Accepted Manuscript

Title: Do Bookmakers Possess Superior Skills to Bettors in Predicting Outcomes?

Authors: Michael A. Smith, David Paton, Leighton Vaughan Williams

PII: S0167-2681(09)00083-3
DOI: doi:10.1016/j.jebo.2009.03.016
Reference: JEBO 2355

To appear in: *Journal of Economic Behavior & Organization*

Received date: 19-7-2007
Revised date: 19-3-2009
Accepted date: 21-3-2009

Please cite this article as: Smith, M.A., Paton, D., Williams, L.V., Do Bookmakers Possess Superior Skills to Bettors in Predicting Outcomes?, *Journal of Economic Behavior and Organization* (2008), doi:10.1016/j.jebo.2009.03.016

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Do Bookmakers Possess Superior Skills to Bettors in Predicting Outcomes?

Michael A. Smith

Senior Lecturer in Economics
Leeds Business School
Leeds Metropolitan University
Headingley Campus
Leeds LS1 3HE
United Kingdom
Tel: +44 113 812 7537
Fax: +44 113 812 8604
Email: m.z.smith@leedsmet.ac.uk

David Paton*

Professor of Industrial Economics
Nottingham University Business School
Wollaton Road
Nottingham NG8 1BB
United Kingdom
Tel: +44 115 846 6601
Fax: +44 115 846 6667
Email: David.Paton@Nottingham.ac.uk

and

Leighton Vaughan Williams

Professor of Economics and Finance
Nottingham Business School
Nottingham Trent University
Burton Street
Nottingham NG1 4BU
United Kingdom
Tel: +44 115 848 6150
Email: Leighton.Vaughan-Williams@ntu.ac.uk

* corresponding author

Acknowledgements

We would like to thank participants at the Royal Economic Society Conference (Warwick, March 2008); the Scottish Economic Society Conference (Perth, April 2008); and a staff seminar at the University of Kent at Canterbury (November 2007) for constructive advice on our methodology. We would also like to thank two anonymous referees for a number of useful comments and suggestions.

Do Bookmakers Possess Superior Skills to Bettors in Predicting Outcomes?

Abstract

In this paper we test the hypothesis that bookmakers display superior skills to bettors in predicting the outcome of sporting events by using matched data from traditional bookmaking and person-to-person exchanges. Employing a conditional logistic regression model on horse racing data from the UK we find that, in high liquidity betting markets, betting exchange odds have more predictive value than the corresponding bookmaker odds. To control for potential spillovers between the two markets, we repeat the analysis for cases where prices diverge significantly. Once again, exchange odds yield more valuable information concerning race outcomes than the bookmaker equivalents.

Keywords: betting exchanges; market efficiency; prediction.

JEL Classification: D82, G12, G14.

1. Introduction

There have been many studies of the efficiency of horse race betting markets, based either on bettor determined prices (as in pari-mutuel markets) or bookmaker determined prices, the latter mainly based on UK data (see, for example, Smith et al., 2006). Most studies of bookmaker markets indirectly infer the superiority of bookmaker skills over bettor skills from the existence of persistent negative returns to bettors in aggregate. However, if bettors receive consumption utility from placing wagers in addition to utility from monetary returns, bettor superiority may be consistent with aggregate negative returns. Furthermore these studies tell us nothing about the abilities of bettors who choose to refrain from entering the market when they judge that bookmaker prices overstate the true chances of race entrants.

In this paper we use matched data from traditional bookmaking and person-to-person exchanges to test the hypothesis that bookmakers display skills superior to bettors in predicting the outcome of sporting events. One might expect on the basis of the already extensive literature in the economics of auctions (e.g. Klemperer, 1999, 2004) that the decentralised nature of the decision-making processes characteristic of betting exchanges would accomplish the aggregation of dispersed information in a very efficient manner, whereas the bookmaker (however well informed) may fail to match as efficiently the information revealed through such decentralised bidding. The point here is that the decentralised market aggregates information in a way that no-one is able to do individually. Indeed, there is a growing body of literature which shows that decentralised exchange markets are very efficient in providing forecasts of the probability, the mean and median outcomes, and the correlations among a range of future events. Such markets have been used very successfully to predict uncertain outcomes ranging from the box office prospects of Hollywood movies, through vote shares in elections, to the sales of Hewlett-Packard printers (e.g. Wolfers and Zitzewitz, 2004; Snowberg et al., 2005; Gruca and Berg, 2007).

The structure of the paper is as follows. Section 2 gives relevant background information relating to the betting markets analysed in our study. Section 3 describes the data drawn from bookmaker and betting exchange markets. Section 4 outlines the methodology employed. Our results are presented in Section 5, with discussion. Section 6 concludes.

2. Bookmakers and betting exchanges

A study by Levitt (2004) evaluates the relative assessments of bookmakers and bettors with reference to data from a handicapping competition based on US National Football League matches. Levitt characterises the difference between conventional financial asset markets

and betting markets as follows: in the former the complexity of information affecting the value of assets is such that market makers cannot gain an advantage through superior processing of information to the market as a whole. In contrast, Levitt claims, market makers in betting markets (bookmakers) possess skills in assessing the true chance of various outcomes superior to most bettors, and at least as good as the subset of most skilful bettors. He suggests that the structural consequences of this differential degree of sophistication are that spot markets equalising supply and demand prevail in conventional financial assets markets, with market makers earning the bid-ask spread, whereas profit maximising bookmakers set prices to exploit bettor biases, constrained only by the presence of the smaller number of unbiased bettors. Bookmakers therefore earn the equivalent of a bid-ask spread (known as over-round) and an additional return accruing from their exploitation of bettor biases. One consequence of this tendency of bookmakers to act as price makers is that individual books will expose them to positive risk, as bookmakers assume long and short positions exploiting bettor biases.

