
15	0.025	40	–0.065	15	0.016	40	–0.060	15	0.049*	40	–0.081
16	0.006	41	0.005	16	–0.001	41	–0.029	16	0.022	41	0.084*
17	–0.011	42	–0.068	17	–0.092	42	–0.001	17	0.043*	42	0.035
18	–0.021	43	0.021	18	–0.006	43	–0.008	18	–0.043	43	0.029
19	0.009	44	–0.039	19	–0.005	44	0.004	19	0.005	44	0.066
20	–0.038	45	0.034	20	–0.003	45	–0.084	20	–0.023	45	0.048
21	–0.006	46	–0.044	21	–0.011	46	–0.045	21	–0.027	46	–0.008
22	–0.014	47	0.029	22	–0.036	47	–0.143	22	–0.006	47	–0.012
23	0.020	48	0.030	23	–0.071	48	–0.094	23	–0.052	48	–0.012
24	0.008	49	0.046	24	0.042	49	–0.052	24	0.015	49	–0.133
25	–0.002	50	–0.101	25	–0.002	50	–0.27	25	–0.052	50	–0.402
		Total	–0.003			Total	–0.030			Total	–0.006

Compared to the returns of equally sized groups, the returns of groups with an equal probability interval shown in Table 14 indicate more clearly that unit bets with higher subjective probability yield higher and often statistically significant positive returns, while those with lower subjective probability generate much worse returns. Thus, also this observation indicates a favorite-longshot bias, the phenomenon found in many sections of this study. What is interesting about the FLB here is that while it was discovered in the statistical tests of efficiency using average market odds, it is now discerned also using highest odds on the individual odds level.

The results of twenty groups with equal size and twenty groups with an equal probability interval are, by and large, similar to their respective counterparts with fifty groups. For convenience, the results of these specifications are given in Tables C.1 and C.2 in

Table 14

Returns of unit bets of 50 groups with equal interval

This table presents the returns of unit bets of 50 odds groups with an equal probability interval, employing the one-sided t-test defined in Eq. (28). The lower the group number, the higher the subjective probability and vice versa. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Home wins				*Draws*				*Away wins*			
h	$\bar{R}_h$	h	$\bar{R}_h$	h	$\bar{R}_h$	h	$\bar{R}_h$	h	$\bar{R}_h$	h	$\bar{R}_h$
1	–0.035	26	0.000	1	0.220**	26	–0.314	1	–0.121	26	0.040*
2	0.030*	27	0.001	2	0.285**	27	–0.117	2	0.059***	27	0.040*
3	0.047***	28	0.021	3	0.300	28	0.122	3	0.037	28	–0.007
4	0.008	29	0.010	4	n/a	29	–0.090	4	0.056*	29	–0.003
5	0.054***	30	–0.011	5	n/a	30	0.232*	5	0.065**	30	–0.010
6	0.042***	31	0.001	6	0.437***	31	–0.022	6	0.013	31	–0.003
7	0.025*	32	–0.005	7	n/a	32	0.014	7	0.014	32	0.024
8	0.012	33	0.000	8	–0.300	33	0.009	8	0.001	33	0.023
9	0.012	34	–0.022	9	–1.000	34	–0.011	9	0.006	34	–0.023
10	–0.003	35	–0.046	10	0.600	35	–0.023	10	–0.004	35	–0.024
11	–0.008	36	–0.002	11	–0.477	36	–0.052	11	0.051*	36	–0.003
12	0.011	37	–0.011	12	–1.000	37	–0.046	12	0.058*	37	–0.013
13	–0.001	38	0.030	13	–1.000	38	–0.027	13	0.038	38	–0.021
14	–0.008	39	–0.085	14	–0.160	39	0.016	14	0.023	39	–0.039
15	0.000	40	0.016	15	0.690**	40	–0.089	15	0.058**	40	0.054**
16	0.000	41	0.113**	16	–0.427	41	–0.158	16	0.009	41	0.043*
17	–0.030	42	–0.013	17	0.220	42	–0.034	17	0.070**	42	–0.010
18	0.001	43	0.104	18	0.922***	43	–0.053	18	–0.046	43	0.053*
19	–0.003	44	–0.024	19	0.488	44	–0.058	19	–0.041	44	0.047
20	–0.021	45	–0.092	20	–0.731	45	–0.189	20	0.007	45	0.016
21	0.001	46	0.147	21	–0.008	46	–0.238	21	0.027	46	–0.026
22	0.005	47	–0.070	22	0.381	47	–0.398	22	–0.010	47	–0.065
23	0.006	48	–0.369	23	–0.450	48	–0.335	23	0.024	48	–0.255
24	–0.011	49	–0.729	24	0.421	49	–0.564	24	0.027	49	–0.702
25	–0.015	50	1.374	25	–0.081	50	–0.149	25	0.007	50	–0.473
		Total	–0.003			Total	–0.030			Total	–0.006

Appendix C. Less surprisingly, the most notable difference is that the average returns of the specifications with 20 groups have less variation, due to the considerably larger group size. Moreover, the number of groups whose returns are statistically significantly positive is smaller. In any case, when considering the signs of the returns of these groupings displayed in Table 15, the occurrence of positive (negative) returns of the groups with highest (lowest) subjective probability is clearly visible. The difference in the number of pluses and minuses between the two methods in Table 15 also demonstrates that unit bets with highest subjective probability are more profitable but that there is only a limited number of them.

Table 15

Signs of unit bet returns of groupings with 20 groups

This table shows the combined signs of unit bet returns of the two odds grouping methods: equal group size and equal probability interval, for 20 groups. A given sign for a given group denotes that at least two (out of three) of the returns of the different outcomes (home wins, draws, and away wins) had that sign in that group. The lower the group number, the higher the subjective probability of a given outcome and vice versa.

	Group																			
Method	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**
Equal size	+	+	+	–	–	+	–	–	–	–	+	–	–	–	–	–	–	+	–	–
Equal interval	+	+	+	+	–	+	+	–	+	–	–	+	+	–	–	–	+	–	–	–

In terms of the favorite-longshot bias, these results are parallel to the majority of earlier relevant studies. The phenomenon in which bets on favorites generate higher returns than those on longshots has been observed at least to a certain extent by Cain et al. (2000; 2003), Malarić et al. (2008), and Direr (2013). Moreover, the above results shed light on the Cain et al. (2003) question whether the FLB occurs only at the high- and low-probability ends of the odds spectrum or whether we might expect to see a continuous decrease in returns when progressing towards longshots; based on Tables 13 and 14, the answer appears to be the former. On the other hand, the results contradict with two studies performed in English soccer. Deschamps and Gergaud (2007) find mixed evidence regarding the FLB, showing that there is a FLB for away wins, a reverse FLB for draws, and no bias at all for home wins, whereas Dixon and Pope (2004) detect a reverse FLB for all the outcomes.

In terms of economic efficiency, the results differ more from those of the earlier studies. Regardless of some biases, the absence of positive returns is discovered by Cain et al. (2000; 2003), Kuypers (2000), Dixon and Pope (2004), Deschamps and Gergaud (2007), and Malarić et al. (2008), while only Direr (2013) is able to demonstrate statistically significant positive returns when betting on clear favorites with odds less than 1.21, consistent with the results in Tables 13 and 14. Thus, when exploring the returns of unit bets of odds groups, this study indicates more signs of economic inefficiency than the earlier literature, which can be explained by the considerably higher odds available in the current online betting market.

9.2.3. Returns of naïve betting strategies

None of the naïve betting strategies yield statistically significant positive returns, though with a couple of mild exceptions. Nevertheless, the returns in Table 16 bring out some other interesting aspects. First, as suggested earlier, betting on favorites performs better than betting on longshots, consistent with Vlastakis et al. (2009) whose returns are anyway more negative than those in this study. In fact, when considering both the staking strategies, out of the five betting strategies betting on favorites performs best and betting on longshots worst. Second, betting on home wins produces higher returns than betting on the other outcomes (away wins with the unit wins staking strategy being the exception). These results are in line with Vlastakis et al. (2009) but conflict with Kuypers (2000) and Deschamps and Gergaud (2007) who find that betting on draws yields highest returns; however, also the highest returns in all these studies are negative and worse than those presented in Table 16. In addition, as indicated already by Tables 13 and 14, unit betting on all outcomes generates a slightly negative return, beating the return of –7% by Forrest et al. (2005) and the return of –18% by Kuypers (2000) but losing to the minimally positive return of 0.4% by Franck et al. (2013). In any case, in comparison with the earlier studies, the returns of the naïve betting strategies presented here are higher but do obviously not point to economic inefficiency.

Third, as demonstrated by the negative values of the Wilcoxon test statistic in all cases in Table 16, the staking strategy of unit wins yields consistently higher returns than that of unit bets, which suggests that it pays off to modify the stake according to the odds separately for each match. This feature is completely in line with Hvattum and Arntzen (2010) who also underline it as an additional evidence of a favorite-longshot bias, due to the fact that the staking strategy of unit wins places smaller bets when subjective probability is low.

Altogether, thus far all the results of the economic tests of efficiency indicate the absence of profitable betting strategies if employing only simple staking plans and not making any distinction between different matches and the odds relating to them. Thus, in search for economic inefficiencies, we need to apply more sophisticated strategies that involve diligent

Table 16

Returns of naïve betting strategies

This table presents the returns obtained with five different naïve betting strategies and two simple staking plans, employing the one-sided t-test defined in Eq. (29). *Unit bets* means betting a stake of one on every match, while *Unit wins* stands for determining the stake so that the potential winnings are equal to one. *Home wins*, *Draws*, and *Away wins* refer to placing a bet on the given outcome in every match in the sample, while *Favorite* and *Longshot* denote betting on the outcome with the lowest and highest odds, respectively, in every match. Besides reporting the returns, the table shows the z-values for the Wilcoxon signed-rank test that tests the equality of a given naïve strategy between the two staking plans, the null hypothesis being that both distributions are the same. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	2009–10	**2010–11**	**2011–12**	**2012–13**	**2013–14**	**Total**
Unit bets						
Home wins	–0.013	–0.006	–0.021	0.007	0.018*	–0.003
Draws	–0.040	–0.038	–0.025	–0.025	–0.020	–0.030
Away wins	–0.010	–0.016	0.004	–0.001	–0.006	–0.006
Favorite	0.001	–0.001	–0.007	0.008	0.000	0.000
Longshot	–0.023	–0.041	–0.006	–0.015	0.003	–0.016
Unit wins						
Home wins	–0.003	0.000	–0.006	0.002	0.004	–0.001
Draws	–0.009	–0.008	–0.005	–0.005	–0.003	–0.006
Away wins	–0.001	–0.002	0.005*	0.002	0.001	0.001
Favorite	0.000	0.000	–0.002	0.004	0.002	0.001
Longshot	–0.004	–0.009	0.003	–0.004	0.001	–0.003
Wilcoxon: unit bets vs. unit wins						
Home wins	–19.468***	–18.860***	–21.337***	–17.238***	–19.742***	–43.297***
Draws	–12.243***	–12.734***	–11.907***	–11.685***	–12.769***	–27.427***
Away wins	–31.978***	–31.553***	–29.055***	–27.695***	–29.452***	–66.979***
Favorite	–3.649***	–4.319***	–5.284***	–2.170**	–4.194***	–8.820***
Longshot	–28.745***	–30.638***	–26.563***	–26.997***	–26.997***	–62.578***

criteria to choose the matches and outcomes to bet, besides also using a more advanced staking technique—this is exactly what will be done in the remaining economic tests.

9.2.4. Returns of sophisticated betting strategies

The results of the sophisticated strategies are delivered in three parts. First, we go through the application that is based on quasi-arbitrage. Table 17 lists the number of quasi-arbitrage bets traced from the sample with different thresholds of the value betting edge. As expected, the larger the edge (i.e. the deviation between the highest and average market odds) the smaller the number of available quasi-arbitrage bets. Furthermore, on each threshold, the number of these bets increases season by season, reflecting the intensified competition and price war characteristic of the online betting market during the recent decade.

Table 17

Occurrence of quasi-arbitrage bets

This table displays the number of quasi-arbitrage bets available in the sample with six different thresholds of the value betting edge (r_j): 1.00, 1.05, 1.10, 1.20, 1.30, and 1.50. In other words, on a given threshold, it reveals the number of bets that were placed when following the quasi-arbitrage strategy introduced in Section 7.2.4.

