‘Quarbs’ and Efficiency in Spread Betting: can you beat the book?

January 2002

David Paton*

Nottingham University Business School
Jubilee Campus
Nottingham NG8 1BB
UK
Tel: +00 44 115 846 6601
David.Paton@nottingham.ac.uk

and

Leighton Vaughan Williams

Department of Economics and Politics
Nottingham Trent University
Burton Street
Nottingham NG1 4BU
UK
Tel: 0115 848 5516
Leighton.Vaughan-Williams@ntu.ac.uk

*Corresponding author

Word Count: 7254 (assuming Tables equivalent to 1000 words)

‘Quarbs’ and Efficiency in Spread Betting: can you beat the book?

Abstract

In this paper, we examine a relatively novel form of gambling, index (or spread) betting, that mirrors (and indeed overlaps with) practices in conventional financial markets. In this form of betting, a number of bookmakers quote a bid-offer spread about the result of some future event, and bettors are invited to buy (sell) at the top (bottom) end of the quoted spreads. We hypothesise that the existence of an outlying spread may provide uninformed traders with information that can be used to develop improved trading strategies. Using conditional moment tests on data from a popular spread betting market in the United Kingdom, we find that in the presence of a number of price-setters, the market mid-point is indeed a better predictor of asset values than the outlying price. We further show that this information can be used to develop trading strategies that lead to returns that are consistently positive and superior to those from noise trading and, in some cases, significantly so.

Keywords: Quarbs, market efficiency, betting.

JEL Classification: D82, G12, G14.

'Quarbs' and Efficiency in Spread Betting: can you beat the book[1]

1. Introduction

For a number of years, researchers have studied market efficiency in betting markets, arguing that such markets are good proxies for financial markets. Notably, these markets possess not only many of the usual attributes of financial markets, specifically a large number of investors (bettors) with potential access to widely available information sets, but also the property that each asset (or bet) has a well-defined termination point, characterised by a definite value. This contrasts, for example, with financial securities involving options, where the value of an asset in the present is dependent both on the present value of future cash flows and also on the uncertain price at which it can be sold at some future point. Moreover, by enabling a more productive and clearer learning process, a delineated end-point might be expected to promote information efficiency. Evidence of inefficiency in betting markets may therefore be of special significance.

There is, indeed, an array of evidence to support the contention that fixed-odds and parimutuel betting markets may be subject to systematic biases.[2] Most notable amongst these biases are the favourite-longshot bias, which is the observed tendency for the expected return to bets placed at lower odds to exceed that to bets placed at higher odds, and the 'hot hand' effect. The hot hand effect is a tendency by bettors to overestimate the extent to which a team or individual's performance is positively autocorrelated. However, such work may be unrepresentative of general financial markets in two key respects. Firstly, the satisfaction gained from making the wager itself and/or jointly consuming the associated event is relatively more significant for the majority of bettors than for financial traders, whose utility functions are more likely to be dominated by wealth and risk considerations. Secondly, a typical wager (either on a lottery or a horse race) involves the bettor risking a fixed, small proportion of their wealth. In contrast, many financial decisions involve risking a more variable and potentially much larger proportion of wealth.

---

[1] The authors would like to thank participants at the 2nd Equine Industry Academic Conference, at Louisville, Kentucky, the 3rd International Conference on Money Investment and Risk at Nottingham, and at research seminars at Nottingham University Business School and Nottingham Trent University. We would particularly like to thank Donald Siegel, Paul Fenn, Kevin Dowd, Ian Walker and Steve Heasell for useful comments and suggestions that have helped to clarify the ideas in this paper.

[2] For examples of this literature, see Thaler and Ziemba (1988), Camerer (1989), Brown and Sauer (1993), Schnytzer and Shilony (1995), Vaughan Williams and Paton (1997), Golec and Tamarkin (1998), Cain, Law and Peel (2000). An extensive review is given in Sauer (1998).

In this paper, we consider a relatively novel form of wagering, namely 'Index (or Spread) Betting', which much more closely resembles the operation of conventional financial markets and, indeed, is in some cases indistinguishable from them. In this sense, we are echoing the work of Avery and Chevalier (1999), who look at sports betting from a perspective which has a natural analogue in the stock market.