A disadvantage of the Levitt approach is that, for his data, bookmakers set the terms of the transaction, and bettors respond with a simple decision whether to bet or not. The most skilful players in this situation may be exercising their talents most effectively in cases where they leave specific games alone, but these decisions are not measured in the Levitt study. A more comprehensive test of the relative sophistication of bookmakers and bettors in assessing the true chances of a range of outcomes would permit bettors to express alternative prices to bookmakers so that we can observe the distribution of revealed preferences of both groups.

We are fortunate that this experiment can now be observed to occur spontaneously over many events in a set of parallel betting markets that has developed in the UK in recent years. The first of these markets is the competitive array of bookmaker fixed odds for specific races available to bettors on the internet. The second is to be found in the person-to-person markets, or betting exchanges, which have revolutionised the betting industry in the UK in recent years (Jones et al., 2006).

Betting exchanges exist to match people who want to bet on a future outcome at a given price with others who are willing to offer that price. The person who bets on the event happening at a given price is the backer. The person who offers the price to an identified sum of money is known as the layer of the bet. The advantage of this form of wagering to the bettor is that, by allowing anyone with access to a betting exchange to offer or lay odds, it serves to reduce margins in the odds compared to the best prices on offer with traditional

bookmakers. Exchanges allow clients to act as backers or layers at will, and indeed to back and lay the same event at different times during the course of the market.

The major betting exchanges present clients with the three best odds and stakes for which other members of the exchange are asking or offering. For example, for a horse named Take The Stand to win the Grand National, the best odds on offer might be 14 to 1 to a maximum stake of £80, 13.5 to 1 to a further stake of £100 and 12 to 1 to a further stake of £500. These odds and staking levels may have been offered by one or more other clients who believe that the true odds are longer than they have offered.

An alternative option available to potential backers is to enter the odds at which they would be willing to place a bet together with the stake they are willing to wager at that odds level. This request (say £50 at 15 to 1) will then be shown on the request side of the exchange and may be accommodated by a layer at any time until the event begins. Every runner in the race will similarly have prices offered, prices requested, and explicit bet limits.

The margin between the best odds on offer and the best odds sought tends to narrow as the volume of bets increases so that in popular markets the real margin against the backer (or layer) tends towards the commission levied on winning bets by the exchange. This commission varies up to 5 per cent, depending on the amount of business the client does with the exchange. Clients can monitor price changes, which are frequent, on the Internet website pages of the betting exchange, and execute bets, lay bets, or request a price instantly and interactively.

Bookmakers have also innovated to take advantage of the Internet, and for many races they offer prices competitively, usually for all runners. Bettors can access the array of prices for runners in matrices displayed on sites such as the *Racing Post* or *Oddschecker*. As with the exchanges, bettors place bets instantly and interactively. Unlike the betting exchanges, however, bet limits are generally not stated, and clients cannot lay or request prices.

In this study we used matched odds data from bookmaker and betting exchange markets for 263 UK horse races in order to measure empirically the accuracy of probability assessments implicit in the prices simultaneously offered by bookmakers and the leading exchange, *Betfair*. The bookmaker data are traditional fixed prices¹, whereas betting exchanges, whose clients are generally non-bookmakers², offer a parallel set of fixed odds,

¹ Prices or “odds” are fixed in the sense that once agreed the odds specified in the transaction are locked in; the terms of subsequent transactions are subject to market fluctuations.

² Evidence given by *Betfair* to a UK Parliamentary committee indicates that only 0.71% of its accounts belonged to active customers who made more than £15,000 from trading in the previous year (Joint Committee on the Draft Gambling Bill, 2004). Even so, there is anecdotal evidence that bookmakers set up accounts with

enabling an assessment of which set of prices has the greatest predictive value. In this way it is possible to compare directly the relative evaluative skills of bookmakers and bettors in assessing the outcome of races.

Our study was completed in three stages, in which we utilised conditional logistic regression, a maximum likelihood estimation (MLE) technique. At stage 1 this method was applied to the datasets in aggregate, with odds probabilities normalised to give a rational probability distribution. Non-runners were removed for the purpose of our analysis. At this stage we expected betting exchange prices to be more accurate than bookmaker data as predictors of race results because, whereas both bookmaker and betting exchange markets hold a structural bias known as the *favourite-longshot bias* (whereby low probability runners or longshots are overbet, and high probability runners are relatively underbet), this phenomenon is more extreme in bookmaker markets than in the betting exchanges (Smith et al., 2006).

The favourite-longshot bias is a notable empirical finding of many studies analysing pari-mutuel betting markets (those in which the odds are strictly proportional to the amounts bet on the competing horses; see Snyder 1978 for a major review of the bias in such markets) and bookmaker markets (e.g. Dowie 1976; Crafts 1985; Shin 1991; Smith et al., 2005, 2006; Paton et al., 2008). Early studies explained the bias by reference to demand-side factors relating to bettor rationality and sophistication; notably, “risk loving” behaviour by bettors was cited.

In recent years a number of theoretical and empirical studies of the bias have suggested supply-side explanations, stressing the behavioural characteristics of market makers (e.g. Shin 1991, 1992, 1993), or structural characteristics of the markets themselves, notably Hurley and McDonough (1995), Vaughan Williams and Paton (1997) and Sobel and Raines (2003).

Hurley and McDonough present an information based model of the favourite-longshot bias. They suggest that odds bias is positively related both to the transaction costs faced by bettors as a class in acquiring information concerning the true probabilities of runners and to the magnitude of the ‘take’ or deductions (i.e. the profit margin or administrative costs of market operators). A further implication is that bias is positively related to the proportion of turnover attributable to casual or uninformed bettors in the market.

the exchanges to help manage their liabilities. We took steps to mitigate against the consequent feedback between the two odds sets by including stage 3 of our analysis (outlined in the narrative).

Supporting the Hurley and McDonough hypothesis, Vaughan Williams and Paton (1997) find that the favourite-longshot bias is more pronounced in low-grade races than in high class races. This finding is consistent with a reasonable assumption that the cost of acquiring information relevant to race outcomes is higher for low grade races than high class contests because there is likely to be less public and media scrutiny of low grade runners.

