

EFFICIENCY IN BETTING MARKETS: EVIDENCE FROM ENGLISH FOOTBALL

Bruno Deschamps† and Olivier Gergaud‡

†*University of Bath*

‡*University of Reims*

We analyze the efficiency of English football betting markets between 2002 and 2006. We find evidence of a positive favourite-longshot bias for both home odds and away odds. Draw odds are instead characterized by a negative longshot bias. We also identify a draw bias in the sense that betting at draw odds yields a higher return than betting at home or away odds. Finally, we investigate betting strategies that exploit the variance of odds between bookmakers.

1. INTRODUCTION

There exists a large empirical literature analyzing the existence of weak form efficiency in betting markets.¹ The most widely documented inefficiency is the longshot bias. In horse races betting markets, betting on favourites yields a higher return than betting on the longshots.² Evidence of a favourite-longshot bias has been found in other sports including football. In UK football, the evidence is mixed. Forrest, Goddard and Simmons (2005), using a sample of nearly 10,000 football matches played between 1998 and 2003, find no evidence of a longshot bias. Kuypers (2000) focus on the 1993/1994 and 1994/1995 seasons and analyze the odds for 3382 matches. They do not find evidence of a bias either. Dixon and Pope (2004) use data for three seasons (1993/1994 to 1995/1996) and a total of 6629 matches. They find evidence of a negative longshot bias: betting on the favourites yields a lower return than betting on the longshots. Cain, Law and Peel (2000) analyze betting efficiency for the 1991/1992 season, for a total of 2855 matches. They find evidence of a positive longshot bias.

This paper analyzes the existence of weak form efficiency in English football betting markets. Weak form efficiency implies that no abnormal return can be achieved by using only price information. We base our analysis on 8377 matches played during four seasons (from 2002/2003 to 2005/2006). For each match we have the odds on the match outcomes (home win, draw and away win) for six different bookmakers, and the actual outcomes. We first investigate whether the odds are characterized by a longshot bias. After controlling for the odds status (home win, draw and away win), we find mixed evidence. For both home win odds and away win odds, there is a clear positive longshot bias. Betting at longer odds generates a lower return than betting on the short odds. For draw odds, instead, we find a reverse bias. Betting at longer odds yields a higher return than betting at shorter odds. We show that despite the bias, there exists no betting strategy based solely on the odds that has a positive return. We find another bias, in that betting on draw odds yields a much higher return (-7%) than betting on away win odds (-14%) or home win odds (-11%).

We then investigate whether more sophisticated betting strategies can be profitable and have a positive return. First, we analyze the strategy of choosing systematically the best available odds among the six bookmakers. It turns out that such a strategy improves significantly the return of punters but still has a negative return of -6.8% . Combining that strategy with betting only on draw odds offers a return of -4% , which is significantly higher but still far from being profitable.

Finally, we consider a more elaborate strategy. The idea is to consider that the variance of the odds between bookmakers may act as a signal that the best available odd is too generous. Our hypothesis is that if a bookmaker offers much better odds than other bookmakers for the same event, it might be that this bookmaker has underestimated the probability that this event occurs. In that case, a betting strategy consisting of betting on the best available odds if and only if the variance of odds across bookmakers is high enough could be profitable. We find that the variance of odds, or "disagreement" between bookmakers is a variable that can be used to generate a higher return. However, the return is still negative.³

The layout of this paper is the following. Section 2 describes the data. Section 3 investigates the longshots bias and the return that can be achieved by choosing the best available odds. Section 4 analyzes how the variance of odds between bookmakers can be exploited to achieve a higher return. And Section 5 concludes.

2. DATA

We have collected the full-time results of 8377 Football League matches played in England during four seasons, from 2002/2003 to 2005/2006. For each match we have the quoted outcome odds (home win, draw and away win) of six bookmakers. The bookmakers are: Bet365 (B365 henceforth), Gamebookers (GB), Interwetten (IW), Ladbrokes (LB), Sporting Bets (SB) and William Hill (WH). We use the following notation. If an odd is q , a successful bet with a size of one yields a profit of $q-1$. For instance, betting on an outcome with an odd of 2.2 will yield a profit of 1.2 and a return of 120% if the bet is successful. Given that we have a total of 8377 matches, three odds for each match and six bookmakers, there is a total of around 150,000 observations. Our dataset is relatively large compared to most existing studies. Table 1 shows some summary statistics. Overall, the average odds are quite stable between 2002 and 2006.