Spread (Index) Betting originated in the United Kingdom in the mid-1970s, but developed rapidly in the late 1980s and the 1990s. It is quite different to point spread betting, as operated in the US, which is essentially a fixed-odds 'handicap' betting system, in which bettors wager at fixed odds on one team to beat the other after points are artificially deducted from one of the teams. In Spread (Index) Betting, bettors instead buy or sell notional assets associated with an event (for example, points in a football game), based on a bid-offer spread set by traders (bookmakers). The bid-offer spread may either increase or decrease until the value of the asset is known with certainty (at the end of the game). At this point a bettor who bought the asset will win or lose the difference between the ex-post value of the asset and the ask price multiplied by their original stake. The bettor may 'close' the trade at any time up to the terminal point based on the current bid-offer spread, taking either a loss or a win as appropriate.

The size of the bid-offer spread involves a trade-off between setting a large spread, so as to minimise the profit of insider traders (defined as those who possess superior information to the market-maker) and setting the optimal spread against noise or liquidity traders. Papers by Copeland and Galai (1983), Glosten and Milgrom (1985) and Kyle (1985) have provided formal analyses of this issue and this methodology has been adapted for use in fixed-odds betting markets by, for example, Shin (1993). The findings of this body of work reveal that market prices are indeed set on the basis that there is a significant incidence of insider activity. In any case, in an efficient market it should not be possible at quoted prices to make systematic abnormal returns, based on publicly available information.

The scope of 'assets' that are traded in spread betting markets is wide, ranging from the price of gold to the number of goals in a football match. Indeed, the 'asset' that is the subject of the trade can include almost any clearly quantifiable feature of an array of sports, political and financial markets. There are also some so-called 'speciality' indices, for example the number of 'Oscars' won by a given film at the Academy Awards ceremony. Standard political trades include the number of seats gained by a political party in an election. Financial trading is exceptionally diverse, and includes the values of the Wall Street, DAX, FTSE, Hang Seng and Nikkei share indices, the price of individual shares traded on these markets, the price of a variety of commodities, as well as bond and currency futures.[3]

Spread betting markets are characterised by a low level of transaction costs (at least relative to traditional betting markets). Therefore, they are an attractive option both to small traders, motivated primarily by wealth considerations, and to larger traders using financial spread betting markets as part of a more general risk management strategy. In particular, spread trading is often used to hedge against, for example, a potential short-term fall in the market. The low transactions costs also make it possible for potential arbitrageurs to profit from relatively small mispricings in the market (see, for example, Hurley and McDonough, 1995). Jackson (1994) and Haigh (1999) have also provided formal probabilistic treatments of such trading markets.

The main purpose of this paper is to examine efficient pricing within spread betting markets, notably whether it is possible for rational 'arbitrageurs' to exploit mispricing available in the quoted odds across the market-place. In this sense, we are seeking to add to the literature contributed by several authors who have already shown conditions in more traditional markets where 'arbitrageurs' fail to eliminate mispricing.[4] In particular, we consider the setting of a price by one company that is everywhere outside the mid-point of the spreads of all market-makers combined. We term such a price a 'quasi-arbitrage' or 'Quarb' and ask whether it can provide useful information to bettors.[5]

Most previous work on mispricing in sports betting markets has focused on the final line (see Gandar et al., 1988 for a number of references; see also Woodland and Woodland, 1994, Gray and Gray, 1997). In this paper, we consider the opening line explicitly and critically, as well as the line once prices have settled. In this way we add to work by Gandar et al. (1988, 1998) and Avery and Chevalier (1999), albeit in a different betting arena. We also extend the literature away from the focus on the favourite-longshot bias and the hot-hand bias.

Using data from the popular spread market in bookings points (corresponding to an index of fair play) for football matches in the UK, we employ a variety of tests to distinguish between alternative predictors of asset prices. We then investigate whether the information from these tests can be used to develop profitable (or at least improved) trading strategies.

---

[3] Paton, Vaughan Williams and Siegel (2002) provides further discussion of the distinction between financial and sports spread betting.

[4] See De Long et al. (1990, 1991), Shleifer and Vishny (1997), Avery and Chevalier (1999) for examples of this work in a different context.