Sobel and Raines (2003) offer further supporting evidence for an information-based explanation, identifying a lower bias in high volume betting markets, assumed to be better informed, than low volume markets, assumed to be proportionately more heavily populated by casual bettors, consistent with Hurley and McDonough. Smith et al. (2006) similarly find bias to be positively related to market operators' deductions and/or profit margin, and positively related to information search costs.

Shin instead explains the bias observed in bookmaking markets as the consequence of bookmakers' response to asymmetric information where some bettors have privileged information concerning the true probability of one or more horses winning a race by virtue of their insider status. Shin's subsequent theoretical model of the bias as an optimal risk minimising response of bookmakers to the above situation is developed as the solution to a strategic game whereby they set equilibrium prices that maximise profits whilst mitigating losses arising from the presence of insiders.

Shin (1993) adapts the theoretical model in order to estimate "Shin's z ", an empirical estimate of insider activity in the markets associated with specific samples of races. He also shows that z is a proxy measure for the favourite-longshot bias. Shin's estimate of z , the proportion of betting turnover attributable to insider trading, is 2.46% for his sample. Comparable values have subsequently been found for much larger samples of UK races using bookmaker *SP* data. Vaughan Williams and Paton (1997) and Law and Peel (2002), for example, estimate values of 2.03% (481 races) and 2.7% (971 races) respectively.

An independent test of the validity of z as a measure of bias is found in Cain et al. (2001a), who employ both Shin's methodology and a separate method initially adopted by Crafts (1985) to measure the degree of insider trading in relation to a sample of 1,568 horse races and 936 greyhound races. They find the two measures to be significantly associated in a χ^2 contingency test, with a significant positive relationship between them subsequently confirmed in regression analysis.

Shin argues that the bias evident in bookmaker prices is their response to asymmetric information and adverse selection due to the presence of insiders rather than a fundamental

inability of bookmakers to evaluate true probabilities. Bookmaker prices are those that maximise profits rather than representing the bookmakers' estimate of true probabilities. Therefore a fairer test of the relative skills of odds makers would be, first, to adjust odds for the favourite-longshot bias before making a direct comparison. However, Shin is primarily concerned with measuring an aggregate value of z for a sample of races, whereas we are interested here in calculating values of Shin's z for individual races and thereby deriving adjusted probabilities for each runner in a race. We used a method attributable to Jullien and Salanié (1994), later modified by Cain et al. (2001b), to compute individual probabilities adjusted for bias. The Jullien and Salanié model specification is described in section 4.

At stage 2 of the analysis we therefore derived Shin probabilities from the nominal bookmaker and exchange odds before again applying the MLE model described above in relation to stage 1 of our analysis.

Our datasets comprised prices set very early in the market so as to minimise the risks of feedback between the two markets. However we could not eliminate the possibility of prices having already converged somewhat through bettors using the two markets as benchmarks in arbitrage processes and bookmakers operating within the exchanges to hedge their liabilities on specific horses. Given this possibility, at stage 3 the MLE procedure was repeated for the subset of horses for which the divergence between bookmaker and exchange odds probabilities was greatest. The logic dictating the choice of this class of horses for special attention was that, if bookmakers are responsible for offering the exchange odds in such cases, the same subjective probabilities should be evident in their own odds. These horses are arguably subject to the least feedback between the two markets and permit us legitimately to attribute the exchange odds to bettors.

3. Data

To facilitate the study we required two sets of odds for the same races: one set attributable to bookmakers and the other to bettors. The first set of prices collected were fixed odds offered by bookmakers. Unlike pari-mutuel prices, once accepted these odds do not vary with subsequent fluctuations in the market. The only exceptions are when there are withdrawals of runners in the race, in which case a differential reduction is applied based on the probability of success of the withdrawn runner or runners implied in the odds.

Bookmakers' prices were gathered for 700 horse races run in the UK during 2002. Sample races were drawn from the 2001-02 National Hunt season, the 2002 Flat season, and

the 2002-03 National Hunt season. In order to minimise liquidity issues, sampling was restricted to Saturdays and other days where overall betting turnover was likely to be vigorous. One advantage of sampling over the full calendar year is that our data should not suffer in aggregate from seasonal bias. Prices were taken from the Internet site of the *Racing Post*, the major daily publication dealing with horse racing and betting in the U.K. Taking prices from the Internet site allowed a direct comparison with betting exchange data to be made.

To ensure that bookmaker prices were not merely nominal, a trial was conducted whereby bets were placed to establish that the prices stated could be obtained. Actual bets were small (ranging from £5 to £20), but enquiries were also made with individual bookmakers as to whether much larger bets would be accepted. There was evidence of some limits to bet size set by bookmakers on occasions, but not frequently enough to raise concerns about the integrity of bookmaker prices in general.

We were anxious to avoid including low liquidity races in our final race sample as odds are less likely to be representative of informed opinion in such markets. In addition to restricting the sample to the most active race days, two further methods were employed for excluding low liquidity races. The first involved the use of a set of qualitative decision rules for categorising races according to their liquidity initially employed in Smith et al. (2006); the decision rules are reproduced in Table 1. Class 1 races are the least liquid, with Class 4 races having the greatest liquidity. The rules described in Table 1 provided a sound proxy for the volume of betting turnover at the time of data collection³. Qualitative descriptors were necessary as, although trading volumes and bet limits are explicitly stated for all races on the *Betfair* website, they are not normally published by bookmakers; this procedure avoided sample bias arising from the exclusive use of exchange data to determine betting volumes.