	2009–10	**2010–11**	**2011–12**	**2012–13**	**2013–14**	**Total**
$r_j > 1.00$	16,026	16,524	17,119	17,345	17,660	84,674
$r_j > 1.05$	8,760	8,674	9,368	9,446	10,218	46,466
$r_j > 1.10$	3,789	3,730	4,373	4,445	4,921	21,258
$r_j > 1.20$	912	985	1,310	1,322	1,638	6,167
$r_j > 1.30$	313	389	619	600	791	2,712
$r_j > 1.50$	47	115	240	190	282	874

The annual returns of the quasi-arbitrage strategy support economic weak form efficiency. Available in Table 18, the majority of these returns converge to minus one and do certainly not generate positive returns. This is in a stark contrast to Paton and Vaughan Williams (2005) who are able to devise a consistently profitable strategy based on exploiting the outlier odds. To be exact, however, the context of Paton and Vaughan Williams is bookings spreads, not 1X2 odds, while their sample size is also tiny in relation to this study. The few exceptions in which the strategy yields profits include some higher thresholds during the season 2009–2010 and the most recent season with lower values of fractional Kelly. Moreover, considering the whole five-year period, the only statistically significant positive returns are made with the most conservative staking (when $g = 0.05$) and on the lowest edge thresholds (1.00, 1.05, and 1.10) that, as shown in Table 17, involve placing a greater number of bets. Overall, apart from these rare cases, the results imply that by following the quasi-arbitrage strategy the bettor would ultimately lose her wealth. The more aggressive the staking and/or the lower the edge threshold, the more negative the return. Conversely, the higher the threshold, the higher are the returns of the strategy, even though they are still negative; in this regard, the results align with Deschamps and Gergaud (2007).

As explained earlier, given that the probabilities implied by the average market odds are unbiased estimators of true probabilities, the quasi-arbitrage strategy should generate positive returns since it places only bets with a positive expected value. Thus, the obvious lack of profits perceived here implies that the average market odds are not such unbiased estimators, which is fully consistent with the findings of statistical efficiency presented in Section 9.1. More precisely, the absolute majority of the conceivable value bets that the strategy detects from the sample are either moderate or strong longshots, not favorites. Keeping in mind that we have earlier found a FLB, the quasi-arbitrage strategy assigns too

Table 18

Annual returns of the quasi-arbitrage strategy

This table presents the annual returns obtained with simulating the quasi-arbitrage strategy, as explained in Section 7.2.4, employing the one-sided t-test defined in Eq. (29). For each period, the simulation has been repeated 1,000 times, letting the sequence of matches during the given period to vary randomly. The simulations have been performed separately for each period, level of fractional Kelly staking (g), and threshold of the value betting edge (r_j). In terms of the fractional Kelly staking, four different values for g have been applied: 1.00, 0.50, 0.25, and 0.05. With respect to the threshold of the betting edge, six different levels have been employed: 1.00, 1.05, 1.10, 1.20, 1.30, and 1.50. *** and * indicate significance at the 1% and 10% levels, respectively.

	2009–10	**2010–11**	**2011–12**	**2012–13**	**2013–14**	**Total**
$g = 1.00$						
$r_j > 1.00$	–1.000	–1.000	–1.000	–1.000	–0.999	–1.000
$r_j > 1.05$	–1.000	–1.000	–1.000	–1.000	–0.996	–1.000
$r_j > 1.10$	–0.995	–0.999	–1.000	–1.000	–0.999	–1.000
$r_j > 1.20$	0.217***	–0.994	–1.000	–1.000	–0.995	–0.995
$r_j > 1.30$	0.222***	–0.775	–0.999	–0.999	–0.986	–0.978
$r_j > 1.50$	–0.196	–0.873	–0.968	–0.956	–0.690	–0.866
$g = 0.50$						
$r_j > 1.00$	–0.966	–0.969	–0.942	–0.998	0.732***	–0.954
$r_j > 1.05$	–0.950	–0.937	–0.934	–0.990	1.776***	–0.910
$r_j > 1.10$	–0.615	–0.878	–0.895	–0.980	–0.522	–0.863
$r_j > 1.20$	1.012***	–0.857	–0.957	–0.956	–0.699	–0.825
$r_j > 1.30$	0.443***	–0.335	–0.949	–0.944	–0.711	–0.761
$r_j > 1.50$	–0.053	–0.616	–0.782	–0.762	–0.094	–0.557
$g = 0.25$						
$r_j > 1.00$	–0.602	–0.625	–0.411	–0.893	3.092***	–0.478
$r_j > 1.05$	–0.552	–0.504	–0.420	–0.763	3.798***	–0.319
$r_j > 1.10$	–0.024	–0.454	–0.420	–0.740	0.537***	–0.342
$r_j > 1.20$	0.680***	–0.549	–0.733	–0.727	–0.143	–0.457
$r_j > 1.30$	0.298***	–0.097	–0.743	–0.731	–0.284	–0.434
$r_j > 1.50$	–0.010	–0.365	–0.503	–0.493	0.134***	–0.291
$g = 0.05$						
$r_j > 1.00$	–0.055	–0.068	0.044***	–0.254	0.615***	0.021***
$r_j > 1.05$	–0.044	–0.026	0.028***	–0.136	0.647***	0.064***
$r_j > 1.10$	0.074***	–0.044	–0.011	–0.154	0.258***	0.016***
$r_j > 1.20$	0.142***	–0.120	–0.198	–0.192	0.056***	–0.072
$r_j > 1.30$	0.068***	–0.002	–0.220	–0.213	–0.009	–0.083
$r_j > 1.50$	0.001*	–0.082	–0.120	–0.121	0.065***	–0.054

high probabilities for the longshots. Consequently, due to the FLB, the strategy is not able to trace value bets correctly, which explains the considerably negative returns in Table 18, in line with the discussion in Section 2.4.5 about the challenges of value betting if the estimations of true probabilities are somehow wrong.

Nevertheless, the results of the quasi-arbitrage strategy underline two important aspects. First, to increase the suitability and attractiveness of sports betting as an investment alternative and to decrease the risk of losing a significant part of one’s wealth if experiencing a longer sequence of losses and/or if for some reason obtaining inaccurate probability estimations, the results strongly recommend the use of conservative staking. Only the lowest level of fractional Kelly seem to provide any sensible balance between security and long term capital growth. Second, in search for profits with value betting, we should evidently find an alternative method to determine the true probabilities; luckily, this is precisely the idea that has been employed in the final economic tests of efficiency that utilize our logit model.

Similar to Table 17 for quasi-arbitrage bets, Table 19 provides the number of value bets when determining the true probabilities with the coefficients of our binary logit model presented in Section 9.1.2. Compared to those of quasi-arbitrage, the logit model discovers a somewhat smaller number of value bets, especially on the higher thresholds. In this respect, to obtain a sufficient number of bets and thereby reduce the impact of random variation, in practice it would not be advisable to employ the highest thresholds with the strategy.

Table 19

Occurrence of value bets discovered by the logit model

This table displays the number of value bets detected in the sample when determining the objective probabilities with the binary logit model. Similar to the above case of quasi-arbitrage, six thresholds of the value betting edge (r_j) have been used: 1.00, 1.05, 1.10, 1.20, 1.30, and 1.50. In other words, on a given threshold, it reveals the number of bets that were placed when following the value betting strategy based on the binary logit model.

	2009–10	**2010–11**	**2011–12**	**2012–13**	**2013–14**	**Total**
$r_j > 1.00$	14,143	14,411	14,979	15,899	16,324	75,756
$r_j > 1.05$	3,630	3,213	3,772	4,100	4,398	19,113
$r_j > 1.10$	395	521	841	925	1,054	3,736
$r_j > 1.20$	45	99	230	221	282	877
$r_j > 1.30$	21	52	121	103	152	449
$r_j > 1.50$	9	18	40	36	64	167

As the result of the second sophisticated strategy, the logit-based strategy performs notably better than the quasi-arbitrage strategy, but since all the specifications do not yield consistent positive returns, we cannot still fully reject economic weak form betting market efficiency. Portrayed in Table 20, for the whole five-year period as well as for the majority of the single seasons, the strategy generates statistically significant positive returns on all levels of fractional Kelly staking, except on the most aggressive one, and on the three lowest thresholds of the betting edge. For the thresholds higher than 1.10, the number of available value bets is already low, therefore involving more random variation. The substantial variation in the returns of the more aggressive staking plans is also worth noting, further accentuating the need for applying very conservative staking if using the strategy in real life. The sky-high returns of the season 2013–2014 stand out as an outlier but with an obvious

Table 20

Annual returns of the logit-based strategy

This table presents the annual returns obtained with simulating the value betting strategy based on the binary logit model, as explained in Section 7.2.4, employing the one-sided t-test defined in Eq. (29). For each period, the simulation has been repeated 1,000 times, letting the sequence of matches during the given period to vary randomly. The simulations have been performed separately for each period, level of fractional Kelly staking (g), and threshold of the value betting edge (r_j). In terms of the fractional Kelly staking, four different values for g have been applied: 1.00, 0.50, 0.25, and 0.05. With respect to the threshold of the betting edge, six different levels have been employed: 1.00, 1.05, 1.10, 1.20, 1.30, and 1.50. *** and * indicate significance at the 1% and 10% levels, respectively.

	2009–10	2010–11	2011–12	2012–13	2013–14	Total
$g = 1.00$						
$r_j > 1.00$	–0.994	–0.997	3.225***	–1.000	210.189***	–0.959
$r_j > 1.05$	–0.966	–0.987	2.145***	–0.998	8.207***	–0.880
$r_j > 1.10$	0.018***	1.348***	0.014***	–0.956	–0.695	–0.497
$r_j > 1.20$	–0.256	0.139***	–0.829	–0.734	0.731***	–0.418
$r_j > 1.30$	–0.256	–0.146	–0.755	–0.650	–0.088	–0.451
$r_j > 1.50$	–0.133	0.290***	–0.497	–0.461	0.710***	–0.123
$g = 0.50$						
$r_j > 1.00$	–0.447	–0.692	12.262***	–0.976	176.352***	0.579***
$r_j > 1.05$	–0.287	–0.657	4.762***	–0.826	14.434***	0.305***
$r_j > 1.10$	0.379***	1.108***	0.423***	–0.687	–0.045	0.043***
$r_j > 1.20$	–0.119	0.125***	–0.560	–0.439	0.662***	–0.164
$r_j > 1.30$	–0.137	–0.046	–0.497	–0.392	0.122***	–0.223
$r_j > 1.50$	–0.069	0.172***	–0.290	–0.265	0.500***	–0.031
$g = 0.25$						
$r_j > 1.00$	0.225***	–0.128	4.834***	–0.733	24.391***	1.115***
$r_j > 1.05$	0.187***	–0.223	2.242***	–0.411	5.025***	0.604***
$r_j > 1.10$	0.271***	0.579***	0.307***	–0.378	0.144***	0.133***
$r_j > 1.20$	–0.056	0.078***	–0.325	–0.233	0.392***	–0.060
$r_j > 1.30$	–0.071	–0.013	–0.287	–0.214	0.121***	–0.104
$r_j > 1.50$	–0.035	0.093***	–0.157	–0.142	0.288***	–0.004
$g = 0.05$						
$r_j > 1.00$	0.128***	0.046***	0.535***	–0.163	1.124***	0.264***
$r_j > 1.05$	0.093***	–0.005	0.328***	–0.049	0.539***	0.162***
$r_j > 1.10$	0.063***	0.111***	0.071***	–0.075	0.057***	0.043***
$r_j > 1.20$	–0.011	0.018***	–0.073	–0.048	0.085***	–0.007
$r_j > 1.30$	–0.015	–0.001	–0.065	–0.046	0.035***	–0.019
$r_j > 1.50$	–0.007	0.020***	–0.034	–0.030	0.063***	0.002*

explanation. When using the logit-based strategy, the season in question embodies substantially many winning bets that, together with aggressive staking and a very large number of bets to be placed (16,324 on the lowest edge threshold, for example), produces an extremely high return. On the other hand, during the season 2013–2014 even the specification according to which $g = 1.00$ and $r_j > 1.00$ implies a return of only 0.03% per bet. At the same time, conversely, the notably large number of losing bets accounts for the returns that are negative with all the specifications during the season 2012–2013.

Another aspect that catches our attention in the results portrayed in Table 20 (as well as in the results displayed below) concerns the fact that the absolute majority of the positive returns are also statistically significant at the 1% level. This can be justified by the way in which the simulations have been performed for each period, staking, and threshold. Even though the simulations have allowed the match order to vary randomly in each case, as explained in Section 7.2.4, the distribution of these returns was often narrow, denoting that the match order does eventually not play any key role in the formation of the returns. Thus, although the size of a Kelly stake is dependent on the size of the bankroll at a given moment, actual match outcomes (that naturally remained the same throughout the simulations) drive the returns to a very large extent when the total number of bets to be placed is large.

