University of Zurich UZH

Center for Research in Sports Administration (CRSA)

Working Paper Series

Prediction accuracy of different market structures – bookmakers versus a betting exchange

Egon Franck[a], Erwin Verbeek[a,*], Stephan Nüesch[a]

[a] Institute of Strategy and Business Economics,
University of Zurich, Plattenstrasse 14, 8032 Zürich, Switzerland

Abstract

There is a well-established literature on separately testing the prediction power of different betting market settings. This paper provides an inter-market comparison of the forecasting accuracy between bookmakers and a major betting exchange. Employing a dataset covering all football matches played in the major leagues of the "Big Five" (England, France, Germany, Italy, Spain) during three seasons (5478 games in total), we find evidence that the betting exchange provides more accurate predictions of the same underlying event than bookmakers. A simple betting strategy of selecting bets for which bookmakers offer lower probabilities (higher odds) than the bet exchange generates above average and, in some cases, even positive returns.

---

* Corresponding author: University of Zurich, Institute for Strategy and Business Economics, Plattenstrasse 14, 8032 Zurich, email: erwin.verbeek@isu.uzh.ch, phone: +41 44 634 29 61

1. Introduction

Similar to financial securities, betting markets trade contracts on future events. The price of a contract reflects the owner's claim, which is tied to the event's outcome. Therefore, the market price can be interpreted as a prediction of the future event. According to Vaughan Williams (1999), betting markets are particularly well suited to investigate forecasting accuracy because – in contrast to most financial markets – the contracts have a definite value that becomes observable after a clear termination point.

The traditional form of gambling on sports events is bookmaker betting. In this market setting, the bookmaker acts as a dealer announcing the odds against which the bettor can place his bets. In recent years, a different market structure has evolved: exchange betting. Whereas the bookmaker defines the odds ex ante, the prices in the bet exchange are determined by a multitude of individuals trading the bets among themselves. This form of person-to-person betting has lately experienced rapid growth.

Empirical research on the prediction accuracy of bookmaker odds is well established in the literature. While some papers document a high forecasting performance of bookmaker odds (e.g., Boulier & Stekler, 2003; Forrest, Goddard, & Simmons, 2005), other research provides evidence of biases in bookmaker predictions. These biases, however, turn out to be rather small, as they hardly provide opportunities to systematically beat the odds (e.g., Cain, Law, & Peel, 2000; Dixon & Pope, 2004; Goddard & Asimakopoulos, 2004).

Furthermore, there is a growing body of literature concerned with the prediction power of bet exchange markets. It is found that these markets exhibit high prediction accuracy as they regularly outperform non-market forecasting methods (e.g., Forsythe, Nelson, Neumann, & Wright, 1992; Wolfers & Leigh, 2002; Spann & Skiera, 2003; Berg, Nelson, & Rietz, 2008; and Snowberg, Wolfers, & Zitzewitz, 2008, for an excellent review).

The coexistence of different betting markets offering quotes on the very same event enables us to compare their prediction power. Surprisingly, examples of this kind of re-

search are rare.[2] To the best of our knowledge, this paper is the first to contrast the fore-casting accuracy of the bookmaker market with a major betting exchange. Using a dataset covering all football matches played in the major leagues of the "Big Five" (England, France, Germany, Italy, Spain) during three seasons (5478 matches in total), we compare the prediction accuracy of eight different bookmaker odds with the forecasting power of the corresponding odds traded at *Betfair*, a common bet exchange platform. Our results indicate that the prices of the bet exchange market exhibit higher prediction power than the bookmaker odds. Furthermore, we develop a simple betting strategy in order to test the economic relevance of our findings. We show that a strategy of selecting bets for which the bookmaker announced lower probabilities (and thus, offered higher odds) than the person-to-person market is capable of yielding above average and, in some cases, even positive returns. This betting strategy is not restrictive in terms of betting opportunities.

Our findings contribute to the ongoing discussion about the predictive properties of different market structures by providing empirical evidence of the superiority of exchange betting in delivering more accurate forecasts of the outcome of sporting events.

2. Different betting market structures

In this section, we present some preliminary background information on how to interpret betting odds as outcome probabilities. We then outline the structures of the bookmaker market and the bet exchange and present literature on their forecasting effectiveness.

2.1. Betting odds and outcome probabilities

In football matches, there are three different outcomes $e \in \{h, d, a\}$ – a home win, a draw and an away win – on which a bet can be placed. The market prices of these out-

---

[2] Comparing bookmaker odds and bet exchange odds in UK horse racing, Smith, Paton, and Vaughan Williams (2006) discover that person-to-person betting is more efficient, as it lowers transaction costs for consumers. Spann and Skiera (2008) compare the predictions of a bookmaker (*Oddset*) with the prices of a virtual football stock exchange market (www.*bundesligaboerse.de*) and find that they perform equally well.

comes are typically presented as 'decimal odds' o_e, which stand for the payout ratio of a winning bet. The inverse of the decimal odds $1/o_e$ can be interpreted as the probability of the underlying event occurring, which is offered to the betting audience. These market probabilities on all possible outcomes of an event usually sum up to greater than one because of the transaction costs, the so-called 'overround'. Thus, $\sum_e 1/o_e \geq 1$ holds. In order to get the market's prediction of the outcome, we assume that the overround is equally distributed over the outcome probabilities.[3] Therefore, we obtain the market's 'implicit probabilities' by a linear transformation,

$$Prob_e = \frac{1}{o_e} \frac{1}{\sum_e \frac{1}{o_e}}$$

In what follows, we refer to this expression as the market's prediction of a future event.

2.2. The bookmaker market

Bookmaker betting is among the most popular forms of sports gambling. In this setting, the bookmaker acts as a market maker. He determines the odds on a given event and takes the opposite side of every transaction.[4] The bettor is left with a take-it-or-leave-it decision: he can hit the market quotes or refrain from participating. This is why bookmaker markets are sometimes denoted as quote-driven markets, by analogy to the same setting in financial markets.

The prediction accuracy is determined by the price-setting behavior of the bookmaking firm. If the odds are fixed such that they reflect true outcome probabilities, the bookmaker will, on average, earn a profit margin equal to the commission charged to the bettors. Alternatively, the bookmaker can 'balance his book' by setting the odds to attract equal relative betting volumes on each side. In this case, he is able to pay out the winners with the stakes of the losers and earn the overround independently from the outcome of the event. Levitt (2004) has pointed out that the bookmaker might use a combination of

---

[3] This assumption is in line with the literature. See, for example, Forrest, Goddard, and Simmons (2005).