The bookmaker odds data were matched with corresponding betting exchange prices, both acquired at 10.30 a.m. on the morning of the races. This time was chosen as it allowed sufficient time for markets to achieve acceptable levels of liquidity. However there are instances where key runners are withdrawn from a race shortly before 10.30, necessitating reformation of the market. When this occurs there is insufficient time for the betting exchange markets to regain the liquidity levels commensurate with the Class of race, and therefore the decision rules in Table 1 would lead to a classification that did not accurately represent trading volume. A second, quantitative rule was therefore applied to mitigate this

³ Confirmed by cross-referencing to press reports on betting volume, where available. We also wish to thank Lawrence McDonough for constructive advice on the formulation of the decision rules.

problem in the small number of races affected. Where bet limits were small, the prices offered were ignored, and races where overall betting volume was trivially low were excluded from the sample of races on the grounds that the market did not have sufficient liquidity to warrant treating such observations as representative. A minimum acceptable aggregate turnover threshold (£2000 per race, by 10.30 am) was applied as a filter to the races in the sample in respect of *Betfair* prices; races where this turnover threshold was not met were screened out of the analysis. After exclusion of races on grounds of low turnover, 267 races remained for analysis, four of which were subsequently abandoned due to adverse weather conditions, leaving a final study sample of 263 races.

4. Empirical models employed

Our test of predictive accuracy in relation to the two markets involved conditional logistic regression, a maximum likelihood technique giving log likelihood calculations which can be judged against the χ^2 distribution. It is a technique that is frequently employed in empirical studies of predictive models in horse race wagering (for good examples, see Figlewski, 1979; Bolton and Chapman, 1986; and more recently Sung et al., 2005).

For an individual race j with n runners, the conditional logistic regression model can be expressed as

$$p_i = \frac{e^{Z_i\beta}}{\sum_{j=1}^N e^{Z_j\beta}} \quad (2)$$

where p_i is the conditional probability of horse i winning race j ; Z is the vector of predictors, in this case the subjective probabilities implied by the odds pertaining to horse i (numerator), and all race entrants, $i = 1 \dots N$ (denominator); and β is the vector of coefficients attached to the predictors.

As a maximum likelihood technique, the estimated coefficients β are those which maximise the likelihood function:

$$L = \prod_{k=1}^M P_w^k \quad (3)$$

where P_w^k is the estimated probability associated with the horse winning the k^{th} race in a sample of M races (using the notation employed by Figlewski, 1979).

L is initially calculated with restrictions on the variables Z (i.e. the coefficients are all set to zero). In this initial calculation $p_i = 1/N$ for all runners in the race, representing the

situation where there is no information about the race entrants. After iterative estimation of the coefficients, L is re-calculated by dropping the restriction(s) on one or more of the predictor variables Z . The interpretation of the change in L is that the closer its value is to zero, the more closely are the race outcomes explained by the information held in the predictors, in this case the odds.

McFadden (1973) showed that the value $2(Lu - Lr)$, where Lu is the likelihood function calculated with unrestricted predictors and Lr is that with restrictions, closely follows the χ^2 distribution. This test therefore permits us to judge whether a set of odds holds significant information about the outcome of races in our sample.

The vector of predictors, Z , in our sample comprises probabilities corresponding to the odds in our bookmaker and betting exchange data. Bookmaker prices in the UK are expressed as fractional odds values (e.g. $2/1 = 0.33$ odds probability; $4/1 = 0.2$, and so on). The recorded value for each race entrant i was based on the mean of observed bookmakers' odds for that horse expressed as a probability. The mean of the odds array was chosen rather than the outlier on the basis that the former better represents the consensus of bookmaker opinion.⁴ The betting exchange odds adopted were the maximum available to significant bet limits. They are expressed on the website as the return inclusive of one unit stake in decimal rather than fractional format. Increases in odds at the high probability end of the odds scale are expressed in "ticks" of 0.1 point. For low probability runners, prices increase in 1 point increments.

At stage 1 of the analysis, we normalised odds probabilities to sum to unity for each race in the respective markets, removing withdrawn runners, so that

$$p_{ij}^n = \frac{p_{ij}^o}{\sum p_{ij}^o} \quad (4)$$

where p_{ij}^n is an estimate of the race specific true probability of horse i winning race j .

This method of proportional normalisation, commonly adopted in earlier studies (for example, Tuckwell, 1983; Bird and McCrae, 1987), ensures that the odds pertaining to each race observe a rational probability distribution, but does not remove the effects of the favourite-longshot bias as all odds probabilities are adjusted in the same proportion.

⁴ We repeated our three stage analysis using the outlier instead of the mean and found there to be very little difference in the outcome.

The initial estimation of the likelihood function (3) involves restricting to zero the coefficients of all predictors, from both bookmaker and betting exchange normalised odds probabilities. The resulting log likelihood, Lr , represents the model's best fit with no prior information. Further estimations of the likelihood function were then carried out as in Table 2. The statistic Lm , if significantly lower than the baseline value, would provide strong evidence that mean bookmaker prices hold valuable information about the outcome of races. We expected this to be the case, as bookmaker favourites won 25.48% of the races in the sample, representing a large improvement on the expected success rate of a randomly chosen runner in each race, 6.84% (263 winners divided by total number of runners in the sampled races, 3,843). Similarly, we expected the betting exchange data alone to add valuable information, evidenced by Lb .

Lm , $L(b,m)$ shows how the log likelihood for Lm changes by including the exchange odds in addition to the existing predictor of bookmaker odds. The estimation Lb , $L(b,m)$ reverses the order of predictor additions. These two estimations are key to judging the relative information held by the two odds sets and, hence, their predictive value.

Shin was primarily concerned with measuring an aggregate value of z for a sample of races, whereas for stages 2 and 3 we wished to calculate values of Shin's z for individual races and derive adjusted probabilities for each runner in a race. We used a method attributable to Jullien and Salanié (1994) to compute individual probabilities adjusted for bias.