As the third and final sophisticated strategy, I replicated the logit-based strategy by utilizing the favorite-longshot bias and the apparent underpricing of favorite bets found in this study even more directly by placing bets only on favorites. I did this by scanning value bets with the logit model as above but only with the betting edge threshold of $r_j > 1.00$, thereafter excluding bets whose odds exceeded a given odds threshold and then performing the simulations. Table 21 lists the number of value bets detected by this advanced model. As expected, the number of bets that meet the criteria is much smaller than in the above specifications visible in Tables 17 and 19 for $r_j > 1.00$. In any case, according to Table 21, the odds thresholds of 1.30 and above are applicable for implementing this strategy.

When including only favorites, the sophisticated strategy founded on the logit model provides unanimous evidence of economic weak form betting market inefficiency. As portrayed in Table 22 (page 88), all the simulations that consider the whole sample period generate returns that reject the null hypothesis of unprofitability. The same holds for the absolute majority of the individual seasons as well. For example, when involving only clear favorites whose odds are not higher than 1.50, combined with the most conservative staking, the strategy yields an annual return of 7.6% during the five-year period, while being negative during none of the single seasons.

Moreover, similar to the earlier sophisticated strategies, the volatility of the returns is higher among the higher levels of fractional Kelly. At the same time, however, there is no clear trend with regard to the magnitude of the returns on different odds thresholds on each level of fractional Kelly. Based on the results, in practice it would be advisable to apply a specification of the strategy with an odds threshold not lower than 1.40 (to find an adequate number of value bets) and with a level of fractional Kelly not larger than 0.25 (to reduce the

Table 21

Occurrence of value bets discovered by the logit model with favorites only

This table displays the number of value bets detected in the sample when determining the objective probabilities with the binary logit model and including only bets for favorites whose odds (δ) are below a specific threshold. The table lists the figures for seven different odds thresholds: 1.70, 1.60, 1.50, 1.40, 1.30, 1.20, and 1.10. In other words, on a given threshold, it reveals the number of bets that were placed when following the value betting strategy that is based on the logit model and includes only favorites.

	2009–10	**2010–11**	**2011–12**	**2012–13**	**2013–14**	**Total**
$\delta < 1.70$	1,556	1,840	1,921	1,801	2,050	9,168
$\delta < 1.60$	1,125	1,294	1,372	1,257	1,441	6,489
$\delta < 1.50$	670	769	837	738	912	3,926
$\delta < 1.40$	314	385	421	371	462	1,953
$\delta < 1.30$	93	122	148	119	154	636
$\delta < 1.20$	5	6	19	21	22	73
$\delta < 1.10$	0	0	0	0	0	0

volatility of the returns). However, as this study aims only at exploring whether the betting market is economically weak form efficient when employing various strategies, specification of the optimal choice of staking and bets is out of scope here and left as a future exercise.

As explained in Section 7.2.4, there is a theoretical possibility that the returns of the logit-based sophisticated strategies in Tables 20 and 22 are subject to in-sample bias, even though the risk is not considerable in practice. The presumption of no bias is further supported by the results in Section 9.1.2 that show how the coefficients produced by the binary logit model for each season remain relatively stable throughout the sample period, indicating a permanent favorite-longshot bias. In any case, for robustness, I performed the corresponding out-of-sample simulations, using the seasons 2009–2012 as the in-sample period and then operating the seasons 2012–2014 as the out-of-sample period. As expected, the out-of-sample returns behave similarly to those above; in fact, these returns are even slightly higher than the equivalent in-sample returns portrayed in Tables 20 and 22.

Therefore, the results of the more advanced logit-based strategy appear to hold very well also out-of-sample. The strategy that employs value betting, Kelly staking, and a logit model that captures the statistical bias found in the betting market generates consistent positive returns, which allows us to reject Hypothesis II. For convenience, the out-of-sample results are exhibited in Appendix C, in Tables C.3 and C.4 for the basic logit-based strategy (corresponding to Tables 19 and 20 above), and in Tables C.5 and C.6 (being comparable to Tables 21 and 22 above) for the advanced logit-based strategy that includes only favorites.

Table 22

Annual returns of the logit-based strategy with favorites only

This table presents the annual returns obtained with simulating the value betting strategy based on the binary logit model including only bets for favorites whose odds are below a specific threshold, employing the one-sided t-test defined in Eq. (29). For each period, the simulation has been repeated 1,000 times, letting the sequence of matches during the given period to vary randomly. The simulations have been performed separately for each period, level of fractional Kelly staking (g), and odds threshold. In terms of the fractional Kelly staking, four different values for g have been applied: 1.00, 0.50, 0.25, and 0.05. With respect to the odds threshold, five different levels have been employed: 1.70, 1.60, 1.50, 1.40, and 1.30. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	2009–10	**2010–11**	**2011–12**	**2012–13**	**2013–14**	**Total**
$g = 1.00$						
$\delta < 1.70$	–0.859	–0.336	0.846***	–0.755	46.539***	0.150***
$\delta < 1.60$	–0.785	0.407***	0.586***	0.621***	97.812***	1.383***
$\delta < 1.50$	0.310***	–0.338	0.203***	3.257***	19.025***	1.454***
$\delta < 1.40$	1.844***	0.644***	0.642***	7.970***	3.773***	2.187***
$\delta < 1.30$	0.986***	–0.105	–0.220	1.380***	2.792***	0.658***
$g = 0.50$						
$\delta < 1.70$	–0.496	0.111***	0.894***	–0.247	10.722***	0.564***
$\delta < 1.60$	–0.434	0.471***	0.615***	0.692***	13.624***	1.016***
$\delta < 1.50$	0.274***	–0.078	0.266***	1.400***	4.709***	0.828***
$\delta < 1.40$	0.769***	0.370***	0.397***	2.292***	1.566***	0.956***
$\delta < 1.30$	0.421***	–0.031	–0.086	0.587***	1.030***	0.323***
$g = 0.25$						
$\delta < 1.70$	–0.238	0.137***	0.493***	–0.039	2.896***	0.371***
$\delta < 1.60$	–0.210	0.278***	0.350***	0.394***	3.199***	0.515***
$\delta < 1.50$	0.158***	–0.010	0.165***	0.607***	1.534***	0.403***
$\delta < 1.40$	0.346***	0.190***	0.207***	0.856***	0.665***	0.429***
$\delta < 1.30$	0.195***	–0.010	–0.036	0.268***	0.440***	0.158***
$g = 0.05$						
$\delta < 1.70$	–0.042	0.038***	0.097***	0.008**	0.340***	0.081***
$\delta < 1.60$	–0.039	0.059***	0.072***	0.080***	0.352***	0.098***
$\delta < 1.50$	0.034***	0.003*	0.037***	0.106***	0.216***	0.076***
$\delta < 1.40$	0.063***	0.038***	0.042***	0.136***	0.114***	0.078***
$\delta < 1.30$	0.037***	–0.001	–0.006	0.050***	0.077***	0.031***

When comparing the results of the sophisticated strategies based on the logit model with the earlier literature we must be careful because, to the best of my knowledge, there are no previous studies in soccer that would be directly comparable. In particular, Koning (2012) and Nyberg (2014), both of whom reject statistical weak form betting market efficiency with logit models similar to this study, do not utilize their models to explore economic efficiency at all. When considering the variety of sophisticated strategies more loosely and allowing also for semi-strong form studies, the results in the earlier literature mostly align to those in this study in the sense that substantial improvement in returns can be achieved compared to less sophisticated strategies, as demonstrated by Pope and Peel (1989), Dixon and Pope (2004), Vlastakis et al. (2009), and Franck et al. (2010).

At the same time, however, the previous studies show less consistency in terms of positive returns, i.e. in the rejection of economic efficiency. The strategies of Pope and Peel (1989) are not translatable into profitable post-tax profits (bettor taxes were the norm at that time). Vlastakis et al. (2009) suggest that econometric models, such as the multinomial logit model, can be employed to form profitable betting strategies, but in the context of 1X2 odds they cannot show positive returns. Nevertheless, Dixon and Pope (2004) and Franck et al. (2010) are able to generate marginal statistically significant profits with their advanced models for betting edge thresholds above 1.1 and 1.2, respectively, somewhat in line with the results presented in this chapter.

In summary, the four tests of economic efficiency give a more ambiguous view on efficiency than those of statistical efficiency. On the aggregate level, tobit regression shows no profit opportunities but proves that all bets do not yield the same expected return, implying a FLB. Neither the naïve betting strategies nor the quasi-arbitrage strategy demonstrate chances for profitable betting, whereas some of the odds groups involving clear favorites appear to generate moderate profits. In any case, the sophisticated strategy that takes into account the earlier observed favorite-longshot bias yields consistent profits, which means that we ultimately found a profitable strategy in European 1X2 soccer betting.

Altogether, since we have now rejected both of our hypotheses introduced in Chapter 5, finally including also the stricter hypothesis of economic efficiency, we may rather safely conclude that the online betting market in European soccer is weak form inefficient. The next chapter provides a discussion on these results.

10. Discussion

The results above clearly indicate that the online sports betting market in European soccer is weak form inefficient both statistically through an explicit favorite-longshot bias and economically when applying a sophisticated betting strategy that takes into account that bias. In the light of this study, the observed betting market inefficiency appears to be robust and resilient. Regardless of the developments in the current online betting market, the FLB has not vanished, which appears to give rise to profitable betting opportunities. Consequently, it is important to contemplate why the FLB still can persist, without sharp bettors completely exploiting the irrationality similarly to what sharp investors and arbitrageurs would do in financial markets to eliminate mispricing. I find two streams of intertwining explanations: institutional arrangements and market immaturity.

Regarding institutional arrangements, it might be that bookmakers' option to unilaterally limit the stakes of specific customers or to close their accounts entirely can at least hinder sharp bettors from exploiting the bias. As reminded by Franck et al. (2013), unlike betting exchanges, bookmakers are not trading against anonymous submitters of market orders; they are capable of and allowed to follow the trading history of each client and discriminate against skilled bettors, thereafter skewing the odds to optimize profits from the remaining, less skilled customer base. This is naturally in a stark contrast to financial marketplaces, but its real life significance is difficult to estimate. Franck et al. (2013, 321) also mention that this practice of "hidden price discrimination" has not been investigated thus far, but that it might have important consequences in the way that bookmakers' odds setting decisions should be modelled and potential market anomalies interpreted.

On the other hand, the importance of the explanation around institutional arrangements could be reduced in the future due to two reasons. First, if assuming that the population of bettors would become sharper, the size of the less sophisticated customer base would decrease, forcing bookmakers to alter their strategies to maintain and enhance revenues at the expense of profitability. In other words, to remain in the business, it might be that bookmakers' only option is to accept even sharp bettors. In this respect, the evolution of the bettor population along the continuum between sophistication and unsophistication plays a pivotal role. Second, even if this population would not become sharper, sophisticated bettors could anyway apply some specific methodologies to avoid becoming limited by bookmakers, or compromise their returns slightly by taking advantage of bookmakers that are less keen on limiting successful customers. Similar to the question around price discrimination, there is thus far no academic evidence of the degree of economic efficiency when employing betting strategies that dodge limitations or a subset of bookmakers that do not enforce stake limits.

The second stream of explanations for the persistence of market inefficiency not discerned in financial markets relates to the immaturity of betting markets. Compared to financial markets, present-day betting markets include more participants whose motives are less financial and whose stakes are relatively small. As emphasized by Paton and Vaughan Williams (2005), sports betting encompasses entertainment value that is significant for a large proportion of bettors, while the utility functions of financial traders are more dominated by wealth and risk considerations. Moreover, for the majority of bettors, a typical bet involves only a small proportion of wealth, whereas many financial decisions necessitate risking a much larger proportion. Taken together, these factors might lead to betting decisions that are not profit maximizing. Consequently, as formalized by Levitt (2004) and later discussed for example by Vlastakis et al. (2009) and Franck et al. (2010), bookmakers would set prices strategically (statistically less efficiently) to exploit bettor biases, as long as the institutional arrangements introduced above restrain sharp bettors from utilizing the biases too much.

At this point, it is challenging to give accurate estimations on how the population of bettors will evolve along the sophistication or maturity continuums. As sports betting markets are constantly expanding, most likely attracting mainly amateur players, it might well be that the proportion of the volume generated by these less sophisticated bettors will remain at least on the same level, if not increasing. Even if and when the pace of this expansion becomes slower, it might be that a significant proportion of the population of bettors will always have different utility functions than those of sophisticated bettors, meaning that betting markets would never reach maturity in the wealth maximizing sense. If this would be the case, the markets would remain inefficient and, in terms of exploiting the biases, bypassing the potential institutional arrangements becomes the key. If, on the other hand, betting markets would become more mature, signified by a situation in which supply matches demand by market forces with more efficient equilibrium prices (odds), performing even more subtle tests of efficiency and capitalizing on them—or completely switching over to other forms of making money in sports betting—would become the key.