[4] The bookmakers have the right to change the odds after the market has opened, but they rarely make adjustments (Forrest, Goddard, & Simmons, 2005). The bettor's claim is tied to the initially taken odd and does not depend on subsequent price changes. We therefore speak of 'fixed-odds betting'.

the two extreme cases in order to increase his profits. Forrest and Simmons (2008) and Franck, Verbeek and Nüesch (2008) demonstrate, for example, that bookmakers actively shade prices to attract betting volume evoked by sentiment in Spanish and English football, respectively. Here, bookmakers offer more (less) favorable terms for bets on teams with a comparably large (small) fan base in order to attract disproportionately more betting volume. Thus, the bookmaker odds are likely to be influenced by the true outcome probabilities as well as the bettors' demand. The latter can lead to deviations from the true outcome probability. Nevertheless, such deviations are clearly limited. In practice, it might be a difficult task to balance the book by shading the odds, and therefore, the bookmaker is exposed to substantial risk if his prices deviate from true outcome probabilities.

The empirical literature on the prediction accuracy of bookmaker odds is mixed. Forrest, Goddard, and Simmons (2005) compare the prediction accuracy of published bookmaker odds for English football games with the forecasts of a benchmark statistical model that incorporates a large number of quantifiable variables relevant to match outcomes. They find evidence that bookmaker odds are more effective in predicting game outcomes than the statistical model. A longstanding empirical regularity that challenges the suitability of bookmaker odds as predictors is the 'longshot bias'. It refers to the observation that the odds often underestimate high-probability outcomes (favorites to win the game) and overestimate low-probability outcomes (underdogs to win the game). As a result, it is frequently found that bets on low-probability outcomes ('longshots') yield lower average returns than bets on high-probability outcomes (e.g., Cain, Law, & Peel, 2000). Dixon and Pope (2004) as well as Goddard and Asimakopoulos (2004) find that the bookmaker's odds are weak-form inefficient as they do not impound all information that has proven to be significantly related to the game's outcome according to a statistical forecasting model.

2.3. The bet exchange market

In recent years, person-to-person exchange betting has evolved as a different betting market structure. Here, individuals contract their opposing opinions with each other. On an online platform, they can post the prices under which they are willing to place a bet – on or against – a given event. The latent demand for wagers is collected and presented in

the order book. The order book displays the most attractive odds with the corresponding available volume in a canonical manner. This is why such a market design is often referred to as an order-driven market. The bettor has the choice to either submit a limit order and wait for another participant to match his bet or to submit a market order and directly match an already offered bet. As a result, there is a continuous double auction process taking place at the online platform. If two bettors with opposing opinions agree on a price, their demands are automatically translated into a transaction. After the bets have been matched, both individuals hold a contract on a future cash flow. The size of the cash flow is determined by the price of the contract, and the direction of the cash flow is tied to the outcome of the underlying event. The provider of the platform charges a commission fee, typically lower than the bookmaker's overround, on the bettors' net profits.

Online betting exchanges have experienced a fast boom. The odds analyzed in this paper are from *Betfair,* which is one of the most prominent bet exchange platforms. With a weekly turnover of more than $50m and over two million registered users, *Betfair* accounts for 90% of all exchange-based betting activity worldwide (Croxson & Reade, 2008; *www.betfaircorporate.co.uk*). It has been online since 2000 and claims to process five million trades a day.

From a theoretical perspective, bet exchanges should yield accurate forecasts. First, the betting exchange provides incentives to gather and process information. Traders who have superior knowledge are able to generate higher average returns than naïve bettors. Second, the betting exchange provides incentives for the truthful revelation of information. Based on their knowledge, traders put money at stake and, in doing so, they reveal their expectations of the outcome's probability. Third, through the price mechanism, the betting exchange provides an efficient algorithm to collect and aggregate diverse information in a dynamic way (Wolfers & Zitzewitz, 2004; Berg, Nelson, & Rietz, 2008; Snowberg, Wolfers, & Zitzewitz, 2008). As a matter of fact, empirical studies have shown that bet exchanges provide highly accurate predictions. They routinely render better predictions of political elections than opinion polls (Forsythe, Nelson, Neumann, & Wright, 1992; Wolfers & Leigh, 2002; Berg, Nelson, & Rietz, 2008) and outperform expert opinions in forecasting future business outcomes (Spann & Skiera, 2003; Pennock, Lawrence, Nielsen, &

Giles, 2001). In addition to the specific prediction literature, there are a few papers examining the efficiency of *Betfair* prices in particular. Smith, Paton, and Vaughan Williams (2006) used matched data on UK horse racing from *Betfair* and from traditional bookmakers to test the well-documented longshot bias. They find that the bet exchange is significantly more efficient than the bookmaker market, as the tendency to overvalue underdogs is less pronounced in person-to-person betting. Croxon and Reade (2008) employ high-frequency *Betfair* data to test efficiency in relation to the arrival of goals. They conclude that prices swiftly and fully impound the relevant news, indicating the high efficiency of *Betfair* odds.

3. Prediction accuracy of the different markets

3.1. The data

Our data cover all football games of the English *Premier League*, the Spanish *Primera Division*, the Italian *Serie A*, the German *Bundesliga* and the French *Ligue 1* during three seasons (2004/05 to 2006/07), summing up to 5478 games in total. We analyze the odds of eight different bookmakers[5] from *www.football-data.co.uk*, on which they are recorded on Friday afternoons for weekend games and on Tuesday afternoons for midweek games. In addition, we matched the bookmaker data with corresponding betting exchange prices from *www.betfair.com*, collected at the same time.[6] The decimal odds from *Betfair* and the bookmakers are converted into implicit probabilities according the procedure described in chapter 2.1.

Correlations between the implicit probabilities of *Betfair* and a random bookmaker are 0.917 for draw bets, 0.978 for away win bets and 0.981 for home win bets. Thus, the odds traded on the betting exchange are very similar to the bookmaker odds. Figure 1 graphically relates *Betfair* probabilities to the (random) bookmaker probabilities for all

---

[5]   The bookmakers are *B365*, *Bet&Win*, *Gamebookers*, *Interwetten*, *Ladbrokes*, *William Hill*, *Stan James* and *VC Bet*.