Jullien and Salanié restated Shin's model to show that, for an individual race, the true probability of winning p_i for horse i can be expressed as

$$p_i = \frac{\sqrt{\left[z^2 + 4 \frac{\pi_i^2}{\Pi} (1-z) \right]} - z}{2(1-z)} \quad (5)$$

where z is Shin's measure of insider trading for that race, π_i is the nominal odds probability associated with horse i , and Π is the sum of π_i in the race. Jullien and Salanié showed that z can be estimated using the following equation:

$$\sum p_i \left(\frac{\pi_i}{\sqrt{\Pi}}, z \right) = 1 \quad (6)$$

Through an iterative procedure similar to that employed by Shin, the observed values of π_i are substituted into equations (5) and (6) to derive race specific values of z that will yield probabilities from equation (5), adjusted for insider trading and which sum to unity.

The idea is that we can derive probabilities for our bookmaker and betting exchange datasets using this technique, which should remove the systematic bias in nominal odds associated with the strategic behaviour modelled by Shin arising from asymmetric behaviour, described in section 1 above. The resulting adjusted probabilities then represent the underlying estimates of true probabilities, attributable to bookmakers and exchange clients respectively, implicit in the odds.

A question arises as to whether it is legitimate to apply the Shin's z measure of bias, which is derived from a model of bookmaker behaviour, to betting exchange data. Shin models bookmaker competition, whereas the exchange consists of individuals who do not need to maintain a credible market structure embracing all runners. Aside from the degree of bias and level of transaction costs, however, the structure of exchange markets resembles those of bookmakers quite closely. For example, for individual runners exchange layers have to offer competitive prices to attract bettors with the sum of price probabilities usually exceeding unity by only a few percent. On the other hand, the occasions when the sum of probabilities falls below one are extremely rare. We felt justified, therefore, in applying the Shin's z measure in this non-bookmaker betting medium.

As the Shin probabilities are estimated independently of results, a useful test of their efficiency in removing bias is to calculate returns for the datasets at odds corresponding to the Shin probabilities themselves. If the Shin adjustments are successful, the distribution of returns arising from these calculations should be equal across different odds values. In order to test the degree to which our adjustments for bias were efficient, we regressed the notional returns to Shin odds equivalents (dependent variable) against odds probabilities corresponding to actual odds (independent variable), with standard errors adjusted for heteroscedasticity, to see if this was the case. A slope coefficient not significantly different from zero would provide evidence that the Shin probabilities were unbiased. Having derived adjusted odds probabilities, we completed stage 2 by repeating the MLE procedures outlined for stage 1, substituting the Shin probabilities for the nominal odds probabilities.

For stage 3 of our analysis we required a suitable measure of price divergence. Cain et al. (2001a) compared alternative means of calibrating odds movements during the course of market trading; although here we were instead interested in odds differentials arising from competing odds values at a point in time, the relative merits of the possible methods of measurement are similar.

Arithmetic differences in odds were ruled out due to the disproportionate importance attached to the odds differentials of outsiders (a one point odds difference between 9/1 and

10/1 does not carry the same significance, in terms of probability or betting turnover, as a similar arithmetic difference between 2/1 and 3/1). Crafts (1985) partially overcame this problem by adopting the ratio of the forecast odds probability to the final (starting price) odds probability. However the trading volume required to change the odds probability from 0.1 to 0.2 (a Crafts ratio of 2) is much greater than that needed to cause a movement from 0.2 to 0.4 (also a ratio of 2) and therefore leads to undue emphasis on longshot observations.

Law and Peel (2002) employed the alternative measure, pm , such that

$$pm = \log\left(\frac{1}{1-p_1}\right) - \log\left(\frac{1}{1-p_2}\right) \quad (7)$$

where, for an individual runner in a race, p_1 and p_2 are the odds probabilities derived from, for example, starting odds and forecast odds respectively. Unlike the Crafts ratio, equation (7) weights price movements from initially low odds with greater emphasis than those from initially high odds, reflecting the greater trading volumes required to cause odds to change at low odds. For similar reasons we adopted equation (7) as our measure of divergence, pd , between bookmaker mean and exchange odds for each horse, with p_1 being the highest odds probability value, and p_2 the lowest odds probability value. For illustrative purposes Table 3 indicates the divergence for different levels of odds required to yield specific values of pd .

To complete stage 3 we then repeated the MLE procedures employed in the previous stages for the subset of horses exhibiting the greatest odds divergence, allowing for possible sensitivity of results to our choice of pd value constituting high divergence.

5. Results and discussion

The estimation identifiers used in this section follow the descriptions in Table 2 for the MLE iterations based on nominal odds probabilities (stage 1). The corresponding identifiers for the stages 2 and 3 iterations based on Shin probabilities differ only in the use of a subscript, a . Thus the term “ ma ” indicates Shin adjusted mean bookmaker odds probabilities; the term “ ba ” similarly indicates Shin adjusted exchange odds probabilities.

Table 4 summarises the stage 1 results of the initial log likelihood estimates for normalised odds probabilities derived from bookmaker mean and exchange odds. The χ^2 test statistics in Table 4 correspond to the various restrictions on predictors outlined in Table 2.

The log likelihood values for bookmaker odds alone (Lm) and exchange odds alone (Lb) indicate that each set of odds individually contributes significant information in

predicting the outcomes of the races in our sample; the χ^2 test statistic for each is significant at $p = 0.01$. Recall that this result was anticipated as each odds set reflects information on the respective merits of runners not contained in the baseline model, which assumes a uniform probability distribution across race competitors.

The further measures Lm , $L(b,m)$ and Lb , $L(b,m)$ permit us to judge whether either of the nominal odds sets holds valuable information in addition to the other set alone. The mean bookmaker odds unadjusted for the favourite-longshot bias yield significant additional information to that contained in the exchange data, with the χ^2 statistic associated with Lb , $L(b,m)$ being significant at $p = 0.05$. When the order of addition is reversed, the exchange data add significantly to the amount of information concerning race outcomes held in the unadjusted mean odds alone, evidenced by a χ^2 value for Lm , $L(b,m)$ significant at $p = 0.01$. Whilst not conclusive, the evidence suggests a marginal superiority of exchange odds as predictors. This is consistent with the Shin's z values for these datasets; z for the bookmaker mean odds is 2.17%, significantly greater than that for the exchange odds, at 0.09% (see Smith et al., 2006, for confirmation of the independence of these Shin's z results for the same datasets).