Taken together, the focal implication of this study for finance is to show that even a competitive market that resembles any financial market, characterized by an ocean of market participants, considerable volumes, and low transaction costs, can be weak form inefficient under some specific conditions. By these conditions I refer to arrangements that prevent sophisticated investors from exploiting the biases and, at least partly due to these arrangements, to a situation in which a significant proportion of the total volume is brought about by less sophisticated investors. As such, these conditions differ from those under which rational arbitrageurs do not always eliminate sentimental mispricing in financial markets (see, e.g., DeLong et al., 1990; Shleifer and Vishny, 1997; Avery and Chevalier, 1999).

The results of this study contribute to the literature on weak form betting market efficiency in three dimensions. First, with a more comprehensive data set than in any previous soccer betting study in terms of the number of both matches and bookmakers considered, it is able to carry out more accurate odds level tests and explore the profitability of different betting strategies with a greater number of value bets, enabling more precise specifications as well as smaller random variation and more consistency in returns. In this respect, the study confirms the earlier findings around statistical weak form inefficiency that manifests itself in the favorite-longshot bias, but amplifies the literature regarding the location of the FLB on the odds spectrum. Moreover, compared to the existing literature, the study demonstrates more consistent economic weak form inefficiency by the existence of profitable betting strategies that utilize the FLB with a statistical model.

As the second empirical contribution, the study performs various tests of efficiency in the online betting era, which, in comparison to the betting market contexts in the majority of earlier studies, is characterized by an increased number of bookmakers, thinner bookmaker margins, and intensive competition. While the statistical biases found in the previous literature have often not resulted in economic inefficiency due to the higher bookmaker margins, the current study shows that similar biases in the more competitive online betting era lead also to economic inefficiency and consistent profits. Third and finally, while some of the otherwise noteworthy papers focus purely either on statistical or economic efficiency, this study covers both the forms of efficiency, which enables the detection of value bets—the ingredients of profitable betting strategies—both from statistical biases and quasi-arbitrage opportunities. In this respect, the study shows how the statistical inefficiencies can be profitably exploited, while the usefulness of quasi-arbitrage is minimal.

The study contains two main limitations that might influence the validity of the results: the assumption of the proportional spread of the bookmaker margin across 1X2 outcomes and the presumption of the absence of bookmaker stake limits. I see the former limitation as less serious, whereas I evaluate the seriousness of the latter limitation as moderate. First, wholly consistent with the practice of the discipline, it has been assumed that the bookmaker margin is spread proportionally across each possible outcome. This is necessary to be able to calculate the subjective probabilities assigned by the market, but if the margin would not be dispersed in the supposed way, the statistical tests of efficiency would be invalid due to inaccurate determination of the subjective probabilities. As long as bookmakers do not announce their policy and/or share the relating data, it is strictly speaking impossible to know for sure how the margin is distributed in reality. In this study, however, the significance of this limitation is diminished by two factors. First, the margins in the current online betting market are smaller than ever, which reduces the overall effect of the margin in probability calculations. Second, also the economic tests of efficiency, which do not need to convert odds to probabilities, indicate the existence of a bias similar to that found by the statistical tests, which supports the presumption that the subjective probabilities determined in the study are correct.

As the second (and somewhat more serious) limitation, the study assumes that the bettor is able to bet any given sum at the highest closing odds quoted in the market. However, because some bookmakers are known (and always allowed) to limit the stakes of their customers, it might be that the strategies performed in the study could have not been consistently implemented in reality. Once again, to be sure, either bookmakers should come out with their policies or we should check whether the bets suggested by the strategies could be placed in practice. Although this limitation might distort our results, it is probably reduced due to three factors. First, the current online betting market includes various bookmakers who publically accept sharp bettors. Second, sharp bettors can also use some specific techniques to minimize the possibility of being identified by bookmakers as such, without compromising the returns of the suggested strategies too much. Third, as the study rests on value betting rather than on arbitrage betting, the risk of becoming limited is less significant.

11. Conclusion

Online sports betting markets have expanded considerably during the recent decade. Meanwhile, due to their similarities with financial markets, the academic literature has viewed sports betting markets as an appropriate empirical setting for various tests of efficiency. Employing a more extensive data set than in any previous relevant study, this thesis has investigated weak form efficiency of online sports betting in European soccer from two perspectives: statistical and economic. Using three different methodologies, the statistical tests have inspected whether subjective probability implied by average market odds is an unbiased estimator of objective outcome probability. Economic tests, which represent the stricter tests of efficiency, have examined whether any betting strategy yields positive returns when utilizing highest odds quoted in the market. While the majority of the pertinent literature explore betting markets with more limited data sets in the pre-online era, this study contributes to the literature by carrying out an extensive analysis of both statistical and economic weak form efficiency in the online betting era with lower transaction costs.

The statistical tests find unanimous evidence of a favorite-longshot bias and therefore of statistical weak form betting market inefficiency. Both the linear and logit regression models, which investigate efficiency on the aggregate level, discover that as subjective probability increases, objective probability increases more than implied by efficiency. On the individual odds level, the method of sorting odds into groups based on subjective probability reveals that the departures between subjective and objective probability occur on both perimeters of the odds spectrum but not in the middle of it. In this respect, the betting market seems to be weak form inefficient at high respective low probabilities, whereas it appears to be efficient in between. In terms of statistical (in)efficiency, the results mostly align with those in earlier studies in soccer (Pope and Peel, 1989; Paton and Vaughan Williams, 1998; Cain et al., 2000, 2003; Deschamps and Gergaud, 2007; Vlastakis et al., 2009; Strumbelj and Sikonja, 2010; Koning, 2012; Direr, 2013; Nyberg, 2014) but also in other sports (see, e.g., Griffith, 1949; Dowie, 1976; Ali, 1977; Snyder, 1978; Asch et al., 1982, 1984; Vaughan Williams and Paton, 1997; Cain et al., 2003; Snowberg and Wolfers, 2010).

The economic tests give a more ambiguous picture of economic weak form betting market efficiency. When involving all the matches in the sample, neither the tobit regression model nor any of the naïve strategies show chances for profitable betting, whereas some of the odds groups with highest subjective probability display moderately positive returns. When only involving matches with a positive expected value and simulating the returns associated with them, the sophisticated strategy based on quasi-arbitrage is unprofitable, but the other sophisticated strategy that takes into account the favorite-longshot bias with a logit model generates consistent profits. For the whole sample period, a strategy specification that employs conservative staking and places bets only on clear favorites yields an annual return of 8%. Thus, ultimately, the betting market appears to be inefficient also in an economic sense. In terms of economic (in)efficiency, the results contradict with the majority of previous 1X2 soccer betting studies that find statistical but not economic weak form efficiency (Pope and Peel, 1989; Paton and Vaughan Williams, 1998; Cain et al., 2000, 2003; Deschamps and Gergaud, 2007; Vlastakis et al., 2009). Moreover, with connection to the more recent soccer studies that discover statistical inefficiency but do not investigate economic efficiency at all (Strumbelj and Sikonja, 2010; Koning, 2012; Nyberg, 2014), the results of this study show how these biases lead also to economic inefficiency.

Altogether, the study concludes that the European online sports betting market in soccer is weak form inefficient. From the statistical efficiency point of view, the results confirm the existence of a favorite-longshot bias commonly observed in the earlier literature, as well as shed more light on the location of this bias at both ends of the odds spectrum, instead of occurring monotonically throughout the spectrum. From the economic efficiency point of view, contrary to the previous literature, the study demonstrates how sophisticated strategies that exploit the FLB can now generate consistent profits, due to the considerably lower transaction costs in the online betting era.

The results therefore imply that the betting market has not become less inefficient during the five-year period of study, regardless of the intensified competition in terms of an increasing number of bookmakers and decreasing bookmaker margins. The study provides two intertwined explanations for the persistence of the observed inefficiencies. On the one hand, institutional arrangements might hinder sharp bettors from exploiting the biases. On the other hand, market immaturity refers to the effects that the less financial motives and smaller stakes of the significant part of the bettor population might have. More generally, the implication of this study is for finance is to portray how even a competitive market, featured by a wealth of market participants, considerable volumes, and low transaction costs, can be weak form inefficient under some distinctive circumstances. In any case, because sports betting markets are not yet as sophisticated as many financial markets, they provide attractive investment opportunities for sharp bettors.

12. Ideas on future research

Online sports betting constitutes an emerging, alluring, and dynamic field of study and plenty of interesting research topics stem also from this thesis. In this final chapter, I will use the opportunity to shortly propose three ideas on future research.

First, as online sports betting markets encompass various market structures, types of bets, sports, as well as time horizons, this study could be replicated in other contexts with other data sets by simply changing the focus in any of the dimensions presented in Section 2.2.2. While this study has focused on 1X2 pre-match bookmaker betting in European soccer, many other noteworthy betting markets are still to be investigated. We could hypothesize that the more immature a specific betting market is, the less statistically efficient it might be; on the other hand, the more mature the market is, the thinner the margins would most likely be, therefore providing higher odds and perhaps more signs of economic inefficiency.

Second, the context of this study could be employed with a slightly modified methodology. As the other of the two main limitations of the study relates to conceivable stake limits imposed by bookmakers, we could include only those bookmakers who publicly welcome winners and always allow large volumes, perhaps incorporating also betting exchanges. If we could then demonstrate similar deviations from efficiency, we would have even more waterproof evidence of weak form betting market inefficiency. At the same time, we could also investigate whether the opposite deviation at the other end of the odds spectrum could be profitably exploited utilizing betting exchanges. Since odds on low levels of subjective probability appear to be lower than implied by efficiency, these bets are overvalued and strategies based on selling these bets (i.e. betting against a given outcome to occur, which is available in betting exchanges) might well turn out to be profitable.

Third and lastly, as this thesis proposes that present-day sports betting markets provide attractive investment opportunities for sharp bettors, the idea of really considering these markets as an investment alternative among the more traditional alternatives deserves much more attention. While this study has merely inspected whether or not a single betting market in European soccer is weak form efficient, assuming risk neutrality, future research should put effort on determining the optimal betting strategies when bettors' individual risk appetite is taken into account. Since the degree of risk in the context of betting can be controlled by the choice of both staking and bets, both of these variables should be involved in the search for optimal strategies.

13. Appendices

13.1. Appendix A: Additional descriptive statistics

Table A.1

Countries and divisions in the sample

This table shows the number of soccer matches per season in each country and division considered in the study. The name of each division stands for its current name; due to name changes in some divisions, the name mentioned in the table might not correspond to the name during the sample period. All promotion and/or relegation matches in which teams from two adjacent divisions have taken part have been counted in to the number of matches in the higher of these two divisions.

Country and division	2009–10	2010–11	2011–12	2012–13	2013–14	Total
Armenia						
Premier League	108	104	106	160	108	586
Austria						
Tipico Bundesliga	174	174	176	174	168	866
Erste Liga	187	177	178	176	170	888
Azerbaijan						
Premier League	137	188	187	177	162	851
Belarus						
Vysshaya Liga	198	200	166	191	191	946
Belgium						
Jupiler League	278	310	307	309	306	1,510
Proximus League	337	299	300	303	295	1,534
Bosnia and Herzegovina						
Premier League	231	238	237	237	232	1,175
Bulgaria						
A PFG	212	239	231	218	238	1,138
Croatia						
1. HNL	238	238	231	197	178	1,082
Cyprus						
First Division	202	211	209	208	225	1,055
Czech Republic						
Synot Liga	236	239	229	232	239	1,175
Division 2	239	239	235	237	236	1,186
Denmark						
Superliga	197	198	191	191	190	967
Bet25 Liga	238	237	178	178	196	1,027

Country and division	2009–10	2010–11	2011–12	2012–13	2013–14	Total
England						
Premier League	373	371	374	362	348	1,828
Championship	554	548	552	553	547	2,754
League One	554	554	553	550	549	2,760
League Two	552	551	553	551	543	2,750
Estonia						
Meistriliiga	181	173	177	179	176	886
Faroe Islands						
Premier League	105	134	135	135	135	644
Finland						
Veikkausliiga	183	194	193	195	188	953
Ykkönen	176	148	131	125	135	715
France						
Ligue 1	368	374	369	366	361	1,838
Ligue 2	373	372	375	373	369	1,862
National	377	385	362	369	296	1,789
FYR of Macedonia						
First League	103	195	185	191	191	865
Germany						
Bundesliga	301	304	296	296	290	1,487
2. Bundesliga	301	303	301	304	305	1,514
3. Liga	375	379	371	371	366	1,862
Greece						
Super League	242	244	241	244	300	1,271
Hungary						
OTP Bank Liga	236	236	237	233	231	1,173