[6]   We used the *Betfair* odds on which bets were actually matched.

three possible match outcomes separately; the black line indicates the cases for which the probabilities of the two markets are equal.

< Figure 1 >

It can be seen that the probabilities of the two markets are closely aligned. At first glance, the differences are somewhat unsystematically distributed. Upon a closer look, the bookmaker probabilities appear to be higher (lower) than *Betfair* probabilities in the area of low (high) probability outcomes for home and away win bets.

Table 1 presents some summary statistics in order to provide a first impression of the two markets' effectiveness in forecasting the outcomes. The first column outlines the observed overall proportions of the three possible outcomes of a game. The second and the third columns contain the predicted probabilities implied by the odds of the betting exchange and the bookmaker market[7], respectively.

< Table 1 >

Table 1 suggests that the average probabilities of *Betfair* are closer to the overall proportions of home wins, draws and away wins compared to the bookmaker. Both of them, however, underestimate the occurrence of home wins against away wins. These numbers provide only a rough picture of the markets' prediction accuracy. In the following, we will test how well the markets' implicit probabilities correspond to the actual outcome of every single game.

3.2. Goodness-of-fit of discrete response models

We estimate the following model to explain the actual outcome (win or loss) of a certain bet $Y_{ei} \in \{0,1\}$ for a given match i by the implicit probabilities of the different markets $Prob_{eij}$

$$Y_{ei} = G\big(\alpha_{ej} + \beta_{ej}\, Prob_{eij} + \varepsilon_{eij}\big).$$

---

[7] For the sake of clarity, we solely report the probabilities of one bookmaker, who is randomly picked for each match from our set of eight bookmakers.

For each event e (home wins, draws and away wins) and every market j (eight bookmakers and *Betfair*), the coefficients $\hat{\beta}_{ej}$ are estimated by a probit model. The probit model relates the probability of occurrence of discrete events to some set of explanatory variables, where $G(\cdot) = \phi(\cdot)$ is the standard normal cumulative distribution.[8]

The prediction accuracy is examined by various goodness-of-fit measures. Whereas the first three are common goodness-of-fit measures proposed for discrete choice models, the fourth indicator, the Brier Score, is a descriptive measure often used in the literature on prediction accuracy (e.g., Boulier & Stekler, 2003; Forrest, Goddard, & Simmons, 2005).

In a linear model, the percentage of variance of the dependent variable explained by the model, R^2, would be the obvious measure. In non-linear discrete response models, however, the R^2 measure is not directly applicable, as proper variance decomposition is not possible. A number of so-called *pseudo R^2* measures have been suggested for discrete response models.[9] The most common was proposed by McFadden (1974) and is defined in the following way:

$$R^2_{McFadden,ej} = 1 - \frac{\log L_{ur,ej}}{\log L_{0,ej}},$$

where $\log L_{ur,ej}$ is the value of the (maximized) log-likelihood function for the estimated model for a given event and market, and $\log L_{0,ej}$ is the value of the (maximized) log-likelihood function in the model with only an intercept for a given event and market. As the value of the log-likelihood function is always negative, $\log L_{ur}/\log L_0$ = $|\log L_{ur}|/|\log L_0|$ holds. Further $|\log L_{ur}| \leq |\log L_0|$, which implies that the *pseudo R^2* is always between 0 and 1. The *pseudo R^2* of McFadden is 1 if the model is a perfect predictor and zero if the model has no explanatory power.

---

[8]   An alternative to the probit model is the logit model or the linear probability model (*LPM*). The logit model assumes a logistic distribution and is therefore, like the probit, a non-linear model, whereas the *LPM* is based on ordinary least squares and assumes constant marginal effects. In order to test the robustness of our results, we also ran logit and *LPM* estimations. The results, however, are not sensitive to alternative estimation procedures.

[9]   Winkelmann and Boes (2005) and Wooldridge (2002) provide excellent reviews of the goodness-of-fit measures in non-linear discrete response models.

Another *pseudo R²* measure was proposed by McKelvey and Zavoina (1975). Unlike McFadden's R^2, it is based on a linear model $Y_{ei} = \beta_{ej}Prob_{eij} + \varepsilon_{eij}$. So, the goodness-of-fit is defined as:

$$R^2_{McKelvy\&Zavoina,ej} = \frac{\sum_{i=1}^{I}(\hat{y}_{eij} - \bar{Y}_{ej})^2}{I + \sum_{i=1}^{I}(\hat{y}_{eij} - \bar{Y}_{ej})^2},$$

where $\sum_{i=1}^{I}(\hat{y}_{eij} - \bar{Y}_{ej})^2$ in the numerator denotes the explained sum of squares and $I + \sum_{i=1}^{I}(\hat{y}_{eij} - \bar{Y}_{ej})^2$ in the denominator is the model's total sum of squares. *I* denotes the number of observations that are used for estimating the model, which corresponds to the number of games in our context.

An often-used alternative measure of prediction accuracy is the *percentage of correct predictions* (e.g. Spann & Skiera, 2008; Berg, Nelson, & Rietz, 2008). So if $\hat{\beta}_{ej}Prob_{eij} > t$, $\hat{y}_{ei}$ is predicted to be unity, and if $\hat{\beta}_{ej}Prob_{eij} \leq t$, $\hat{y}_{ei}$ is predicted to be zero. Usually, the cut-off value *t* equals 0.5. If the distribution of the dependent variable is skewed, however, the percentage correctly predicted can be misleading as a measure of prediction accuracy (Wooldridge, 2002). In such cases other cut-off values have to be chosen in order to minimize the forecasting errors. All combinations of a given sample and possible cut-off values can be summarized in the so-called *Receiver Operating Characteristic (ROC)* curve. The area under the *ROC* curve indicates the goodness-of-fit of a certain discrete response model. The *ROC* area varies between 0.5, indicating no prediction power at all and 1, which means perfect prediction.

As a fourth measure of prediction accuracy, we use the *Brier Score* (Brier, 1950), which is defined as the mean squared difference between the actual outcome and the predicted outcome:

$$Brier\ Score_{ej} = \frac{\sum_{i=1}^{I}(Y_{ei} - Prob_{eij})^2}{I}.$$

Unlike the other goodness-of-fit measures, a small *Brier Score* indicates not low but high forecasting accuracy. If the predictions are perfectly accurate, the *Brier Score* is 0 and vice versa for a *Brier Score* of 1.