We then performed stage 2 of the analysis, repeating the conditional logistic regression in relation to odds adjusted for the favourite-longshot bias, which was removed using the Jullien and Salanié method outlined above. If the Shin adjustments are successful, the distribution of returns arising from regressions of returns on odds probabilities should be equal across different odds values and should also approximate zero. Table 5 summarises the estimated coefficients of these regressions for the unadjusted and adjusted bookmaker mean and exchange odds.

The pre-adjustment slope coefficients in Table 5 are consistent with the Shin values reported above. The table shows that before Shin adjustment, the bookmaker odds contain an appreciable bias: the slope coefficient β indicates that for every 1% increase in odds probabilities, returns increase by a highly significant 1.57%. After adjustment the coefficient estimate reduces to 0.7, insignificant at conventional significance levels. The β value for the exchange odds before the Shin adjustment is 0.52. This is not significantly different from zero, implying little initial bias. The adjustment of exchange odds decreases the estimate of β to a value very close to zero and insignificant at any level. We conclude therefore that the Shin adjustments successfully remove bias from the data.

Figures 1 and 2 confirm this outcome, although it is evident that there remains a trend in the direction of regular bias in both sets of odds, more pronounced in the bookmaker data, if statistically insignificant. Because at stage 3 we used directly the adjusted probabilities upon which these regressions were based, we cannot entirely eliminate the possibility that our results at that stage partly reflected inefficiencies in the Shin probabilities.

At stage 2 of the analysis, the initial conditional regressions carried out at stage 1 were repeated with the Shin probabilities. The results are summarised in Table 6. From the test $Lma, L(ba,ma)$, it is apparent that following adjustment for bias there is highly significant evidence that betting exchange odds continue to add further useful information to that contained in the bookmaker odds, with a χ^2 value significant at $p = 0.01$. In contrast, $Lba, L(ba,ma)$ yields a corresponding χ^2 statistic significant only at $p = 0.1$. For the full sample of races with probabilities adjusted for bias, it appears that the betting exchange odds still have greater predictive accuracy than the bookmaker equivalents, and the margin of advantage is, in fact, greater than for nominal odds.

At stage 3 the conditional logistic regressions were repeated for subsets of horses with varying levels of price divergence. In order to avoid an arbitrary choice we performed regressions for alternative pd filter levels, beginning with $pd \geq 0.01$, then $pd \geq 0.02$, and $pd \geq 0.03$.

The outcomes of the stage 3 conditional logistic regressions are summarised in Table 7, organised by pd filter value. For the purpose of comparison the results for the category $pd < 0.01$ are also included. The key values to consider in the current context are those relating to the $Lma, L(ba,ma)$ and $Lba, L(ba,ma)$ regressions. By and large, the results in Table 7 confirm the superior predictive value of the exchange odds found at earlier stages for the subset of horses exhibiting greater odds divergence than the norm. Changes in the log likelihood values support this conclusion, being significant at $p = 0.05$ for $pd \geq 0.01$ and $pd \geq 0.03$. Our main caveat to this inference is that the result may be partially attributable to the inefficient removal of bias by the use of Shin probabilities.

We attribute this greater predictive power to bettor influence, since if bookmakers are responsible for offering exchange odds in these cases, why are these probability estimates not reflected in their own prices, in which case the divergence would not arise. We also rule out pricing error as the cause of high divergence observations as a class because errors would be unsystematic, not yielding significant χ^2 values. Whether such bettor superiority is the result of skill or the possession of inside information is open to question.

Figure 3 restates the results from all three stages of our analysis in terms of changes in log likelihood arising from the key models Lm , $L(b,m)$, our test of exchange odds holding additional information to the bookmaker odds, and Lb , $L(b,m)$, the equivalent test for bookmaker odds yielding information in addition to the exchange odds.

6. Conclusions

The visual evidence in Figure 3 supports the conclusion, drawn from our χ^2 significance tests, that the exchange odds tend to have greater predictive value than the bookmaker equivalents in our sample. Our expectation that nominal betting exchange odds have more predictive value than bookmaker odds due to the lower degree of favourite-longshot bias in the former is confirmed. After adjusting for bias the exchange odds continue to hold more information concerning race outcomes than bookmaker odds, and this result is also broadly confirmed in those instances where price divergence is higher than the norm; although the advantage at $pd \geq 0.02$ disappears and is only marginal at $pd \geq 0.03$, the number of cases in these categories is relatively few. In the main, therefore, the exchange odds prove to be superior predictors of the results of the sampled races.

Our principal finding contrasts with that of Levitt, who found that bookmakers exhibited superior skills in evaluating objective outcomes in the handicapping contest that was the medium for his study. The observation was made in the introduction that Levitt's methodology made it likely that the preferences of the most skilled or informed bettors might not be revealed if they decided that the terms of the wagers set by bookmakers were unfavourable and, in consequence, chose not to trade. The same would be true of the bookmaker markets studied here. In contrast the betting exchanges offer opportunities for these bettors to trade which are not available in bookmaker markets. For example, skilled traders, insiders, and bettors seeking hedging opportunities are all able to lay odds on the exchanges which may as a result more accurately reflect the chances of the horses concerned than those offered by bookmakers. In these circumstances we might expect the proportion of turnover attributable to casual bettors to be lower in the exchanges than in bookmaker markets, with a consequent tendency for odds to reflect more closely objective probabilities. This account of the differences between the two markets is consistent with recent transaction cost explanations of the structure of betting markets, outlined in section 2. It is also consistent with the empirical evidence suggesting that decentralised markets are efficient predictors.