We analyze how the margins of the bookmakers have changed between 2002 and 2006. Note that, contrary to horse races, football odds are fixed, meaning that they are set before the match and are not affected by betting volumes. Hence, the margin in football betting is not a fixed percentage of the total amount bet. The theoretical margin can nevertheless be estimated by the 'over-round' implied in the odds. The over-round λ is defined as the

EVIDENCE FROM ENGLISH FOOTBALL

TABLE 1
SUMMARY STATISTICS

Year	2002/2003	2003/2004	2004/2005	2005/2006
Matches	1919	1985	1995	2478
Average home win odds	2.11	2.15	2.15	2.16
Average draw odds	3.30	3.33	3.32	3.30
Average away win odds	3.55	3.62	3.63	3.55

difference between the sum of the inverse of the odds and one.

$$(1) \quad \lambda_{mi} = \sum_j \frac{1}{q_{mij}} - 1$$

where m stands for match m , i stands for bookmaker i and j refers to the odd status ($j = 1, 2$ or 3 for home, draw or away odds respectively). The shorter the odds, the higher the over-round λ , and the higher the margin of the bookmaker i .⁴ Table 2 shows that the average margin has been reduced between 2002 and 2006, meaning that bookmakers offer better odds in 2006 than in 2002. The mean value has actually decreased from 12.32% in 2002/2003 to 11.56% in 2005/2006. This trend could be the result of greater competition in the UK betting market, in particular given the growing popularity of internet betting.

3. ODDS EFFICIENCY

In this section we start by investigating whether the longshot bias can be observed in our dataset. In Section 3.1 we find mixed evidence of the bias by pooling all 150,000 observations together. In Section 3.2 we show that the direction of the bias depends on the odds status. Section 3.3 documents the existence of a draw bias. Finally, Section 3.4 analyzes the betting strategy of selecting the best available odds.

TABLE 2
EVOLUTION OF THE MARGINS

	2002/2003	2003/2004	2004/2005	2005/2006
B365	11.65%	10.75%	10.44%	10.34%
GB	10.28%	9.95%	9.82%	10.46%
IW	15.76%	16.44%	14.68%	14.10%
LB	12.39%	12.36%	12.36%	12.31%
SB	11.36%	11.57%	10.31%	9.71%
WH	12.50%	12.50%	12.47%	12.49%
Mean	12.32%	12.26%	11.68%	11.56%

3.1 Longshot bias

In order to find a longshot bias, positive or negative, we first pool all seasons (from 2002/2003 to 2005/2006) and all odds status together. If the betting market was efficient, the odds should reflect the outcome probability. No betting strategy based on the odds level should have abnormal return. In order to check the existence of a bias, we calculate the implicit probability ϕ_{mij} for every single odds level.

$$(2) \quad \phi_{mij} = \frac{1}{q_{mij}} \frac{1}{\sum_j \frac{1}{q_{mij}}}$$

where m stands for match m , i stands for bookmaker i and j is the odd status. The implicit probability formula is such that for each match, the sum of the implicit probabilities $\sum_j \phi_{mij}$ is equal to one for each bookmaker.⁵ We then rank the 150,000 observations according to their implicit probability and we split them into 20 categories of equal size. The first odds category includes the odds whose implicit probability is the lowest (i.e. the longshots), while the 20th odds category includes the shortest odds, and so on. For each category, we compute the actual return of a strategy consisting of betting £1 at each odds level in the category. The results are shown in Table 3 and in Figure 1.