In the following, we illustrate the prediction accuracies of *Betfair* and various book-makers for home win bets, draw bets and away win bets separately.

< Table 2 >

Table 2 suggests that the implied probabilities of *Betfair* explain the actual outcomes better than the bookmakers' probabilities do. The model with the implicit probabilities of *Betfair* (first column) as the explanatory variable renders better goodness-of-fit scores than the regressions using the average of bookmaker probabilities (second column) and the probability of every single bookmaker (third to tenth columns). With the exceptions of *VC Bet* for home and draw bets and *Stan James* for draw bets, the two R^2 measures and the *ROC area* are always higher and the *Brier Score* is always lower for the bet exchange probabilities.

There are some additional patterns that are worth mentioning. First, the prediction accuracy of draws is remarkably worse than for home and away wins. The goodness-of-fit measures in the middle row are considerably lower (and higher regarding the *Brier score*). This observation is in line with Dobson and Goddard's (2001) conclusion that, in football matches, draws appear to be almost random events. Second, the marginal effects of the implied probability are substantially above unity for all bookmakers and events. Thus the actual winning probability increases disproportionately with the implied book-maker probabilities. This indicates the presence of a longshot bias. Hence the odds underestimate high-probability outcomes (e.g., favorites to win the game) and overestimate low-probability outcomes (e.g., underdogs to win the game). This effect is strongest for draw bets and weakest for away win bets. Most important, the marginal effects of the *Bet-fair* probabilities are closer to unity than the probabilities of any other bookmaker. Therefore, at least part of the better prediction accuracy of the bet exchange is the consequence of the weaker longshot bias in person-to-person betting compared to the bookmaker market, which confirms the findings of Smith, Paton, and Vaughan Williams (2006).

3.3. Direct comparison of prediction accuracy

The differences of the goodness-of-fit measures between the two markets in Table 2 are rather small. In the following, we include the predictions of the two betting markets in

the same model. In doing so, we are able to test whether the probabilities of the bet exchange contribute additional explanatory power beyond the bookmaker's forecasts.

We rerun the regressions described in the previous section, but we include the ratio of bet exchange to bookmaker probability $R_{eij} = \frac{Prob_{ei,BF}}{Prob_{eij}}$ as a variable capturing the difference between the two markets' predictions.[10] Thus, we estimate the following probit model for every single bookmaker j and all three events e:

$$Y_{ei} = G\big(\alpha_{ej} + \beta_{1,ej}\, Prob_{eij} + \beta_{2,ej}\, R_{eij} + \varepsilon_{eij}\big)\,.$$

If $\hat{\beta}_{2,ej} \neq 0$, the prediction of the bet exchange provides some relevant information that is not fully captured by the odds of the bookmaker.

<center>< Table 3 ></center>

Table 3 shows that the coefficient $\hat{\beta}_{2,ej}$ is significantly positive in each case. Thus, the inclusion of the exchange market's predictions improves the forecasting accuracy of the bookmaker odds. This demonstrates that the odds offered by the bookmakers fail to impound some relevant information delivered by the odds traded at *Betfair*.[11]

4. A simple betting strategy

Our results suggest that the betting exchange market predicts future outcomes more accurately than the bookmakers do. Next, we set out to test the economic relevance of this observation. We look at the return of a bet as a combination of its price and its winning probability. If the exchange market provides better forecasts of this probability than the

---

[10] Another possibility would be to directly include both, the bookmaker's probability and the probability of the bet exchange, into the same model. A potential problem of this procedure is the high multicollinearity between the two variables. Nevertheless, this method renders results that could be interpreted in the same way as the results of the method reported in the paper: The coefficient of every single bookmaker probability loses its significance against the *Betfair* probability in the home win and away win regressions. In the regressions for draws, both coefficients lose their statistical significance.

[11] Furthermore, comparing Table 3 with the results of Table 2, one can see that the *pseudo R^2* values increase with the inclusion of R_{eij} in all cases (except for draw bets with *VC Bet* for which the *pseudo R^2* remains the same).

bookmaker does, a betting rule exploiting forecasting differences between the two markets should yield above average returns.

In a first step, we compare the mean return of a simple betting strategy with normal returns. The trading rule is to place a bet against a given bookmaker in all the cases in which the implicit probability of *Betfair* exceeds the average implicit probability of the bookmakers. Thus, we use the prices of the exchange market as a source of information in order to detect favorable bookmaker odds. We place a bet at the bookmaker market whenever the (average) odds offered by the bookmakers are higher than the odds traded at *Betfair*. Table 4 presents the mean returns when following this betting strategy; the number of available bets is given in parentheses.

<center>< Table 4 ></center>

The results in Table 4 are broken down by the events on which to place a bet (columns) and the bookmakers (rows). The first row presents the results for a randomly chosen bookmaker and the second row for the bookmaker offering the most favorable odds. It can be seen that the strategy enables above-average returns in all cases, as the mean returns following the trading rule (right hand side of each column) are less negative than the average return of all bets on a given event (left hand side of each column) and, in some cases, they are even positive. The markup is strongest for away win bets in which the trading rule is capable of generating positive returns, except for *William Hill* bets.

As a second step, we compute the observed average returns for different levels of disagreement between the two markets. In doing so, we get a closer picture of the findings documented in Table 4. If the implicit probabilities of *Betfair* are closer to the true outcome probabilities, the expected return of a bet against a given bookmaker increases with the difference of the two markets' probabilities. Thus, in line with our previous findings, we expect a positive relation between the observed returns and the ratio of bet exchange to (average) bookmaker probability. To investigate this relation by means of our data, we rank all bets according to their ratio of *Betfair* to average bookmaker probability defined as

<center>12</center>

$$R_{ei}^* = \frac{Prob_{ei,BF}}{\frac{1}{J_i} \sum_j Prob_{eij}},$$

where J_i is the number of participating bookmakers in match i.[12] We then plot the observed mean returns of the bets against different categories of R_{ei}^*. The categories are specified by a bandwidth of 0.05 of R_{ei}^* and at least 50 observations per group are required. Furthermore, we run a locally weighted polynomial regression (Fan, 1992; Fan & Gijbels, 1996). In doing so, we have to make no assumption about the functional form of the relation between the returns and R_{ei}^*.[13] Figure 2 graphs the results of this procedure for bets against a random bookmaker on all events; Figure 3 presents the results for home win, draw and away win bets separately.