Differences in the nature of traders and trading activities may therefore explain the greater relative efficiency of the exchanges in reflecting objective outcome probabilities observed in the current study. Similarly, the results presented here may not so much contradict Levitt's findings as reflect a different composition of traders engaged in the respective betting media studied.

As the betting exchanges continue to expand in size and liquidity, it will be interesting to monitor how well they continue to predict the outcomes of events for which they offer markets.

Tables

Table 1: Decision Rules for Classifying Horse Race Markets by Liquidity.

<p style="text-align: center;">Class 1: Races with low betting volume.</p> <p>Races run at low grade racetracks and for small monetary prizes; often unexposed or unknown form for a number of runners; minimal media coverage.</p>
<p style="text-align: center;">Class 2: Races with moderate betting volume.</p> <p>Races for middle ability horses; form is more exposed than Class 1 races; average prize money.</p>
<p style="text-align: center;">Class 3: Races with higher than average betting volume.</p> <p>Competitive races with a high degree of betting interest, generated by characteristics of the race or its contenders likely to attract public interest and enhanced media coverage; higher than average reported betting volumes in the press.</p>
<p style="text-align: center;">Class 4: Races with very high betting volume.</p> <p>High profile and top class races; high profile contending horses; high degree of competition and media interest, speculation on runners often extending weeks before the contest.</p>

Note:

The class numbers and definitions here are those employed by the authors, and do not correspond to classes as specified by the UK racing authorities.

Table 2: Estimations of the likelihood function with various degrees of restriction for race outcome predictors (odds)

Estimation of likelihood function	Restrictions on predictors	Alternative hypothesis tested
L_r	All predictors restricted	
L_m	Only mean bookmaker odds unrestricted	Bookmaker odds hold useful information concerning race outcomes
L_b	Only betting exchange odds unrestricted	Exchange odds hold useful information concerning race outcomes
$L_m, L(b,m)$	Both predictors unrestricted (relative to log likelihood of bookmaker odds alone)	Exchange odds hold useful information <i>additional</i> to that contained in bookmaker odds
$L_b, L(b,m)$	Both predictors unrestricted (relative to log likelihood of exchange odds alone)	Bookmaker odds hold useful information <i>additional</i> to that contained in exchange odds

Note:

The test statistic in all cases is χ^2 with 1 degree of freedom

Table 3: Odds differences corresponding to increasing levels of price divergence, pd

Higher odds	Lower odds, yielding specified pd values				
	$pd = 0.01$	$pd = 0.02$	$pd = 0.03$	$pd = 0.04$	$pd = 0.05$
2	1.94	1.89	1.83	1.78	1.73
5	4.72	4.46	4.23	4.02	3.82
10	9.00	8.18	7.49	6.90	6.39
20	16.51	14.04	12.20	10.77	9.63

Notes:

- (i) pd is a measure of the divergence between the odds probabilities equivalent to the mean of bookmaker array of odds for an individual horse and the greatest *Betfair* (betting exchange) odds on offer to non trivial stakes for the corresponding horse.
- (ii) pd is measured as in equation (7); see also accompanying narrative.
- (iii) All odds expressed to a 1 unit stake e.g. “2-1”, “1.94 to 1” and so on.

**Table 4: Conditional logistic regression results for whole dataset:
nominal odds probabilities**

Model restrictions	- 2 log likelihood
Lr	1368.309
Lm	1247.120*** (121.189)
Lb	1239.801*** (128.508)
$Lm, L(b,m)$	1235.272*** (11.848)
$Lb, L(b,m)$	1235.272** (4.529)
N	3843

Notes:

(i) ***p=0.01, **p=0.05, *p=0.1

(ii) Figures in parentheses are the relevant χ^2 statistics.

Table 5: Coefficients of returns regressed on odds probabilities, bookmaker mean & exchange odds, unadjusted & Shin adjusted

Coefficients	Bookmaker mean odds		Exchange odds	
	Unadjusted	Shin adjusted	Unadjusted	Shin adjusted
c	-0.2756*** (0.0581)	0.0037 ((0.0898)	-0.0916 (0.0845)	-0.0027 (0.0980)
α	-0.3821*** (0.0873)	-0.0434 (0.1421)	-0.1268 (0.1327)	-0.0084 (0.1568)
β	1.575*** (0.5903)	0.6971 (0.9212)	0.5211 (0.8346)	0.0843 (0.9850)
R^2	0.0008	0.0001	0.0000	0.0000
N	3843			

Notes:

(i) *** p = 0.01 ** p = 0.05 * p = 0.1.

(ii) Figures in parentheses are robust standard errors.

(iii) c is the coefficient of returns (measured as net profit to a unit stake to the specified odds) regressed on a constant.

(iv) Shin adjusted returns are those to notional odds corresponding to Shin probabilities.

(v) α and β are the constant term and slope coefficients respectively of returns regressed on probabilities corresponding to the specified odds.

**Table 6: Conditional logistic regression results for whole dataset:
Shin probabilities**

Model restrictions	- 2 log likelihood
<i>Lr</i>	1368.309
<i>Lma</i>	1246.456*** (121.852)
<i>Lba</i>	1239.962*** (128.347)
<i>Lma, L(ba,ma)</i>	1236.161*** (10.296)
<i>Lba, L(ba,ma)</i>	1236.161* (3.802)
N	3843

See notes to Table 4.

**Table 7: Conditional logistic regression results for subset of horses with
greatest price divergence between mean bookmaker & exchange odds
(based on Shin probabilities)**

Model restrictions	<i>pd</i> < 0.01	<i>pd</i> ≥ 0.01	<i>pd</i> ≥ 0.02	<i>pd</i> ≥ 0.03
<i>Lr</i>	737.398	302.964	42.296	5.545
<i>Lma</i>	682.924*** (54.474)	263.545*** (39.419)	36.033** (6.262)	5.493 (0.052)
<i>Lba</i>	681.985*** (55.503)	260.663*** (42.301)	35.642*** (6.653)	5.400 (0.145)
<i>Lma, L(ba,ma)</i>	681.485 (1.438)	259.667** (3.878)	35.594 (0.440)	1.755** (3.739)
<i>Lba, L(ba,ma)</i>	681.485 (0.410)	259.667 (0.996)	35.594 (0.490)	1.755* (3.645)
N	2852	991	207	63

Notes:

- (i) This subset includes horses for which $ma > ba$ and those where $ba > ma$
- (ii) See also notes to Table 4.