(Continued)

Table A.1 – *Continued*

Country and division	2009–10	2010–11	2011–12	2012–13	2013–14	Total
Iceland						
Pepsideild	128	128	123	126	129	634
Ireland						
Premier League	178	177	174	199	194	922
Division 1	196	165	110	110	106	687
Italy						
Serie A	366	360	366	368	360	1,820
Serie B	456	456	459	459	459	2,289
Lega Pro C1/A	305	294	307	279	244	1,429
Lega Pro C1/B	307	295	312	246	261	1,421
Kazakhstan						
Premier League	140	188	178	179	168	853
Latvia						
Virslīga	134	142	172	131	159	738
Lithuania						
A Lyga	134	186	160	119	151	750
Moldova						
Divizia Nationala	178	263	188	178	134	941
Montenegro						
Prva Crnogorska Liga	184	195	194	189	197	959
Netherlands						
Eredivisie	325	321	321	321	320	1,608
Eerste Divisie	331	300	298	277	371	1,577
Northern Ireland						
NIFL Premiership	221	226	228	224	222	1,121
Norway						
Tippeligaen	242	235	241	238	234	1,190
OBOS-ligaen	207	231	233	238	240	1,149
Poland						
Ekstraklasa	235	206	233	233	284	1,191
Division 1	303	296	305	291	299	1,494
Portugal						
Primeira Liga	235	235	236	236	236	1,178
Segunda Liga	238	236	237	454	450	1,615

Country and division	2009–10	2010–11	2011–12	2012–13	2013–14	Total
Romania						
Liga I	301	298	294	294	292	1,479
Russia						
Premier League	232	230	339	230	232	1,263
Division 1	356	377	469	267	310	1,779
Scotland						
Premiership	225	227	224	225	230	1,131
Championship	184	183	185	184	182	918
League One	184	186	185	184	184	923
Serbia						
Super Liga	240	239	234	230	228	1,171
Slovakia						
Fortuna liga	196	192	197	195	193	973
Slovenia						
Prva liga	171	178	177	178	178	882
Spain						
Primera Division	370	371	371	366	339	1,817
Segunda Division	457	455	454	451	459	2,276
Sweden						
Allsvenskan	240	238	240	231	225	1,174
Superettan	241	239	240	232	237	1,189
Switzerland						
Super League	178	178	163	174	172	865
Challenge League	228	226	231	176	176	1,037
Turkey						
Super Lig	270	298	321	297	294	1,480
PTT 1. Lig	307	273	307	297	315	1,499
Ukraina						
Pari-Match League	240	236	233	233	216	1,158
Wales						
Premier League	303	196	194	192	187	1,072
Europe						
Champions League	203	204	204	199	182	992
Europa League	448	454	458	424	424	2,208
Total	19,183	19,369	19,361	18,939	18,937	95,789

13.2. Appendix B: Results of statistical tests with alternative specifications

Table B.1

OLS estimates and efficiency testing results of 100 equally sized groups

This table presents the OLS estimates and the efficiency testing results of the standard linear regression model, as defined in Eq. (13), of 100 equally sized odds groups. The regressions have been run separately for home wins, draws, and away wins. The groups are sorted by subjective probability, as defined in Eq. (3), so that the group size is 956 for 99 groups and 1,145 for the remaining group. Robust standard errors of the estimated coefficients are reported in parentheses. In the case of the coefficient β, the null hypothesis is that $\beta = 1$. The sixth and seventh rows give the F-statistics for the joint test that $\alpha = 0$ and $\beta = 1$, which denotes the test of weak form betting market efficiency. *** indicates significance at the 1% level.

	Home wins	**Draws**	**Away wins**
α	–0.0337***	–0.0740***	–0.0301***
	(0.0049)	(0.0085)	(0.0033)
β	1.0928***	1.2644***	1.0929***
	(0.0104)	(0.0339)	(0.0102)
R^2	0.991	0.901	0.991
F	49.44***	46.03***	43.38***
(p-value)	(0.000)	(0.000)	(0.000)
n	100	100	100

Table B.2

OLS estimates and efficiency testing results of 50 groups with equal interval

This table presents the OLS estimates and the efficiency testing results of the standard linear regression model, as defined in Eq. (13), of 50 odds groups with variable size. The groups are sorted by subjective probability, as defined in Eq. (3), so that the interval of each group is of the same size in terms of subjective probability. The regressions have been run separately for home wins, draws, and away wins. Robust standard errors of the estimated coefficients are reported in parentheses. In the case of the coefficient β, the null hypothesis is that $\beta = 1$. The sixth and seventh rows give the F-statistics for the joint test that $\alpha = 0$ and $\beta = 1$, which denotes the test of weak form betting market efficiency. *** indicates significance at the 1% level.

	Home wins	**Draws**	**Away wins**
α	–0.0355***	–0.0463	–0.0355***
	(0.0048)	(0.0447)	(0.0065)
β	1.1055***	1.1127	1.1174***
	(0.0118)	(0.1952)	(0.0201)
R^2	0.997	0.460	0.993
F	41.67***	1.74	17.27***
(p-value)	(0.000)	(0.188)	(0.000)
n	50	50	50

Table B.3

ML estimates and efficiency testing results of the multinomial logit model with home wins as the benchmark category

This table presents the maximum likelihood estimates and the efficiency testing results of the multinomial logit model, as defined by Eq. (20). Home wins are used as the benchmark category. Robust standard errors of the estimated coefficients are reported in parentheses. In the case of the coefficients β_0 and β_{-1}, the null hypothesis is that $\beta_0 = \beta_{-1} = 1$. Wald's χ^2_4 refers to the Wald's chi-square test statistic for the estimated model, while $\chi^2_4(\alpha = 0; \beta = 1)$ stands for the test statistic for the Wald's test of coefficient restrictions, the null hypothesis of weak form efficiency being that $\alpha_j = 0$ and $\beta_j = 1$ for $j = -1, 0$. Pseudo-R^2 refers to McFadden's pseudo R-squared. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	2009-10	**2010-11**	**2011-12**	**2012-13**	**2013-14**	**Total**
α_0	0.0510	0.0499	0.1175**	0.1596***	0.1841***	0.1019***
	(0.0403)	(0.0407)	(0.0472)	(0.0542)	(0.0524)	(0.0204)
β_0	1.2371	1.2784	1.4372**	1.7948***	1.7976***	1.4658***
	(0.1755)	(0.1742)	(0.2110)	(0.2397)	(0.2220)	(0.0889)
α_{-1}	–0.0060	–0.0306	0.0618	–0.0461	0.0136	–0.0013
	(0.0473)	(0.0467)	(0.0509)	(0.0536)	(0.0511)	(0.0222)
β_{-1}	1.1881*	1.1978*	1.0973	1.1823	1.0522	1.1436***
	(0.1092)	(0.1049)	(0.1178)	(0.1239)	(0.1129)	(0.0505)
Wald's χ^2_4	2,498.37***	2,498.38***	2,574.00***	2,263.92***	2,538.24***	12,360.37***
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\chi^2_4(\alpha = 0; \beta = 1)$	15.76***	15.61***	24.31***	25.07***	22.65***	90.11***
(p-value)	(0.003)	(0.004)	(0.000)	(0.000)	(0.000)	(0.000)
pseudo-R^2	0.0745	0.0747	0.0765	0.0681	0.0786	0.0745
n	19,183	19,369	19,361	18,939	18,937	95,789

Table B.4

ML estimates and efficiency testing results of the multinomial logit model with away wins as the benchmark category

This table presents the maximum likelihood estimates and the efficiency testing results of the multinomial logit model, as defined by Eq. (20). Away wins are used as the benchmark category. Robust standard errors of the estimated coefficients are reported in parentheses. In the case of the coefficients β_1 and β_0, the null hypothesis is that $\beta_1 = \beta_0 = 1$. Wald's χ^2_4 refers to the Wald's chi-square test statistic for the estimated model, while $\chi^2_4(\alpha = 0; \beta = 1)$ stands for the test statistic for the Wald's test of coefficient restrictions, the null hypothesis of weak form efficiency being that $\alpha_j = 0$ and $\beta_j = 1$ for $j = 0, 1$. Pseudo-R^2 refers to McFadden's pseudo R-squared. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	2009-10	**2010-11**	**2011-12**	**2012-13**	**2013-14**	**Total**
α_1	0.0060	0.0306	–0.0618	0.0461	–0.0136	0.0013
	(0.0473)	(0.0467)	(0.0509)	(0.0536)	(0.0511)	(0.0222)
β_1	1.0645	1.0527	1.1640	0.9998	1.1361	1.0855*
	(0.1032)	(0.0993)	(0.1094)	(0.1118)	(0.1024)	(0.0469)
α_0	0.0569	0.0805	0.0557	0.2057***	0.1705***	0.1032***
	(0.0510)	(0.0501)	(0.0545)	(0.0591)	(0.0551)	(0.0239)
β_0	1.3606	1.4235**	1.3705	1.9773***	1.7136***	1.5239***
	(0.2242)	(0.2133)	(0.2377)	(0.2521)	(0.2226)	(0.1017)
Wald's χ^2_4	2,498.37***	2,498.38***	2,574.00***	2,263.92***	2,538.24***	12,360.37***
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\chi^2_4(\alpha = 0; \beta = 1)$	13.30***	20.07***	13.23**	28.62***	29.06***	93.98***
(p-value)	(0.010)	(0.001)	(0.010)	(0.000)	(0.000)	(0.000)
pseudo-R^2	0.0745	0.0747	0.0765	0.0681	0.0786	0.0745
n	19,183	19,369	19,361	18,939	18,937	95,789

Table B.5

Subjective vs. objective probability: z-tests of 20 equally sized groups

This table presents the results of the z-tests that compare subjective probabilities ($\overline{\rho}_h$) of 20 odds groups with their objective probabilities ($\overline{\pi}_h$), as defined in Eq. (22) and Eq. (23). The groups are sorted by subjective probability so that the group size is 4,789 for 19 groups and 4,798 for the remaining group. The lower the group number, the higher the subjective probability and vice versa. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Home wins			***Draws***			***Away wins***		
Group	$\overline{\rho}_h$	$\overline{\pi}_h$	**z-value**	$\overline{\rho}_h$	$\overline{\pi}_h$	**z-value**	$\overline{\rho}_h$	$\overline{\pi}_h$	**z-value**
1	0.78	0.83	–9.3***	0.32	0.33	–2.0**	0.66	0.70	–5.7***
2	0.69	0.72	–4.9***	0.30	0.32	–2.3**	0.51	0.53	–2.4**
3	0.64	0.65	–2.4**	0.30	0.31	–1.8*	0.45	0.47	–2.8***
4	0.59	0.62	–3.2***	0.29	0.30	–0.5	0.40	0.41	–0.5
5	0.56	0.57	–1.3	0.29	0.30	–2.0**	0.37	0.37	0.1
6	0.53	0.55	–2.7***	0.29	0.28	0.6	0.35	0.36	–0.9
7	0.51	0.52	–1.7*	0.28	0.28	0.8	0.33	0.33	–0.3
8	0.48	0.48	0.0	0.28	0.29	–0.6	0.31	0.30	1.3
9	0.46	0.47	–0.8	0.28	0.27	0.7	0.30	0.29	1.7*
10	0.45	0.46	–1.6	0.28	0.28	–0.4	0.28	0.27	2.7***
11	0.43	0.44	–1.7*	0.27	0.27	0.1	0.27	0.26	1.3
12	0.41	0.42	–1.8*	0.27	0.26	1.1	0.25	0.24	2.6***
13	0.39	0.40	–0.6	0.27	0.25	2.4**	0.24	0.23	1.4
14	0.38	0.38	0.1	0.26	0.24	3.1***	0.22	0.20	3.9***
15	0.36	0.36	0.5	0.25	0.25	0.8	0.20	0.21	–0.9
16	0.34	0.33	1.0	0.25	0.23	3.0***	0.18	0.17	2.9***
17	0.31	0.30	2.0**	0.23	0.22	1.8*	0.16	0.16	0.6
18	0.27	0.27	1.2	0.22	0.20	2.1**	0.14	0.13	1.4
19	0.22	0.21	1.7*	0.19	0.16	4.6***	0.11	0.10	3.1***
20	0.13	0.12	2.9***	0.14	0.11	6.7***	0.07	0.05	7.8***