< Figure 2 >

< Figure 3 >

It can be seen that for $R_{ei}^* = 1$ the observed mean returns are roughly at the level of normal returns (the dashed horizontal line) and, more importantly, that they increase with R_{ei}^*. The mean returns for the different categories of R_{ei}^* (the dots) as well as the local polynomial smoother (the solid line) increase with R_{ei}^* in the case of home and away win bets (Figure 3) and in the case of all events taken together (Figure 2). The relation is steeper for away win bets than for home win bets. Moreover, the figures demonstrate that the betting strategy enables positive returns for some levels of R_{ei}^*. For example, betting against the random bookmaker on all events in the top 5%-quantile regarding R_{ei}^* (821 bets in total) yields an average return of +10%, betting on all home wins of the top 10%-quantile (547 bets in total) yields +3% and the away win bets of the top 10%-quantile (547 bets in total) yield an average return of +7%.

---

[12] In contrast to R_{eij} in the previous section, R_{ei}^* only varies across matches and events but is constant across bookmakers.

[13] Local polynomial regression involves fitting the response (the observed returns) to a polynomial form of the regressor (R_{ei}^*) via locally weighted least squares. We estimate a local cubic polynomial weighted by the Epanechnikov kernel function. The amount of smoothing is controlled by a bandwidth chosen at 0.2.

This demonstrates that the odds traded at *Betfair* offer information on the outcome probability of the matches that is useful to select under-priced bets in the bookmaker market. Taken together, these findings lend further confirmation for the superiority of the betting exchange in terms of prediction accuracy.

5. Summary and conclusions

Much research has been conducted on separately testing the prediction accuracy of different betting market settings. This paper exploits the coexistence of different market structures offering odds on the same event in order to provide an inter-market comparison of the predictive power of bookmakers and a major betting exchange. We analyze a dataset covering 5478 matches of the major European football leagues containing the odds of eight bookmakers and the corresponding prices of the leading person-to-person betting platform *Betfair*. Our results reveal a clear superiority of the betting exchange over the bookmaker market. First, we estimate a univariate probit regression to explain the actual outcome of a certain bet with the implicit probabilities of the different markets. The goodness-of-fit measures indicate that the bet exchange prices better predict the actual match results. Second, we rerun this regression for all the bookmakers and include a variable capturing the difference of the two different markets' implicit probabilities. The estimated coefficient of this variable suggests that the bet exchange renders additional explanatory power beyond the bookmakers' odds. Finally, we assess the economic relevance of the previous results. A simple betting rule selecting bookmaker bets for which the average bookmaker offers lower probabilities (higher odds) than the bet exchange is capable of generating abnormal and in some cases even positive returns.

However, we are reluctant to interpret these findings as a failure of the bookmakers to process and impound relevant information into the prices. The underlying reasons for the comparably higher prediction accuracy of the bet exchange market are not clear a priori. Bettors endowed with more accurate information and beliefs may self-select into the exchange market while less skilled bettors may place their bets at the bookmaker setting. Alternatively, our findings could be due to the different market structures dealing

with similar but potentially biased demand. Bookmaker odds may not only reflect the dealer's true prediction of the outcome but also his (profit-maximizing) response to the expected (biased) demand. As Levitt (2004), Forrest and Simmons (2008) and Franck, Verbeek and Nüesch (2008) suggest, the bookmakers actively shade prices in the presence of a partly irrational betting audience in order to increase their profit. In regard to our findings, the price impact of a biased demand may be less pronounced in the person-to-person compared to the bookmaker market setting. Nevertheless, a proper examination of these suggestions lies beyond the scope of this paper and needs further research.

Acknowledgements

References

Berg, J., Nelson, F., & Rietz, T. (2008). Prediction market accuracy in the long run. *International Journal of Forecasting,* 24, 285-300.

Boulier, B. & Stekler, H. (2003). Predicting the outcomes of National football league games. *International Journal of Forecasting,* 19, 257-270.

Brier, G. (1950). Verification of weather forecasts expressed in terms of probability. *Monthly Weather Review,* 78, 1-3.

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.

Croxson, K. & Reade, J. (2008). Information and efficiency: Goal arrival in soccer betting. *Working Paper.* Oxford University.

Dixon, M. & Pope, P. (2004). The value of statistical forecasts in the UK association football betting market. *International Journal of Forecasting,* 20, 697-711.

Dobson, S. & Goddard, J. (2001). *The Economics of football*. Cambridge: Cambridge University Press.

Fan, J. (1992). Design-adaptive nonparametric regression. *Journal of the American Statistical Association*, 87, 998-1004.

Fan, J. & Gijbels, I. (1996). *Local Polynomial Modelling and its Applications*. London: Chapman & Hall.

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. (2008). Sentiment in the betting market on Spanish football. *Applied Economics,* 40, 119-126.

Forsythe, R., Nelson, F., Neumann, G. R., & Wright, J. (1992). Anatomy of an experimental stock market. *American Economic Review,* 82, 1142-1161.

Franck, E., Verbeek, E., & Nüesch, S. (2008). The sentiment bias in English soccer betting. *Working Paper,* University of Zurich.

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

Levitt, S. (2004). Why are gambling markets organized so differently from financial markets? *Economic Journal*, 114, 223-246.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In: P. Zarembka (Ed.), *Frontiers in econometrics*, New York: Academic Press, pp. 105-142.

McKelvey, R. & Zavoina, W. (1975). A statistical model for the analysis of ordinal level dependent variables. *Journal of Mathematical Sociology*, 4, 103-120.

Pennock, D., Lawrence, S., Nielsen, F., & Giles, D. (2001). Extracting collective probabilistic forecasts from web games. *Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining,* 174-183.

Smith, M., Paton, D., & Vaugham Williams, L. (2006). Market efficiency in person-to-person betting. *Economica*, 73, 673-689.

Snowberg, E., Wolfers, J., & Zitzewitz, E. (2008). Information (in)efficiency in prediction makets. In: L. Vaughan Williams (Ed.): *Information efficiency in financial and betting markets.* Cambridge: Cambridge University Press, pp. 366-386.

Spann, M. & Skiera, B. (2003). Internet-based virtual stock markets for business forecasting. *Management Science*, 49, 1310-1326.