Figure 1. Fitted values derived from regression of returns on bookmaker odds probabilities

Figure 2. Fitted values derived from regression of returns on *Betfair* odds probabilities

Figure 3. Change in McFadden statistic (-2 log likelihood): Stages 1 (Unadjusted), Stage 2 (Adjusted), & Stage 3 (High divergence)

Note:

$L_m, L(b,m)$ is the test of exchange odds holding additional information to the bookmaker odds;
 $L_b, L(b,m)$ is the equivalent test for bookmaker odds yielding information in addition to the exchange odds.

References

- Bird, R., McCrae, M., 1987. Tests of the efficiency of racetrack betting using bookmaker odds. *Management Science* 33, 1552-1562.
- Bolton, R.N., Chapman, R.G., 1986. Searching for positive returns at the track: A multinomial logit model for handicapping horse races. *Management Science* 32, 1040-1060.
- Cain, M., Law, D., Peel D.A., 2001a. The relationship between two indicators of insider trading in British racetrack betting. *Economica* 68, 97-104.
- Cain, M., Law, D., Peel D. A., 2001b. The incidence of insider trading in betting markets and the Gabriel and Marsden anomaly. *The Manchester School* 69, 197-207.
- Crafts, N.F.R., 1985. Some evidence of insider knowledge in horse race betting in Britain. *Economica* 52, 295-304.
- Dowie, J. 1976. On the efficiency and equity of betting markets. *Economica* 43, 139-50.
- Figlewski, S., 1979. Subjective information and market efficiency in a betting market. *Journal of Political Economy* 87, 75-88.
- Gruca, T.S., Berg, J.E. 2007. Public information bias and prediction market accuracy. *Journal of Prediction Markets* 1, 219-231.
- Hurley, W., McDonough, L., 1995. A note on the Hayek hypothesis and the favorite-longshot bias in parimutuel betting. *American Economic Review* 85, 949-955.
- Joint Committee on the Draft Gambling Bill, 2004. House of Lords, House of Commons, HL Paper 63-I, HC 139-I, Session 2003-2004, p.139. London: The Stationery Office Limited.
- Jones, P., Turner, D., Hillier, D., Comfort, D., 2006. New business models and the regulatory state: A retail case study of betting exchanges. *Innovative Marketing* 3, 112-119.
- Jullien, B., Salanié, B., 1994. Measuring the incidence of insider trading: A comment on Shin. *Economic Journal* 104, 1418-1419.
- Klemperer, P., 1999. Auction theory: A guide to the literature. *Journal of Economic Surveys* 13, 227-286.
- Klemperer, P., 2004. *Auctions: Theory and Practice*. Princeton, NJ: Princeton University Press.
- Law, D., Peel, D.A., 2002. Insider trading, herding behaviour, and market plungers in the British horse race betting market. *Economica* 69, 327-338.
- Levitt, S., 2004. Why are gambling markets organised so differently from financial markets? *Economic Journal* 114, 223-246.

- McFadden, D., 1973. Conditional logit analysis of qualitative choice behavior. In: Zarembka, P. (Ed.). *Frontiers in Econometrics*. New York: Academic Press, 303-328.
- Paton, D., Siegel, D. and Vaughan Williams, L., 2008. 'The growth of gambling and prediction markets: Economic and financial implications, *Economica*, doi:10.1111/j.1468-0335.2008.00753.x, 1-6.
- Shin, H.S., 1991. Optimal betting odds against insider traders. *Economic Journal* 101, 1179-1185.
- Shin, H.S., 1992. Prices of state contingent claims with insider traders, and the favourite-longshot bias. *Economic Journal* 102, 426-435.
- Shin, H.S., 1993. Measuring the incidence of insider trading in a market for state-contingent claims. *Economic Journal* 103, 1141-1153.
- Smith, M.A., Paton, D., Vaughan Williams, L., 2005. An assessment of quasi-arbitrage opportunities in two fixed-odds horse-race betting markets. In: Vaughan Williams, L. (Ed). *Information Efficiency in Financial and Betting Markets*. Cambridge: Cambridge University Press, 159-171.
- Smith, M.A., Paton, D., Vaughan Williams, L., 2006. Market efficiency in person-to-person betting. *Economica* 73, 673-689.
- Snowberg, E., Wolfers, J. and Zitzewitz, E., 2005. Information (in)efficiency in prediction Markets. In: Vaughan Williams, L. (Ed.). *Information Efficiency in Financial and Betting Markets*. Cambridge: Cambridge University Press, 366-86.
- Snyder, W.W., 1978. Horse racing: Testing the efficient markets model. *Journal of Finance* 33, 1109-18.
- Sobel, R.S., Raines, S.T., 2003. An examination of the empirical derivatives of the favorite-longshot bias in racetrack betting. *Applied Economics* 35, 371-85.
- Sung, M., Johnson, J.E.V., Bruce, A.C., 2005. Searching for semi-strong inefficiency in the UK racetrack betting market. In: Vaughan Williams, L.(Ed.). *Information Efficiency in Financial and Betting Markets*. Cambridge: Cambridge University Press, 179-192.
- Tuckwell, R., 1983. The thoroughbred gambling market: Efficiency, equity and related issues. *Australian Economic Papers* 22, 106-108.
- Vaughan Williams, L., Paton, D., 1997. Why is there a favourite-longshot bias in British racetrack betting markets? *Economic Journal* 107, 150-158.
- Wolfers, J., Zitzewitz, E. 2004. Prediction markets. *Journal of Economic Perspectives* 18 (2), 107-26.