Figure 1 shows the relationship between odds and return. It provides mixed evidence on the longshot bias. On the one hand, the two categories with the longest odds (on the left) have the lowest return (−23% and −19.3% respectively), which is consistent with a longshot bias. On the other hand, past the first two categories, the relationship between odds and return is not well defined. In particular, category 5 has the highest return, which is inconsistent with a longshot bias. Overall, the relationship between odds and return is non-monotonic and there is some evidence of a longshot bias for the very long odds. We believe that the non-monotonicity is due to the fact that we have not controlled for the odds status. Indeed, Section 3.2 shows that the relationship

TABLE 3
IMPLICIT PROBABILITY AND RETURN

Implicit probability	Return	Implicit probability	Return
0–0.180	−23%	0.291–0.307	−12.2%
0.180–0.221	−19.3%	0.308–0.325	−11.7%
0.222–0.246	−5.9%	0.326–0.351	−15.3%
0.246–0.260	−7.8%	0.352–0.368	−17.6%
0.261–0.269	−1.5%	0.369–0.396	−12.1%
0.269–0.275	−8.9%	0.397–0.412	−10%
0.275–0.277	−4.7%	0.413–0.459	−8.9%
0.277–0.279	−10.5%	0.460–0.493	−12.2%
0.279–0.283	−8.2%	0.494–0.551	−10.6%
0.283–0.290	−11.3%	0.552–1	−6.9%

EVIDENCE FROM ENGLISH FOOTBALL

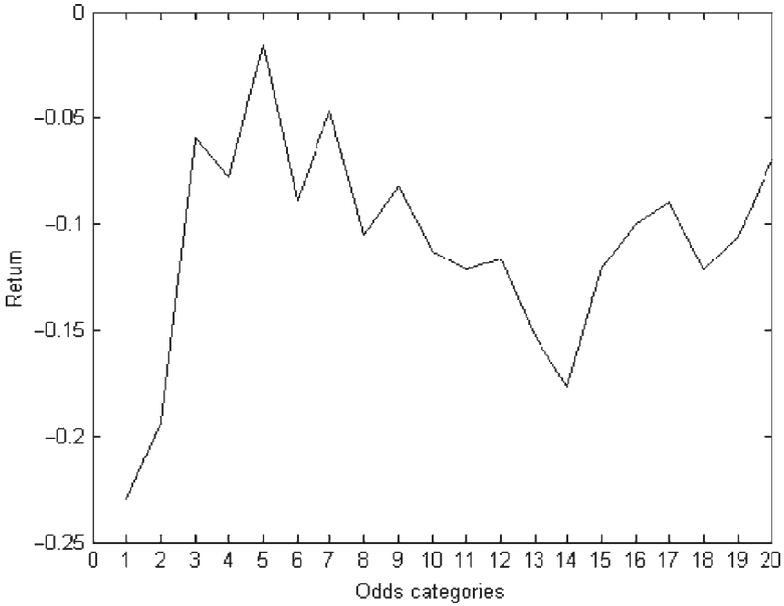


FIGURE 1. Odds categories and Return.

between odds and return is different for each odds status. Pooling all odds status together generates the non-monotonicity. Note that there is no range of odds that has a positive return, and therefore there is no profitable betting strategy based solely on the odds categories.

3.2 Longshot bias and odds status

In this section we investigate how the relationship between odds and return is affected by the odds status. We split the dataset into three groups (home win odds, draw odds and away win odds) and we analyze each group separately. The methodology is similar to the one previously used in Section 3.1. For each odds status, we rank the odds according to their implicit probability and we split them into five categories. The results are shown in Table 4 and in Figures 2–4.