Spann, M. & Skiera, B. (2008). Sports forecasting: A comparison of the forecast accuracy of prediction markets, betting odds and tipsters. *Journal of Forecasting*, 28, 55-72.

Vaughan Williams, L. (1999). Information efficiency in betting markets: a survey. *Bulletin of Economic Research, 51*, 1 - 39

Winkelmann, R. & Boes, S. (2005). *Analysis of microdata.* Heidelberg: Springer.

Wooldridge, J. (2002). *Econometric analysis of cross section and panel data*. Cambridge: The MIT Press.

Wolfers, J. & Leigh, A. (2002). Three tools for forecasting federal elections: lessons from 2001. *Australian Journal of Political Science*, 37, 223-240.

Wolfers, J. & Zitzewitz, E. (2004). Prediction markets. *The Journal of Economic Perspectives*, 18, 107-126.

Table 1
Summary statistics of outcome probabilities and forecasts

	true probabilities	Betfair	bookmaker
home win	0.462 (0.498)	0.456 (0.158)	0.448 (0.139)
draw	0.281 (0.449)	0.280 (0.048)	0.278 (0.038)
away win	0.256 (0.436)	0.263 (0.135)	0.273 (0.122)
observations	5478	5478	5478

Notes: The table presents the outcome probabilities and the forecasts of the exchange market and a randomly picked bookmaker. The mean and standard deviation are given. In terms of these summaries, the exchange market's probabilities are closer to the true outcome probabilities.

Table 2
The prediction accuracy of the bet exchange versus the bookmaker markets

	Betfair	average	B365	B&W	GB	IW	LB	WH	SJ	VC
home win bets										
impl. prob.	1,124 (0.047)	1,264 (0.054)	1,225 (0.052)	1,194 (0.051)	1,231 (0.052)	1,318 (0.057)	1,275 (0.054)	1,275 (0.054)	1,214 (0.063)	1,256 (0.065)
observations	5478	5478	5457	5475	5474	5442	5413	5438	3637	3602
McFadden's R2	0.082	0.079	0.078	0.077	0.079	0.077	0.079	0.079	0.080	0.082
McKelvey/Zavoina's R2	0.167	0.164	0.162	0.159	0.163	0.159	0.163	0.163	0.165	0.170
ROC area	0.6842	0.6800	0.6795	0.6777	0.6794	0.6772	0.6796	0.6790	0.6790	0.6824
Brier score	0.2221	0.2235	0.2236	0.2239	0.2236	0.2248	0.2238	0.2239	0.2227	0.2223
draw bets										
impl. prob.	1,298 (0.135)	1,621 (0.173)	1,451 (0.169)	1,414 (0.155)	1,521 (0.167)	1,491 (0.180)	1,627 (0.199)	1,505 (0.183)	1,461 (0.201)	1,520 (0.222)
observations	5478	5478	5457	5475	5474	5442	5413	5438	3637	3602
McFadden's R2	0.015	0.014	0.012	0.013	0.013	0.011	0.011	0.011	0.013	0.012
McKelvey/Zavoina's R2	0.033	0.032	0.027	0.030	0.030	0.024	0.025	0.025	0.023	0.027
ROC area	0.5688	0.5679	0.5612	0.5657	0.5669	0.5631	0.5635	0.5612	0.5727	0.5631
Brier score	0.1989	0.1993	0.1994	0.1993	0.1995	0.1993	0.1993	0.1993	0.1975	0.1972
away win bets										
impl. prob.	0.998 (0.044)	1,084 (0.049)	1,055 (0.047)	1,017 (0.047)	1,062 (0.048)	1,171 (0.053)	1,072 (0.049)	1,102 (0.049)	1,057 (0.058)	1,073 (0.059)
observations	5478	5478	5457	5475	5474	5442	5413	5438	3637	3602
McFadden's R2	0.088	0.083	0.082	0.070	0.083	0.082	0.082	0.084	0.083	0.084
McKelvey/Zavoina's R2	0.159	0.153	0.152	0.145	0.152	0.152	0.152	0.155	0.153	0.156
ROC area	0.6975	0.6900	0.6897	0.6866	0.6895	0.6887	0.6880	0.6915	0.6891	0.6901
Brier score	0.1714	0.1730	0.1732	0.1735	0.1728	0.1739	0.1736	0.1731	0.1761	0.1760

Notes: The table presents the prediction power for home wins (upper block), draws (middle block) and away win bets (lower block) of the bet exchange versus different bookmakers. The explaining variable is the probability implied by the odds of the different markets. The marginal effects of a probit regression (standard errors in parantheses) and a variety of goodness-of-fit measures are reported. It can be seen that the bet exchange (first column) outperforms every single bookmaker including the average bookmakers' prediction (second column) in terms of forecasting accuracy.

Table 3
The additional explanatory power of the bet exchange forecast

	average	B365	B&W	GB	IW	LB	WH	SJ	VC
home win bets									
impl. prob.	1.139***	1.107***	1.064***	1.110***	1.180***	1.147***	1.150***	1.084***	1.133***
	(0.062)	(0.061)	(0.059)	(0.061)	(0.065)	(0.063)	(0.063)	(0.072)	(0.074)
Betfair prob/ prob	0.409***	0.390***	0.449***	0.399***	0.431***	0.420***	0.405***	0.479***	0.447***
	(0.104)	(0.105)	(0.104)	(0.104)	(0.104)	(0.104)	(0.104)	(0.131)	(0.131)
observations	5478	5457	5475	5474	5442	5413	5438	3637	3602
McFadden's R2	0.081	0.080	0.079	0.080	0.079	0.080	0.080	0.082	0.084
draw bets									
impl. prob.	1.423***	1.223***	1.221***	1.318***	1.222***	1.338***	1.245***	1.225***	1.265***
	(0.194)	(0.193)	(0.174)	(0.189)	(0.209)	(0.224)	(0.207)	(0.229)	(0.255)
Betfair prob/ prob	0.228*	0.254*	0.250*	0.235*	0.266*	0.294**	0.278**	0.274*	0.258*
	(0.102)	(0.104)	(0.102)	(0.103)	(0.106)	(0.103)	(0.103)	(0.126)	(0.127)
observations	5478	5457	5475	5474	5442	5413	5438	3637	3602
McFadden's R2	0.015	0.012	0.014	0.014	0.012	0.012	0.012	0.014	0.012
away win bets									
impl. prob.	0.869***	0.841***	0.801***	0.851***	0.926***	0.844***	0.884***	0.854***	0.873***
	(0.057)	(0.056)	(0.054)	(0.056)	(0.062)	(0.057)	(0.058)	(0.067)	(0.068)
Betfair prob/ prob	0.441***	0.444***	0.467***	0.438***	0.462***	0.474***	0.442***	0.456***	0.445***
	(0.064)	(0.064)	(0.063)	(0.064)	(0.064)	(0.064)	(0.064)	(0.079)	(0.079)
observations	5478	5457	5475	5474	5442	5413	5438	3637	3602
McFadden's R2	0.090	0.090	0.088	0.090	0.090	0.090	0.091	0.090	0.091