Table B.6

Subjective vs. objective probability: z-tests of 20 groups with equal interval

This table presents the results of the z-tests that compare subjective probabilities ($\overline{\rho}_h$) of 20 odds groups with their objective probabilities ($\overline{\pi}_h$), as defined in Eq. (22) and Eq. (23). The groups have been formed so that they have an equal probability interval. The lower the group number, the higher the subjective probability and vice versa. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Home wins			***Draws***			***Away wins***		
Group	$\overline{\rho}_h$	$\overline{\pi}_h$	**z-value**	$\overline{\rho}_h$	$\overline{\pi}_h$	**z-value**	$\overline{\rho}_h$	$\overline{\pi}_h$	**z-value**
1	0.89	0.96	–4.8***	0.73	1.00	n/a	0.88	0.97	–2.6**
2	0.85	0.93	–7.4***	n/a	n/a	n/a	0.84	0.95	–5.1***
3	0.81	0.87	–7.3***	0.67	1.00	n/a	0.80	0.84	–2.1**
4	0.76	0.80	–4.9***	0.63	0.40	1.1	0.75	0.79	–2.1**
5	0.72	0.75	–3.9***	0.59	0.14	3.4***	0.71	0.77	–3.4***
6	0.67	0.70	–3.6***	0.56	0.60	–0.2	0.66	0.72	–3.9***
7	0.63	0.64	–2.1**	0.52	0.69	–1.5	0.62	0.66	–2.7***
8	0.58	0.60	–2.4**	0.48	0.38	0.7	0.58	0.58	–0.5
9	0.54	0.56	–3.5***	0.46	0.60	–1.1	0.53	0.55	–1.8*
10	0.49	0.50	–1.0	0.42	0.39	0.3	0.49	0.51	–2.1**
11	0.45	0.46	–1.3	0.38	0.37	0.1	0.44	0.46	–2.6**
12	0.40	0.42	–2.3**	0.35	0.39	–1.0	0.40	0.40	–0.1
13	0.36	0.36	–0.5	0.31	0.33	–3.4***	0.35	0.36	–1.1
14	0.32	0.30	2.9***	0.28	0.28	–0.2	0.31	0.30	2.5**
15	0.27	0.26	1.2	0.25	0.24	4.5***	0.26	0.25	3.1***
16	0.23	0.21	2.3**	0.22	0.20	3.1***	0.22	0.21	3.1***
17	0.18	0.18	0.1	0.18	0.15	4.9***	0.18	0.17	2.0**
18	0.14	0.12	1.9*	0.15	0.12	3.6***	0.13	0.12	2.3**
19	0.09	0.08	1.9*	0.11	0.07	6.0***	0.09	0.07	5.3***
20	0.06	0.03	3.5***	0.08	0.04	3.6***	0.05	0.02	9.0***

13.3. Appendix C: Results of economic tests with alternative specifications

Table C.1

Returns of unit bets of 20 equally sized groups

This table presents the returns of unit bets of 20 equally sized odds groups based on subjective probability, employing the one-sided t-test defined in Eq. (28). The lower the group number, the higher the subjective probability and vice versa. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Home wins				*Draws*				*Away wins*			
h	$\overline{R}_h$	h	$\overline{R}_h$	h	$\overline{R}_h$	h	$\overline{R}_h$	h	$\overline{R}_h$	h	$\overline{R}_h$
1	0.010*	11	0.015	1	0.011	11	–0.009	1	0.015*	11	0.003
2	0.002	12	0.017	2	0.022	12	–0.030	2	0.010	12	–0.030
3	–0.012	13	–0.003	3	0.013	13	–0.059	3	0.030**	13	0.008
4	0.003	14	–0.011	4	–0.015	14	–0.075	4	–0.001	14	–0.056
5	–0.013	15	–0.011	5	0.021	15	–0.012	5	–0.008	15	0.087***
6	0.008	16	–0.013	6	–0.035	16	–0.058	6	0.018	16	–0.022
7	–0.001	17	–0.021	7	–0.039	17	–0.020	7	0.015	17	0.061**
8	–0.021	18	0.004	8	–0.005	18	–0.024	8	–0.014	18	0.043
9	–0.006	19	0.009	9	–0.029	19	–0.090	9	–0.015	19	–0.002
10	0.009	20	–0.026	10	–0.001	20	–0.156	10	–0.033	20	–0.225
		Total	–0.003			Total	–0.030			Total	–0.006

Table C.2

Returns of unit bets of 20 groups with equal interval

This table presents the returns of unit bets of 20 odds groups with an equal probability interval, employing the one-sided t-test defined in Eq. (28). The lower the group number, the higher the subjective probability and vice versa. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Home wins				*Draws*				*Away wins*			
h	$\overline{R}_h$	h	$\overline{R}_h$	h	$\overline{R}_h$	h	$\overline{R}_h$	h	$\overline{R}_h$	h	$\overline{R}_h$
1	0.025*	11	–0.004	1	0.262***	11	–0.032	1	0.019	11	0.031**
2	0.036***	12	0.011	2	n/a	12	0.086	2	0.055**	12	–0.008
3	0.027***	13	0.003	3	0.437***	13	0.024*	3	0.001	13	0.015
4	0.005	14	–0.032	4	–0.400	14	–0.016	4	0.007	14	–0.016
5	–0.001	15	0.005	5	–0.776	15	–0.045	5	0.038*	15	–0.010
6	–0.002	16	–0.027	6	0.012	16	–0.034	6	0.049**	16	0.000
7	–0.014	17	0.067*	7	0.276	17	–0.103	7	0.021	17	0.030*
8	–0.008	18	–0.026	8	–0.231	18	–0.083	8	–0.021	18	0.030
9	0.006	19	–0.027	9	0.277	19	–0.306	9	0.009	19	–0.061
10	–0.012	20	–0.278	10	–0.087	20	–0.398	10	0.019	20	–0.528
		Total	–0.003			Total	–0.030			Total	–0.006

Table C.3

Occurrence of out-of-sample value bets discovered by the logit model

This table displays the number of out-of-sample value bets detected in the sample when determining the objective probabilities with the binary logit model. The in-sample period consists of seasons 2009–2012, followed by the out-of-sample seasons 2012–2014. Similar to the above case of quasi-arbitrage, six thresholds of the value betting edge (r_j) have been used: 1.00, 1.05, 1.10, 1.20, 1.30, and 1.50. In other words, on a given threshold, it reveals the number of out-of-sample bets that were placed when following the value betting strategy based on the binary logit model.

	2012–13	**2013–14**	**2012–14**
$r_j > 1.00$	15,941	16,397	32,338
$r_j > 1.05$	4,230	4,600	8,830
$r_j > 1.10$	909	1,034	1,943
$r_j > 1.20$	202	259	461
$r_j > 1.30$	94	142	236
$r_j > 1.50$	34	57	91

Table C.4

Annual out-of-sample returns of the logit-based strategy

This table presents the annual out-of-sample returns obtained with simulating the value betting strategy based on the binary logit model, as explained in Section 7.2.4, employing the one-sided t-test defined in Eq. (29). The in-sample period consists of seasons 2009–2012, followed by the out-of-sample seasons 2012–2014. For each period, the simulation has been repeated 1,000 times, letting the sequence of matches during the given period to vary randomly. The simulations have been performed separately for each period, level of fractional Kelly staking (g), and threshold of the value betting edge (r_j). In terms of the fractional Kelly staking, four different values for g have been applied: 1.00, 0.50, 0.25, and 0.05. With respect to the threshold of the betting edge, six different levels have been employed: 1.00, 1.05, 1.10, 1.20, 1.30, and 1.50. *** indicates significance at the 1% level.

	2012–13	**2013–14**	**2012–14**
$g = 1.00$			
$r_j > 1.00$	–1.000	497.851***	–0.897
$r_j > 1.05$	–0.999	56.837***	–0.728
$r_j > 1.10$	–0.977	0.417***	–0.821
$r_j > 1.20$	–0.699	0.608***	–0.304
$r_j > 1.30$	–0.593	0.077***	–0.338
$r_j > 1.50$	–0.430	0.841***	0.025***
$g = 0.50$			
$r_j > 1.00$	–0.954	347.735***	2.991***
$r_j > 1.05$	–0.844	46.072***	1.711***
$r_j > 1.10$	–0.773	1.122***	–0.306
$r_j > 1.20$	–0.410	0.576***	–0.036
$r_j > 1.30$	–0.345	0.214***	–0.109
$r_j > 1.50$	–0.244	0.547***	0.081***
$g = 0.25$			
$r_j > 1.00$	–0.619	36.774***	2.792***
$r_j > 1.05$	–0.426	10.029***	1.516***
$r_j > 1.10$	–0.470	0.717***	–0.046
$r_j > 1.20$	–0.217	0.348***	0.028***
$r_j > 1.30$	–0.185	0.163***	–0.026
$r_j > 1.50$	–0.130	0.303***	0.065***
$g = 0.05$			
$r_j > 1.00$	–0.095	1.321***	0.449***
$r_j > 1.05$	–0.050	0.749***	0.289***
$r_j > 1.10$	–0.104	0.147***	0.014***
$r_j > 1.20$	–0.044	0.077***	0.014***
$r_j > 1.30$	–0.039	0.042***	0.001
$r_j > 1.50$	–0.027	0.065***	0.018***

Table C.5

Occurrence of out-of-sample value bets discovered by the logit model with favorites only

This table displays the number of out-of-sample value bets detected in the sample when determining the objective probabilities with the binary logit model and including only bets for favorites whose odds (δ) are below a specific threshold. The in-sample period consists of seasons 2009–2012, followed by the out-of-sample seasons 2012–2014. The table lists the figures for seven different odds thresholds: 1.70, 1.60, 1.50, 1.40, 1.30, 1.20, and 1.10. In other words, on a given threshold, it reveals the number of out-of-sample bets that were placed when following the value betting strategy that is based on the logit model and includes only favorites.

	2012–13	**2013–14**	**2012–14**
δ < 1.70	1,917	2,198	4,115
δ < 1.60	1,353	1,559	2,912
δ < 1.50	808	1,005	1,813
δ < 1.40	426	536	962
δ < 1.30	149	199	348
δ < 1.20	26	29	55
δ < 1.10	0	0	0

Table C.6

Annual out-of-sample returns of the logit-based strategy with favorites only

This table presents the annual out-of-sample returns obtained with simulating the value betting strategy based on the binary logit model including only bets for favorites whose odds are below a specific threshold, employing the one-sided t-test defined in Eq. (29). The in-sample period consists of seasons 2009–2012, followed by the out-of-sample seasons 2012–2014. For each period, the simulation has been repeated 1,000 times, letting the sequence of matches during the given period to vary randomly. The simulations have been performed separately for each period, level of fractional Kelly staking (g), and odds threshold. In terms of the fractional Kelly staking, four different values for g have been applied: 1.00, 0.50, 0.25, and 0.05. With respect to the odds threshold, five different levels have been employed: 1.70, 1.60, 1.50, 1.40, and 1.30. *** indicates significance at the 1% level.

	2012–13	**2013–14**	**2012–14**
g = 1.00			
δ < 1.70	–0.596	33.406***	2.730***
δ < 1.60	1.811***	85.053***	14.553***
δ < 1.50	5.576***	16.25***	9.651***
δ < 1.40	11.360***	3.242***	6.241***
δ < 1.30	2.266***	2.694***	2.473***
g = 0.50			
δ < 1.70	0.039***	10.208***	2.412***
δ < 1.60	1.342***	13.937***	4.914***
δ < 1.50	2.079***	4.653***	3.172***
δ < 1.40	2.952***	1.537***	2.166***
δ < 1.30	0.872***	1.045***	0.957***
g = 0.25			
δ < 1.70	0.148***	2.918***	1.121***
δ < 1.60	0.659***	3.336***	1.682***
δ < 1.50	0.834***	1.560***	1.167***
δ < 1.40	1.045***	0.674***	0.850***
δ < 1.30	0.380***	0.452***	0.415***
g = 0.05			
δ < 1.70	0.048***	0.347***	0.188***
δ < 1.60	0.121***	0.365***	0.237***
δ < 1.50	0.137***	0.221***	0.178***
δ < 1.40	0.159***	0.117***	0.138***
δ < 1.30	0.068***	0.080***	0.074***

14. References

Aharony, J., Swary, I., 1980. Quarterly dividend and earnings announcements and stockholders' returns: an empirical analysis. Journal of Finance 35, 1–12.

Alberola, J.M., Garcia-Fornes, A., 2013. Using a case-based reasoning approach for trading in sports betting markets. Applied Intelligence 38, 465–477.

Algoet, P.H., Cover, T.M., 1988. Asymptotic optimality and asymptotic equipartition properties of log-optimum investment. Annals of Probability 16, 876–898.

Ali, M.M., 1977. Probability and utility estimates for racetrack bettors. Journal of Political Economy 85, 803–815.

Ali, M.M., 1979. Some evidence of the efficiency of a speculative market. Econometrica 47, 387–392.