TABLE 4
IMPLICIT PROBABILITY AND RETURN BY ODD STATUS

Home		Draw		Away	
Imp. proba	Return	Imp. proba	Return	Imp. proba	Return
0–0.36	– 13.74%	0–0.25	– 5.38%	0–0.20	– 23.20%
0.37–0.40	– 9.64%	0.26–0.275	– 5.95%	0.21–0.25	– 10.58%
0.41–0.46	– 10.43%	0.276–0.279	– 5.85%	0.26–0.30	– 12.24%
0.47–0.52	– 10.33%	0.28–0.284	– 8.02%	0.31–0.35	– 14.10
0.53–1	– 9.42%	0.285–1	– 11.73%	0.36–1	– 13.06%

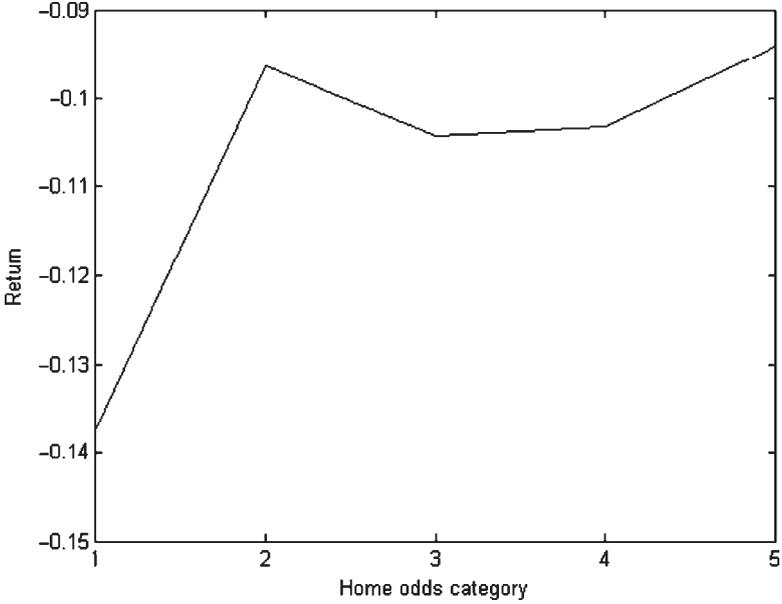


FIGURE 2. Home odds categories and Return.

Let us first consider home win odds. Figure 2 shows the presence of a longshot bias. The first category of odds (the longest odds) has a return close to -14% , which is less than the return for any other category. Note that there is not a large return difference between categories 2, 3, 4 and 5. This longshot

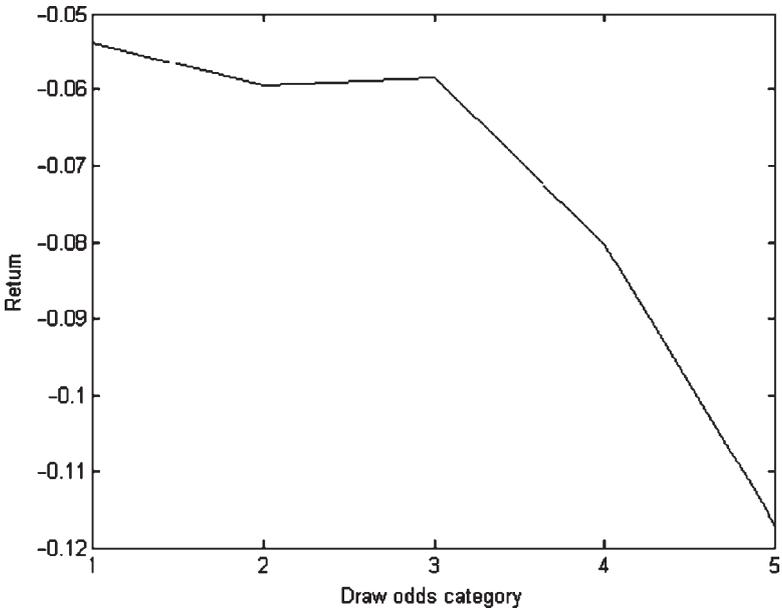


FIGURE 3. Draw odds categories and Return.