Notes: The table presents the additional prediction power which is rendered by the bet exchange for home win bets (upper block), for draw bets (middle block) and away win bets (lower block). The explaining variables are the probabilities implied by the odds of the different bookmakers and the ratio between the bet exchange probability and the bookmaker probability. The marginal effects of a probit regression (standard errors in parantheses) are given. ***, **, * denotes signifcance at the 0.1% , 1%, 5% level respectively. It can be seen that the bet exchange probabilities contain relevant information which is not fully reflected by the bookmakers' odds.

Table 4
The mean returns of a simple betting strategy compared to average returns

bookmaker	all events		home win bets		draw bets		away win bets	
	all	p_BF>p	all	p_BF>p	all	p_BF>p	all	p_BF>p
random	-0.124 (16434)	-0.028 (8234)	-0.084 (5478)	-0.027 (3219)	-0.096 (5478)	-0.053 (3339)	-0.192 (5478)	0.019 (1676)
highest odd	-0.072 (16434)	0.014 (8234)	-0.037 (5478)	0.012 (3219)	-0.055 (5478)	-0.016 (3339)	-0.124 (5478)	0.082 (1676)
B365	-0.109 (16371)	-0.019 (8203)	-0.069 (5457)	-0.019 (3205)	-0.083 (5457)	-0.043 (3329)	-0.177 (5457)	0.027 (1669)
B&W	-0.111 (16425)	-0.020 (8229)	-0.074 (5475)	-0.021 (3217)	-0.085 (5475)	-0.045 (3338)	-0.174 (5475)	0.030 (1674)
GB	-0.109 (16422)	-0.017 (8228)	-0.067 (5474)	-0.015 (3217)	-0.086 (5474)	-0.045 (3337)	-0.175 (5474)	0.032 (1674)
IW	-0.141 (16326)	-0.039 (8177)	-0.084 (5442)	-0.022 (3203)	-0.121 (5442)	-0.081 (3313)	-0.218 (5442)	0.009 (1661)
LB	-0.134 (16239)	-0.034 (8135)	-0.096 (5413)	-0.039 (3185)	-0.102 (5413)	-0.051 (3292)	-0.204 (5413)	0.009 (1658)
WH	-0.137 (16314)	-0.042 8169)	-0.089 (5438)	-0.033 (3196)	-0.115 (5438)	-0.071 (3311)	-0.208 (5438)	-0.002 (1662)
SJ	-0.111 (10911	-0.030 (5524)	-0.88 (3637)	-0.045 (2083)	-0.103 (3637)	-0.063 (2301)	-0.141 (3637)	0.062 (1140)
VC	-0.125 (10806)	-0.039 (5471)	-0.090 (3602)	-0.046 (2068)	-0.108 (3602)	-0.068 (2280)	-0.176 (3602)	0.032 (1123)