Ali, M.M., Thalheimer, R., 1997. Transportation costs and product demand: wagering on parimutuel horse racing. Applied Economics 29, 529–542.

Asch, P., Malkiel, B.G., Quandt, R.E., 1982. Racetrack betting and informed behavior. Journal of Financial Economics 10, 187–194.

Asch, P., Malkiel, B.G., Quandt, R.E., 1984. Market efficiency in racetrack betting. Journal of Business 57, 165–175.

Asch, P., Quandt, R.E., 1987. Efficiency and profitability in exotic bets. Economica 54, 289–298.

Asquith, P., Mullins, D.W.J., 1983. The impact of initiating dividend payments on shareholders' wealth. Journal of Business 56, 77–96.

Avery, C., Chevalier, J., 1999. Identifying investor sentiment from price paths: the case of football betting. Journal of Business 72, 493–521.

Bachelier, L., 1900. Théorie de la speculation. Annales scientifiques de l'École Normale Supérieure 3, 21–86.

Ball, R., Brown, P., 1968. An empirical evaluation of accounting income numbers. Journal of Accounting Research 6, 159–178.

Banks, J., 2013. Edging your bets: advantage play, gambling, crime and victimization. Crime, Media, Culture 9, 171–187.

Bernard, V.L., Thomas, J.K., 1990. Evidence that stock prices do not fully reflect the implications of current earnings for future earnings. Journal of Accounting and Economics 13, 305–340.

Bernoulli, D., 1738. Specimen theoriae novae de mensura sortis. Commentarii Academiae Scientiarum Imperialis Petropolitanae 5, 175–192. Translated by L. Sommer as: Bernoulli, D., 1954. Exposition of a new theory on the measurement of risk. Econometrica 22, 23–36.

Borghesi, R., 2012. The implications of a reverse favourite-longshot bias in a prediction market. Journal of Prediction Markets 6, 12–21.

Boulier, B.L., Stekler, H.O., 2003. Predicting the outcomes of National Football League games. International Journal of Forecasting 19, 257–270.

Breiman, L., 1961. Optimal gambling systems for favorable games. In: Neyman, J. (Ed.), Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics. University of California Press, Berkeley, pp. 65–78.

Brier, G.W., 1950. Verification of forecasts expressed in terms of probability. Monthly Weather Review 78, 1–3.

Brown, W.O., Sauer, R.D., 1993. Fundamentals or noise? Evidence from the professional basketball betting market. Journal of Finance 48, 1193–1209.

Bruce, A.C., Johnson, J.E.V., 1992. Toward an explanation of betting as a leisure pursuit. Leisure Studies 11, 201–218.

Buchdal, J., 2003. Fixed Odds Sports Betting: the Essential Guide. High Stakes, London.

Busche, K., 1994. Efficient market results in an Asian setting. In: Hausch, D.B., Lo, V.S., Ziemba, W.T. (Ed.), Efficiency of Racetrack Betting Markets. Academic Press, New York, pp. 615–616.

Busche, K., Hall, C.D., 1988. An exception to the risk preference anomaly. Journal of Business 61, 337–346.

Cain, M., Law, D., Peel, D., 2000. The favourite-longshot bias and market efficiency in UK football betting. Scottish Journal of Political Economy 47, 25–36.

Cain, M., Law, D., Peel, D., 2003. The favourite-longshot bias, bookmaker margins and insider trading in a variety of betting markets. Bulletin of Economic Research 55, 263–273.

Caruso, R., 2009. The basic economics of match fixing in sport tournaments. Economic Analysis and Policy 39, 355–377.

Charest, G., 1978. Dividend information, stock returns and market efficiency—II. Journal of Financial Economics 6, 297–330.

Church-Sanders, R., 2011. Online Sports Betting: a Market Assessment and Outlook. iGaming Business, London.

Clemen, R.T., Winkler, R.L., 1986. Combining economic forecasts. Journal of Business and Economic Statistics 4, 39–46.

Coleman, L., 2004. New light on the longshot bias. Applied Economics 36, 315–326.

Conlisk, J., 1993. The utility of gambling. Journal of Risk and Uncertainty 6, 255–275.

Constantinou, A.C., Fenton, N.E., 2012. Solving the problem of inadequate scoring rules for assessing probabilistic football forecast models. Journal of Quantitative Analysis in Sports 8, article 1.

Constantinou, A.C., Fenton, N.E., 2013. Profiting from arbitrage and odds biases of the European football gambling market. Journal of Gambling Business and Economics 7, 41–70.

Constantinou, A.C., Fenton, N.E., Neil, M., 2012. pi-football: a Bayesian network model for forecasting association football match outcomes. Knowledge-Based Systems 36, 322–339.

Constantinou, A.C., Fenton, N.E., Neil, M., 2013. Profiting from an inefficient association football gambling market: prediction, risk and uncertainty using Bayesian networks. Knowledge-Based Systems 50, 60–86.

Cowles, A. 3rd, 1933. Can stock market forecasters forecast? Econometrica 1, 309–324.

Croxson, K., Reade, J.J., 2011. Exchange vs. dealers: a high-frequency analysis of in-play betting prices. Department of Economics Discussion Paper 11–19. University of Birmingham, Birmingham.

Croxson, K., Reade, J.J., 2014. Information and efficiency: goal arrival in soccer betting. Economic Journal 124, 62–91.

Dare, W.H., MacDonald, S.S., 1996. A generalized model for testing the home and favorite team advantage in point spread markets. Journal of Financial Economics 40, 295–318.

De Long, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J., 1990. Noise trader risk in financial markets. Journal of Political Economy 98, 703–738.

Demir, E., Danis, H., Rigoni, U., 2012. Is the soccer betting market efficient? A cross-country investigation using the Fibonacci strategy. Journal of Gambling Business and Economics 6, 29–49.

Deschamps, B., Gergaud, O., 2007. Efficiency in betting markets: evidence from English football. Journal of Prediction Markets 1, 61–73.

Direr, A., 2013. Are betting markets efficient? Evidence from European football championships. Applied Economics 45, 343–356.

Dixon, M.J., Coles, S.G., 1997. Modelling association football scores and inefficiencies in the football betting market. Journal of the Royal Statistical Society: Series C (Applied Statistics) 46, 265–280.

Dixon, M.J., Pope, P.F., 2004. The value of statistical forecasts in the UK association football betting market. International Journal of Forecasting 20, 697–711.

Dowie, J., 1976. On the efficiency and equity of betting markets. Economica 43, 139–150.

Dunning, E., 1999. The development of soccer as a world game. In: Dunning, E. (Ed.), Sport Matters: Sociological Studies of Sport, Violence and Civilisation. Routledge, London, pp. 80–105.

Durham, G.R., Hertzel, M.G., Martin, J.S., 2005. The market impact of trends and sequences in performance: new evidence. Journal of Finance 60, 2551–2569.

Edelman, D.C., O'Brian, N.R., 2004. Tote arbitrage and lock opportunities in racetrack betting. European Journal of Finance 10, 370–378.

Epstein, E.S., 1969. A scoring system for probability forecasts of ranked categories. Journal of Applied Meteorology 8, 985–987.

Fama, E.F., 1970. Efficient capital markets: a review of theory and empirical work. Journal of Finance 25, 383–417.

Fama, E.F., 1991. Efficient capital markets: II. Journal of Finance 46, 1575–1617.

Fama, E.F., 1998. Market efficiency, long-term returns, and behavioral finance. Journal of Financial Economics 49, 283–306.

Figlewski, S., 1979. Subjective information and market efficiency in a betting market. Journal of Political Economy 87, 75–88.

Finnigan, M., Nordsted, P., 2010. The Premier Football Betting Handbook 2010/11. Harriman House, Hampshire.

Fomby, T.B., Samanta, S.K., 1991. Application of Stein rules to combination forecasting. Journal of Business and Economic Statistics 9, 391–407.

Forrest, D., Goddard, J., Simmons, R., 2005. Odds-setters as forecasters: the case of English football. International Journal of Forecasting 21, 551–564.

Forrest, D., Simmons, R., 2000. Forecasting sport: the behaviour and performance of football tipsters. International Journal of Forecasting 16, 317–331.

Forrest, D., Simmons, R., 2008. Sentiment in the betting market on Spanish football. Applied Economics 40, 119–126.

Franck, E., Verbeek, E., Nüesch, S., 2010. Prediction accuracy of different market structures—bookmakers versus a betting exchange. International Journal of Forecasting 26, 448–459.

Franck, E., Verbeek, E., Nüesch, S., 2013. Inter-market arbitrage in betting. Economica 80, 300–325.

Friedman, M., Savage, L.J., 1948. The utility analysis of choices involving risk. Journal of Political Economy 56, 279–304.

Gabriel, P.E., Marsden, J.R., 1990. An examination of market efficiency in British racetrack betting. Journal of Political Economy 98, 874–885.

Gainsbury, S., 2012. Internet Gambling: Current Research Findings and Implications. Springer, New York.

Gandar, J.M., Zuber, R.A., Johnson, R.S., 2004. A reexamination of the efficiency of the betting market on National Hockey League games. Journal of Sports Economics 5, 152–168.

Gandar, J.M., Zuber, R.A., Johnson, R.S., Dare, W., 2002. Re-examining the betting market on Major League Baseball games: is there a reverse favourite-longshot bias? Applied Economics 34, 1309–1317.

Gandar, J.M., Zuber, R.A., O'Brien, T., Russo, B., 1988. Testing rationality in the point spread betting market. Journal of Finance 43, 995–1008.

Gil, R., Levitt, S.D., 2007. Testing the efficiency of markets in the 2002 World Cup. Journal of Prediction Markets 1, 255–270.

Goddard, J., Asimakopoulos, I., 2004. Forecasting football results and the efficiency of fixed-odds betting. Journal of Forecasting 23, 51–66.

Golec, J., Tamarkin, M., 1991. The degree of inefficiency in the football betting market. Journal of Financial Economics 30, 311–323.

Golec, J., Tamarkin, M., 1995. Do bettors prefer long shots because they are risk-lovers, or are they just overconfident? Journal of Risk & Uncertainty 11, 51–64.

Golec, J., Tamarkin, M., 1998. Bettors love skewness, not risk, at the horse track. Journal of Political Economy 106, 205–225.

Graham, I., Stott, H., 2008. Predicting bookmaker odds and efficiency for UK football. Applied Economics 40, 99–109.

Gramm, M., Ziemba, W.T., 2008. The dosage breeding theory for horse racing predictions. In: Hausch, D.B., Ziemba, W.T. (Ed.), Handbook of Sports and Lottery Markets. North-Holland, Amsterdam, pp. 307–340.

Gray, P.K., Gray, S.F., 1997. Testing market efficiency: evidence from the NFL sports betting market. Journal of Finance 52, 1725–1737.

Gray, P.K., Gray, S.F., Roche, T., 2005. Efficiency of football betting markets: the economic significance of trading strategies. Accounting & Finance 45, 269–281.

Greene, W., 2012. Econometric Analysis, 7th international edition. Pearson Education, Boston.

Griffith, R.M., 1949. Odds adjustments by American horse-race bettors. American Journal of Psychology 62, 290–294.

Grossman, S.J., Stiglitz, J.E., 1980. On the impossibility of informationally efficient markets. American Economic Review 70, 393–408.

Gruen, A., 1976. An inquiry into the economics of race-track gambling. Journal of Political Economy 84, 169–177.

Guttman, A., 1993. The diffusion of sports and the problem of cultural imperialism. In: Dunning, E., Maguire, J.A., Pearton, R. (Ed.), The Sports Process: a Comparative and Developmental Approach. Human Kinetics, Champaign, pp. 125–137.

Harville, D.A., 1973. Assigning probabilities to the outcomes of multi-entry competitions. Journal of the American Statistical Association 68, 312–316.

Hausch, D.B., Ziemba, W.T., 1990. Arbitrage strategies for cross-track betting on major horse races. Journal of Business 63, 61–78.

Hausch, D.B., Ziemba, W.T., Rubinstein, M., 1981. Efficiency of the market for racetrack betting. Management Science 27, 1435–1452.

Henery, R.J., 1985. On the average probability of losing bets on horses with given starting price odds. Journal of the Royal Statistical Society: Series A (General) 148, 342–349.

Hodges, S.D., Tompkins, R., Ziemba, W.T., 2003. The favorite/long-shot bias in S&P 500 and FTSE 100 index futures options: the return to bets and the cost of insurance. EFA 2003 Annual Conference Paper No. 135; Sauder School of Business, Vancouver.

Hvattum, L.M., Arntzen, H., 2010. Using ELO ratings for match result prediction in association football. International Journal of Forecasting 26, 460–470.

Jaffe, J.F., 1974. Special information and insider trading. Journal of Business 47, 410–428.