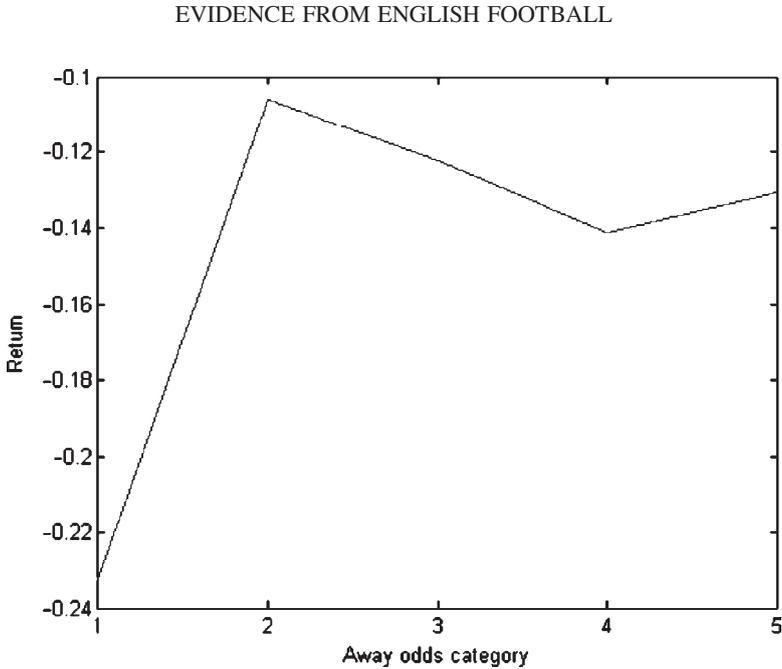


FIGURE 4. Away odds categories and Return.

bias consists thus of a return differential between the very long odds and the shorter odds. Evidence of a longshot bias can also be found among away win odds, as Figure 4 shows. The return for the first odds category is much lower than the return of the other four categories. Note that here also, there is no significant difference between the return of categories 2, 3, 4 and 5.

Interestingly, we find the opposite result for the draw odds, as Figure 3 shows. We find a negative longshot bias: long odds have a relatively high return (-5%), while shorter odds have a much lower return (-11%). Our data suggest a natural explanation for this reverse bias, which is that draw outcomes are extremely difficult to predict. We indeed find that there is virtually no relationship between the draw odds and the probability of a draw outcome. This is unique to draw odds, since home and away odds are strongly correlated with the probability of home and away win. Draw odds, on the opposite, are totally uninformative. This naturally explains why we find a negative longshot bias. Note that the lack of relationship between draw odds and draw outcomes is also documented by Cain, Law and Peel (2000) in a study that analyzes the 1991/1992 football season.

Overall, we find mixed evidence. The relationship between odds and return depends totally on the odds status. For both home and away win odds, very long odds have a lower return than shorter odds. However, besides the very long odds, there is no significant difference in return between odds categories. Draw odds are characterized by a negative bias. All these results suggest that odds are inefficient in the sense that betting strategies based solely on the odds categories and odds status can generate abnormal return.

For instance, betting on long draw odds yields a much higher return (-5%) than betting on long away win odds (return of -23%). According to our results, however, there exists no profitable betting strategy based on the odds level or odds status.

3.3 Draw bias

Additionally to the positive/negative longshot bias, we have found another bias in the odds return, we call it the draw bias. It turns out that between 2002 and 2006, betting on draws has generated a higher return (-7%) than betting on home win (-10.8%) and away win (-14.1%). Such a strong bias has, to our knowledge, not been found in previous studies. This is probably due to the fact that this bias has appeared recently, as shown in Table 5. There was indeed no draw bias during the 2002/2003 season, since the return for draw odds was similar to the return for away and home odds, at around -11% . There is a small draw bias in 2003/2004 and it gets larger in 2004/2005 and 2005/2006. Our data indicate that this bias cannot be explained by more generous odds for draws in 2005/2006. The average draw odds has not changed much between 2002/2003 and 2005/2006. However, the frequency of draw outcomes has increased slightly during the same period, which explains why the return at draw odds has increased. It appears that while draw outcomes have started to become more frequent since 2003/2004, bookmakers have failed to adjust the odds accordingly.