Notes: The table compares the mean returns for a simple betting strategy (right hand side of each column) with normal returns (left hand side of each column). The number of bets is displayed in parentheses. The trading rule is to place a bet at a given bookmaker whenever the probability of *Betfair* is higher than the average probability of the bookmakers. The results are broken down by the events on which to place a bet (columns) and the bookmakers (rows). The first row presents the results for a randomly chosen bookmaker and the second row for the bookmaker offering the most favorable odd. It can be seen that the rule enables above-average returns in all cases and, in some cases, even positive returns.
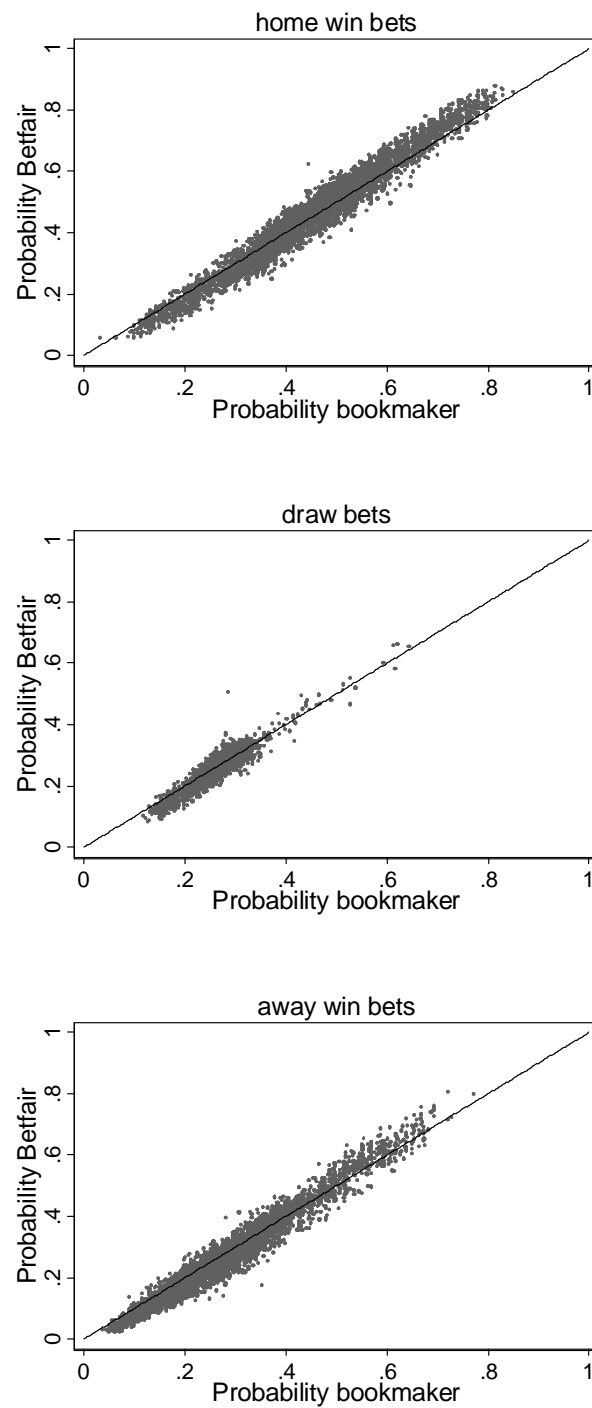
Fig. 1. The implicit probabilities of *Betfair* plotted against a random bookmaker. The black line indicates the cases for which the probabilities of the two markets are equal.
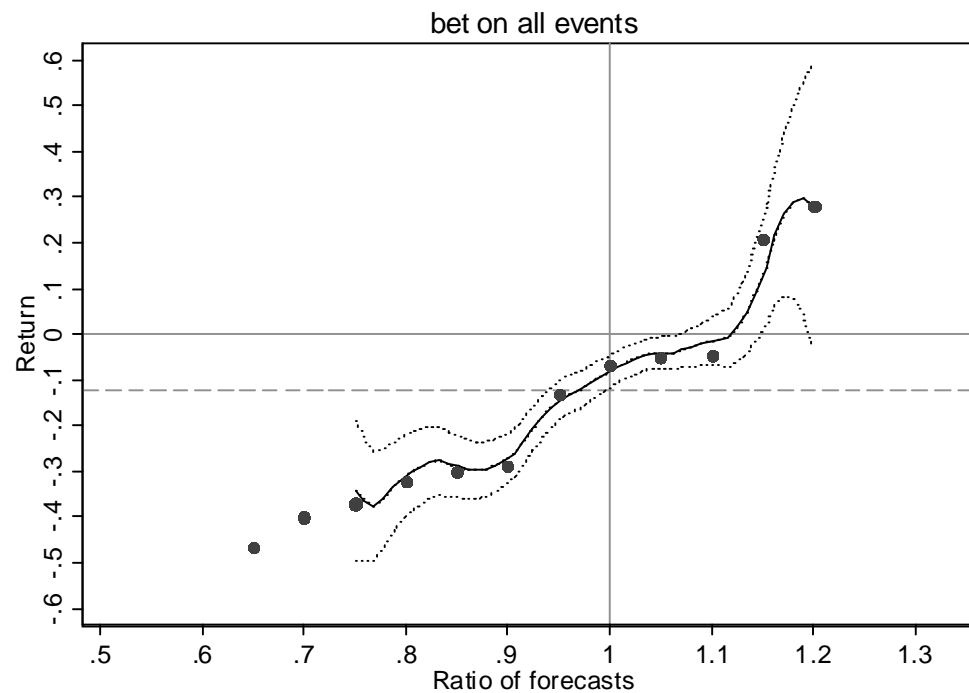
Fig. 2. Observed mean returns of bets placed at a random bookmaker plotted against the ratio of the different markets' implicit probabilities for all events. The lines are for a zero return (solid horizontal line) and the expected return under random betting (broken line). The mean returns for the different categories of the ratio (dots) and the local polynomial smooth (solid line) with 95% confidence intervals (dotted lines) are given.
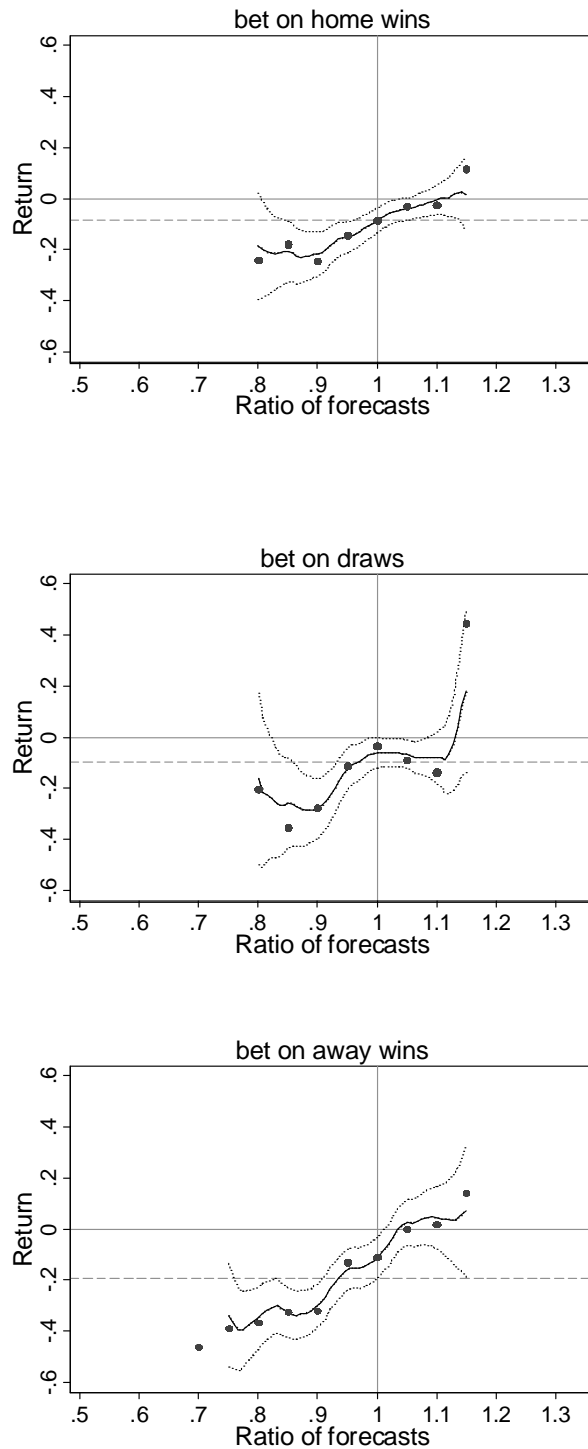
Fig. 3. Observed mean returns of bets placed at a random bookmaker plotted against the ratio of the different markets' implicit probabilities for home win, draw and away win bets separately. The lines are for a zero return (solid horizontal line) and the expected return under random betting (broken line). The mean returns for the different categories of the ratio (dots) and the local polynomial smooth (solid line) with 95% confidence intervals (dotted lines) are given.