Jensen, M.C., 1978. Some anomalous evidence regarding market efficiency. Journal of Financial Economics 6, 95–101.

Jullien, B., Salanié, B., 2000. Estimating preferences under risk: the case of racetrack bettors. Journal of Political Economy 108, 503–530.

Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. Econometrica 47, 263–291.

Kelly, J.L., 1956. A new interpretation of information rate. Bell System Technical Journal 35, 917–926.

Koning, R., 2012. Regression tests and the efficiency of fixed odds betting markets. International Journal of Sport Finance 7, 262–274.

Kotlyar, V., Smyrnova, O., 2012. Betting market: analysis of arbitrage situations. Cybernetics and Systems Analysis 48, 912–920.

Krieger, K., Fodor, A., Stevenson, G., 2013. The sensitivity of findings of expected bookmaker profitability. Journal of Sports Economics 14, 186–202.

Kuypers, T., 2000. Information and efficiency: an empirical study of a fixed odds betting market. Applied Economics 32, 1353–1363.

Laffey, D., 2005. Entrepreneurship and innovation in the UK betting industry: the rise of person-to-person betting. European Management Journal 23, 351–359.

Levitt, S.D., 2004. Why are gambling markets organised so differently from financial markets? Economic Journal 114, 223–246.

Lim, K.P., Brooks, R., 2011. The evolution of stock market efficiency over time: a survey of the empirical literature. Journal of Economic Surveys 25, 69–108.

Lo, A.W., 2004. The adaptive markets hypothesis. Journal of Portfolio Management 30, 15–29.

Lo, A.W., MacKinlay, A.C., 1999. A Non-Random Walk Down the Wall Street. Princeton University Press, Princeton.

Lofthouse, S., 2001. Investment Management, 2nd edition. John Wiley and Sons, New York.

Losey, R.L., Talbott, J.C.J., 1980. Back on the track with the efficient markets hypothesis. Journal of Finance 35, 1039–1043.

MacLean, L.C., Thorp, E.O., Ziemba, W.T., 2010. Long-term capital growth: the good and bad properties of the Kelly and fractional Kelly capital growth criteria. Quantitative Finance 10, 681–687.

MacLean, L.C., Ziemba, W.T., 2006. Capital growth: theory and practice. In: Zenios, S.A., Ziemba, W.T. (Ed.), Handbook of Asset and Liability Management, Vol. 1: Theory and Methodology. North-Holland, Amsterdam, pp. 429–473.

MacLean, L.C., Ziemba, W.T., Blazenko, G., 1992. Growth versus security in dynamic investment analysis. Management Science 38, 1562–1585.

Malarić, R., Katić, T., Sabolić, D., 2008. The market efficiency of the soccer fixed odds internet betting market. Applied Economics Letters 15, 171–174.

Malkiel, B.G., 1992. Efficient market hypothesis. In: Newman, P., Milgate, M., Eatwell, J. (Ed.), New Palgrave Dictionary of Money and Finance. Macmillan, London.

Malkiel, B.G., 2003. The efficient market hypothesis and its critics. Journal of Economic Perspectives 17, 59–82.

Markowitz, H., 1952. The utility of wealth. Journal of Political Economy 60, 151–158.

McGlothlin, W.H., 1956. Stability of choices among uncertain alternatives. American Journal of Psychology 69, 604–615.

Milliner, I., White, P., Webber, D., 2009. A statistical development of fixed odds betting rules in soccer. Journal of Gambling Business and Economics 3, 89–99.

Millman, C., 2001. The odds: One Season, Three Gamblers, and the Death of their Las Vegas. Public Affairs, New York.

Neumann, J.V., Morgenstern, O., 1944. Theory of Games and Economic Behavior. Princeton University Press, Princeton.

Nyberg, H., 2014. A multinomial logit-based statistical test of association football betting market efficiency. HECER Discussion Paper No. 380. Helsinki Center of Economic Research, Helsinki.

Oikonomidis, A., Johnson, J., 2011. Who can beat the odds? The case of football betting reviewed. In: Vaughan Williams, L. (Ed.), Prediction Markets: Theory and Applications. Routledge, Abingdon, pp. 204–220.

Osborne, E., 2001. Efficient markets? Don’t bet on it. Journal of Sports Economics 2, 50–61.

Pankoff, L.D., 1968. Market efficiency and football betting. Journal of Business 41, 203–214.

Paton, D., Vaughan Williams, L., 1998. Do betting costs explain betting biases? Applied Economics Letters 5, 333–335.

Paton, D., Vaughan Williams, L., 2005. Forecasting outcomes in spread betting markets: can bettors use 'quarbs' to beat the book? Journal of Forecasting 24, 139–154.

Paul, R.J., Weinbach, A.P., 2005a. Market efficiency and NCAA college basketball gambling. Journal of Economics and Finance 29, 403–408.

Paul, R.J., Weinbach, A.P., 2005b. Bettor preferences and market efficiency in football totals markets. Journal of Economics and Finance 29, 409–415.

Paul, R.J., Weinbach, A.P., 2012. Sportsbook pricing and the behavioral biases of bettors in the NHL. Journal of Economics and Finance 36, 123–135.

Paul, R.J., Weinbach, A.P., Wilson, M., 2004. Efficient markets, fair bets, and profitability in NBA totals 1995–96 to 2001–02. Quarterly Review of Economics and Finance 44, 624–632.

Pope, P.F., Peel, D.A., 1989. Information, prices and efficiency in a fixed-odds betting market. Economica 56, 323–341.

Preston, I., Szymanski, S., 2003. Cheating in contests. Oxford Review of Economic Policy 19, 612–624.

Quandt, R.E., 1986. Betting and equilibrium. Quarterly Journal of Economics 101, 201–207.

Rosett, R.N., 1965. Gambling and rationality. Journal of Political Economy 73, 595–607.

Rozeff, M.S., Zaman, M.A., 1988. Market efficiency and insider trading: new evidence. Journal of Business 61, 25–44.

Rubinstein, M., 1985. Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active CBOE option classes from August 23, 1976 through August 31, 1978. Journal of Finance 40, 455–480.

Rue, H., Salvesen, O., 2000. Prediction and retrospective analysis of soccer matches in a league. Journal of the Royal Statistical Society: Series D (The Statistician) 49, 399–418.

Ruhm, D.L., 2003. Distribution-based pricing formulas are not arbitrage-free. Casualty Actuarial Society Proceedings 90, 97–129.

Samuelson, P., 1965. Proof that properly anticipated prices fluctuate randomly. Industrial Management Review 6, 41–49.

Sauer, R.D., 1998. The economics of wagering markets. Journal of Economic Literature 36, 2021–2064.

Sauer, R.D., Brajer, V., Ferris, S.P., Marr, M.W., 1988. Hold your bets: another look at the efficiency of the gambling market for National Football League games. Journal of Political Economy 96, 206–213.

Schnytzer, A., Weinberg, G., 2004. Is the NBA betting market efficient? In: Papanikos, G.T. (Ed.), The Economics and Management of Mega Athletic Events: Olympic Games, Professional Sports, and Other Essays. Athens Institute for Education and Research, Athens, pp. 71–85.

Schnytzer, A., Weinberg, G., 2008. Testing for home team and favorite biases in the Australian rules football fixed-odds and point spread betting markets. Journal of Sports Economics 9, 173–190.

Sessford, P., White, P., 2010. A tobit regression model of football performance variables for a profitable fixed odds betting strategy. Journal of Gambling Business and Economics 4, 1–12.

Sewell, M., 2011. History of the efficient market hypothesis. Research note RN/11/04. University College London, London.

Sewell, M., 2012. The efficient market hypothesis: empirical evidence. International Journal of Statistics and Probability 1, 164–178.

Seyhun, H.N., 1986. Insiders' profits, costs of trading, and market efficiency. Journal of Financial Economics 16, 189–212.

Shin, H.S., 1991. Optimal betting odds against insider traders. Economic Journal 101, 1179–1185.

Shleifer, A., Vishny, R.W., 1997. The limits of arbitrage. Journal of Finance 52, 35–55.

Sinkey, M., Logan, T., 2014. Does the hot hand drive the market? Evidence from college football betting markets. Eastern Economic Journal 40, 583–603.

Smith, M.A., Paton, D., Vaughan Williams, L., 2006. Market efficiency in person-to-person betting. Economica 73, 673–689.

Smith, M.A., Paton, D., Vaughan Williams, L., 2009. Do bookmakers possess superior skills to bettors in predicting outcomes? Journal of Economic Behavior and Organization 71, 539–549.

Snowberg, E., Wolfers, J., 2010. Explaining the favorite-long shot bias: is it risk-love or misperceptions? Journal of Political Economy 118, 723–746.

Snyder, W.W., 1978. Horse racing: testing the efficient markets model. Journal of Finance 33, 1109–1118.

Song, C., Boulier, B.L., Stekler, H.O., 2007. The comparative accuracy of judgmental and model forecasts of American football games. International Journal of Forecasting 23, 405–413.

Stekler, H.O., Sendor, D., Verlander, R., 2010. Issues in sports forecasting. International Journal of Forecasting 26, 606–621.

Štrumbelj, E., Šikonja, M.R., 2010. Online bookmakers' odds as forecasts: the case of European soccer leagues. International Journal of Forecasting 26, 482–488.

Surowiecki, J., 2004. The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies, and Nations. Doubleday, New York.

Swidler, S., Shaw, R., 1995. Racetrack wagering and the `uninformed' bettor: a study of market efficiency. Quarterly Review of Economics and Finance 35, 305–314.

Terrell, D., Farmer, A., 1996. Optimal betting and efficiency in parimutuel betting markets with information costs. Economic Journal 106, 846–868.

Thaler, R.H., Ziemba, W.T., 1988. Parimutuel betting markets: racetracks and lotteries. Journal of Economic Perspectives 2, 161–174.

Thorp, E.O., 2006. The Kelly criterion in blackjack, sports betting and the stock market. In: Zenios, S.A., Ziemba, W.T. (Ed.), Handbook of Asset and Liability Management, Vol. 1: Theory and Methodology. North-Holland, Amsterdam, pp. 385–428.

Treynor, J.L., 1987. Market efficiency and the bean jar experiment. Financial Analysts Journal 43, 50–52.

Tryfos, P., Casey, S., Cook, S., Leger, G., Pylypiak, B., 1984. The profitability of wagering on NFL games. Management Science 30, 123–132.

Vaughan Williams, L., 2001. Can bettors win? A perspective on the economics of betting. World Economics Journal 2, 31–48.

Vaughan Williams, L., 2005. Information efficiency in financial markets. In: Vaughan Williams, L. (Ed.), Information Efficiency in Financial and Betting Markets. Cambridge University Press, New York, pp. 5–83.

Vaughan Williams, L., Paton, D., 1997. Why is there a favourite-longshot bias in British racetrack betting markets? Economic Journal 107, 150–158.

Vaughan Williams, L., Paton, D., 1998. Why are some favourite-longshot biases positive and others negative? Applied Economics 30, 1505–1510.

Vecer, J., Ichiba, T., Laudanovic, M., 2006. Parallels between betting contracts and credit derivatives. Unpublished working paper. Department of Statistics, Columbia University, New York.

Vergin, R.C., Scriabin, M., 1978. Winning strategies for wagering on National Football League games. Management Science 24, 809–818.

Vlastakis, N., Dotsis, G., Markellos, R.N., 2009. How efficient is the European football betting market? Evidence from arbitrage and trading strategies. Journal of Forecasting 28, 426–444.

Weitzman, M., 1965. Utility analysis and group behavior – an empirical study. Journal of Political Economy 73, 18–26.

Wilcoxon, F., 1945. Individual comparisons by ranking methods. Biometrics Bulletin 1, 80–83.

Williams, R.J., Wood, R.T., Parke, J., 2012. History, current worldwide situation, and concerns with internet gambling. In: Williams, R.J., Wood, R.T., Parke, J. (Ed.), Routledge International Handbook of Internet Gambling. Routledge, New York, pp. 3–26.

Woodland, L.M., Woodland, B.M., 1994. Market efficiency and the favorite-longshot bias: the baseball betting market. Journal of Finance 49, 269–379.

Woodland, L.M., Woodland, B.M., 2001. Market efficiency and profitable wagering in the National Hockey League: can bettors score on longshots? Southern Economic Journal 67, 983–995.

Woodland, L.M., Woodland, B.M., 2003. The reverse favourite-longshot bias and market efficiency in Major League Baseball: an update. Bulletin of Economic Research 55, 113–123.

Zuber, R.A., Gandar, J.M., Bowers, B.D., 1985. Beating the spread: testing the efficiency of the gambling market for National Football League games. Journal of Political Economy 93, 800–806.