3.4 Best available odds

In this section we consider an alternative betting strategy. Imagine that, instead of betting on all available odds, punters choose to bet on the best available odds in the market. With the growth of internet betting, it has indeed become easier for punters to choose the best available odds in order to maximize their return. We investigate whether such a betting strategy can yield a positive return, and how it affects the longshot bias. For each match and each odd status, we estimate the best available odd by taking the maximum odd among the six bookmakers of our dataset. Obviously, in the real world there are more than six bookmakers, so there might be bookmakers offering slightly better odds for some matches. This should however not affect our result on the longshot bias nor our results on the evolution of the odds.

TABLE 5
DRAW BIAS

	2002/2003	2003/2004	2004/2005	2005/2006
Home win odds return	-11.9%	-9%	-10.8%	-11.8%
Draw odds return	-11.5%	-7.5%	-4.8%	-4.2%
Away win odds return	-10.3%	-15.7%	-15.8%	-14.5%

EVIDENCE FROM ENGLISH FOOTBALL

TABLE 6
BEST AVAILABLE ODDS AND RETURN

Home		Draw		Away	
Average Odd	Return	Average Odd	Return	Average Odd	Return
1.52	-7.7%	3.22	-13.8%	2.21	-8.2%
1.75	-9.6%	3.25	-8.9%	2.66	-9%
1.92	-11.3%	3.26	-0.6%	2.94	-9.8%
2.14	-6.1%	3.32	4.5%	3.36	-9.4%
2.30	-2.3%	3.4	-4.2%	3.90	5.3%
2.51	-10%	3.48	-4.8%	4.62	-15.3%
3.39	-6.8%	3.92	-1.1%	7	-17.5%

For each odds status, we sort the best available odds from the shortest to the highest and we split them into seven categories. The results are shown on Table 6.⁶ We find that for the away win odds, the last two categories have the lowest return, -15.3% and -17% respectively, which is again consistent with the longshot bias. For the draw odds, we find a negative longshot bias, as the first two odds categories have the lowest return. Note that there is no bias for home odds. Overall, bettors choosing the best available odds also face a positive or negative longshot bias, except for home odds.

On the return side, there is no clear strategy that yields a significant positive return even though there are positive returns for two categories of odds. The best strategy is to bet on the best available long draw odds. This yields a return close to zero.

Table 7 shows that the return on the best available odds has not been constant between 2002 and 2006. The return has decreased between 2002/2003 and 2003/2004 and has increased afterwards. This is puzzling at first sight. Indeed, Table 2 shows that the margins of the bookmakers have decreased monotonically between 2002/2003 and 2005/2006. Lower margins should imply higher return, however this is not the case in 2003/2004. The reason for this inconsistency is that the return on the best available odds depends both on the gross margins of the bookmakers but also on the disagreement between bookmakers. Our results suggest that in 2003/2004, the positive effect of the lower margins has been smaller than the negative effect of lower disagreement between bookmakers.

We also find that the return on the best available draw odds has increased dramatically during that period. From -8.2% in 2002/2003, to -2.6% in

TABLE 7
EVOLUTION OF THE RETURN ON THE BEST AVAILABLE ODDS

	2002/2003	2003/2004	2004/2005	2005/2006
Return on the best available odds	-6.9%	-7.2%	-7%	-6.4%

2004/2005 and -1.1% in 2005/2006. Betting on draws has thus been nearly profitable in the last two year of our dataset. A strategy of combining this higher return with the negative longshot bias does not, however, generate a positive return.

4. DISAGREEMENT BETWEEN BOOKMAKERS

We have shown that some betting strategies generate abnormal returns; however none of them is profitable. Betting strategy based on the odds status (home, draw or away) or the odds level (long or short) have at best a return close to zero. Even the best available odds have a negative return. This section investigates a more sophisticated betting strategy. The basic principle is to use the variation of odds across bookmakers, that is the "level of disagreement". This disagreement variable will then be used to identify the odds that might be too generous and have a positive expected return. The strategy consists of betting on the best available odds if and only if the level of disagreement between bookmakers is high enough. The intuition is simple. When bookmakers offer very different odds for a match outcome, then, unless the average bookmaker is wrong, the bookmaker offering the best available odds is likely to have set a wrong price. If, instead, bookmakers offer very similar odds for a given outcome, then the best available odd is close to the average odd and, given the 11–12% margin, no profit can be made. The more the bookmakers disagree, the most likely it is that the best available odds are too long. For instance, if the best available odds level is 10 when the market average is 9.5, we may assume that 10 is not profitable. If instead the market average odds is 5, then 10 looks like a very good opportunity to make a positive return. Hence, our hypothesis is that a strategy based both on the best available odds and on the disagreement between bookmakers can generate a positive return. In this section we check whether this hypothesis is correct. Our hypothesis is strongly related to the contributions of Paton and Vaughan Williams (2005) and Smith, Paton and Vaughan Williams (2005). They examine the concept of Quasi-Arbitrage opportunities (Quarbs) for UK football and horse races. The principle is to assume that the market average price is a good indicator of the objective probability, while the outlier price is not. They show that, in many cases, it is possible to use the outlier price to generate positive return and that betting strategies based on the outlier price can be profitable.

In order to measure the level of disagreement between bookmakers, we calculate the dispersion of the odds for each match and for each odd status by the mean absolute deviation of the odds across the six bookmakers (Mad_{jm}).

$$(3) \quad Mad_{jm} = \frac{1}{6} \sum_i |q_{ijm} - \bar{q}_{jm}|$$

where \bar{q}_{jm} is the average odd for match m and odd status j , and where q_{ijm} is the odd of bookmaker i for match m and odd status j . The higher is Mad_{jm} ,

EVIDENCE FROM ENGLISH FOOTBALL

the more the bookmakers disagree on the odds. As in Section 3.2, we pursue the analysis for each odd status separately, which means that we have a total of 8377 Mad_{jm} for each odd status j . Note that Mad_{jm} is naturally strongly correlated to \bar{q}_{jm} . For instance, if the odds are around 10/1, the mean absolute deviation is naturally higher than if the odds are 2/1. Said differently, the longer the odds, the higher their mean absolute deviation. Therefore, in order to measure the level of disagreement between bookmakers, we need to remove the effect of \bar{q}_{jm} . This is done by regressing (for each odd status) Mad_{jm} on \bar{q}_{jm} and \bar{q}_{jm}^2

$$(4) \quad Mad_{jm} = \beta_0 + \beta_1 \bar{q}_{jm} + \beta_2 \bar{q}_{jm}^2 + \varepsilon_{jm}$$

Note that $\beta_1 > 0$ and $\beta_2 > 0$ for home, draw and away odds. We estimate the level of disagreement by the residuals ε_{jm} . If $\varepsilon_{jm} > 0$ then the level of disagreement between bookmakers is high given \bar{q}_{jm} . If $\varepsilon_{jm} < 0$ then the level of disagreement is low given \bar{q}_{jm} . For each odd status, we rank the matches according to the disagreement level ε_{jm} . We then split the matches in five categories of disagreement⁷: very high, high, average, low and very low. Our hypothesis is that if there is more disagreement, the best available odd is excessively generous and possibly profitable (positive return). Symmetrically, if there is less disagreement, the best available odds are less generous and therefore less profitable (lower return). The results are shown in Table 8 and are overall consistent with our prediction. For each odds status, there tend to be a positive relationship between disagreement and return. For both home odds and away odds, betting on the matches with very high disagreement generates the highest return (-3.1% and -4.3% respectively). For draw odds, betting on the matches with very low disagreement yields a very low return of -10.7% , which is much less than the return for any other category. This results are consistent with our intuition that high disagreement is a signal that the best available odds level on offer is rather generous. Comparing Table 8 to Table 6 shows that a strategy based on the level of disagreement provides a higher return than a strategy based on the odds, in particular for home and away odds. Note that once again, no profitable strategy emerges

TABLE 8
 DISAGREEMENT AND RETURN

Home		Draw		Away	
Disagreement	Return	Disagreement	Return	Disagreement	Return
Very low	-8.7%	Very low	-10.7%	Very low	-14.3%
Low	-7.8%	Low	-2.1%	Low	-7%
Average	-9.9%	Average	-2.7%	Average	-10.8%
High	-7.6%	High	-1.2%	High	-8.6%
Very high	-3.1%	Very high	-3.6%	Very high	-4.3%

from these results. The level of disagreement can be used to increase the return, but this is not sufficient to make a profit.

5. CONCLUSION

This study has analyzed the efficiency of football betting in England between 2002 and 2006. We have found that the betting market is inefficient in that several betting strategies generate abnormal return. Among our results, it turns out that the bias between favourites and longshots depends very much on the odds status. For both home and away odds, there is a clear positive longshot bias. For draw odds, instead, there is a negative longshot bias. Another bias is what we call the draw bias: Draw odds yield a much higher return than home or away odds. This bias is particularly large in 2005 and 2006. We have shown that more elaborate betting strategies may have high abnormal returns. In particular, the variance of odds between bookmakers can be exploited to earn abnormal return, even though this is not sufficient to make profit. We have found that when there is "high disagreement" between bookmakers, the return on the best available odd is usually higher than when bookmakers offer similar odds. Note that despite the existence of better betting strategies, none of them has a significant positive return.

Interestingly, we have found some trends between 2002 and 2006. First, the margins of the bookmakers have decreased slightly. Second, the return on draw odds have increased every year. It would be interesting in the future to analyze whether these trends have persisted. This is particularly important given that several of the betting strategies that we have analyzed had a return close to zero. If the margins of the bookmakers continue to go down, we could expect some betting strategies to have positive return. If instead the margins stop falling and bookmakers react to the existing biases by adjusting their odds, then the return on the most profitable strategies would go down.

NOTES

1. For a literature review of weak form efficiency in betting markets, see Vaughan Williams (2005).
2. See Thaler and Ziemba (1988).
3. Paton and Vaughan Williams (2005) and Smith, Paton and Vaughan Williams (2005) also analyze betting strategies based on the outlier odds. They find that such strategies may generate profit.
4. This assumes that the book is balanced, so that the outcome of the match has no impact on the bookmakers' return.
5. Indeed, $\sum_j \phi_{mij} = \sum_j \frac{1}{q_{mij}} \sum_j \frac{1}{q_{mij}} = 1$
6. Naturally, there might be bookmakers offering slightly better odds on some matches, so the return for a punter searching for the best available odds in the market are higher than in Table 6. But on the other hand, real world punters do not necessarily have the opportunity or the will to search for the best odd among tens of bookmakers.
7. Each category includes 1675 or 1676 matches.

EVIDENCE FROM ENGLISH FOOTBALL

REFERENCES

- M Cain, D Law and D Peel 'The favourite-longshot bias and market efficiency in UK football betting' *Scottish Journal of Political Economy* (2000) 47 25–36.
- M Dixon and P Pope 'The value of statistical forecasts in the UK association football betting market' *International Journal of Forecasting* (2004) 20 697–711.
- D Forrest, J Goddard and R Simmons 'Odds-setters as forecasters: the case of English football' *International Journal of Forecasting* (2005) 21 551–564.
- T Kuypers 'Information efficiency: an empirical study of a fixed odds betting market' *Applied Economics* (2000) 32 1353–1363.
- D Paton and L Vaughan Williams 'Forecasting outcomes in spread betting markets: can bettors use quarks to beat the book' *Journal of Forecasting* (2005) 24 139–154.
- M Smith, D Paton and L Vaughan Williams 'An assessment of quasi-arbitrage opportunities in two fixed-odds horse-race betting markets' in Vaughan Williams (ed) *Information Efficiency in Financial and Betting Markets* (Cambridge University Press, 2005) pp 159–171.
- R Thaler and W Ziemba 'Parimutuel betting markets: racetracks and lotteries' *Journal of Economic Perspectives* (1988) 2 161–174.
- L Vaughan Williams 'Weak form information efficiency in betting markets' in Vaughan Williams (ed) *Information Efficiency in Financial and Betting Markets* (Cambridge University Press, 2005) pp 84–123.