In the next section of the paper we describe the operation of spread betting markets in some detail. In section three, we outline our empirical approach and introduce our data. In section four, we report our results, whilst in section five we make some concluding remarks.

2. Spread Betting Markets

The central feature of a spread betting market is the setting by a market-maker (the bookmaker) of a bid-offer spread for the commodity in question. For example, in a cricket game between England and Australia, the bookmaker might set spread for runs in England's first innings of 240-250. A bettor who believes England's batting is weak may sell total runs at the price of 240 on a stake of, say, £5 per run. If England score 215 runs, this is the termination value of the asset and the bettor will win £125, calculated as the difference between the value and the price (240 - 215) times the stake (£5). On the other hand, if England score 290 runs in the game, the bettor would lose £250, calculated as the difference between 290 and 240 times the £5 stake. Similarly, the bettor could choose to buy total runs at a price of 250. In this case, an England score of 215 would result in a loss of £175 whilst a score of 290 would result in a win of £200.

There are currently four major Spread Betting companies. These are Cantor Index (a subsidiary of Cantor Fitzgerald), IG Index, Sporting City Index and Spreadex. All of the companies are based in the UK (but offer trades to customers from overseas) and are regulated by the UK Financial Services Authority. Each company may offer a different quote about the same market. If the top end of the spread quoted by one company lies below the bottom end of the spread quoted by another, there is potential for arbitrage, in the sense of a riskless profit. Say, for example, Cantor Index offered the spread of 240-250 in the above game and Sporting Index offered a quote of 229-239. This is an arbitrage position, since it is possible to buy points at 239 with Sporting Index and sell points at 240 with Cantor Index, and win whatever the result. Much more common than an outright arbitrage, however, is where the average or mid-point of all the quoted spreads lies outside the top (or bottom) end of the spread quoted by at least one market-maker. We call these circumstances a 'Full Quarb' (Vaughan Williams, 2001) and they are the focus of discussion in this paper.[6] The issue is whether, in the circumstances in which Full Quarbs (henceforth Quarbs) occur, it is the average market position or the outlying market position which provides most information.

[5] See Vaughan Williams (2001) for the first use of this terminology.

If the market is assumed to process all known information in the most efficient manner, it might be expected that the average market position will provide the best predictor of the actual outcome. If, on the other hand, one market-maker (or more) possesses privileged information, or at least has a superior ability to process public information, an outlying market quote might provide a more accurate predictor. The outlying position is taken as the mid-point of the quote offered by the market-maker most out of line with the average market position. We seek to ask, firstly, whether either of these positions is systematically superior as a predictor of the actual outcome and, secondly, whether it is possible to implement a trading strategy based on this, using publicly available information, in such a way as to earn returns that are either abnormal or that are superior to those from noise trading.

3. Data and Empirical Approach

3.1 Data

The data that we use to illustrate and test our hypotheses are taken from the market for 'bookings' in English Premier League football matches during the 1999/2000 and 2000/2001 seasons. This market is one of the most heavily traded of all the spread betting markets on offer. The asset which is traded is an index of 'fair play' based on the number of disciplinary cards issued by referees to players in any given football match. It works on the basis of awarding ten points for each 'yellow card' (caution) and 25 for each 'red card' (dismissal). Two cautions in a match earn a player an automatic dismissal from the field of play for the duration of the match. 35 points is the maximum for any individual player, comprising 10 points for the first yellow card and 25 points for dismissal bought about by the second caution. The average bookings points score for games in our data set is in the region of 40. The width of the spread set by companies in the bookings market varies from three to four. The outcome is, of course, bounded from below by zero (no cards). In theory an upper bound also exists but, in practice, it is never approached. For example, in our data, the highest bookings score for a game is 150, well below the theoretical maximum.

Spreads for the bookings market were collected for up to five companies that existed during this period. The opening prices are usually announced one or two days before each game. These were collected from television text services throughout the two seasons. Data

[6] This can be distinguished from a 'Simple Quarb' (Vaughan Williams, 2001) which occurs when the ask price of one quote is equal to the bid price of another.

are available on 207 matches for the 1999/2000 season and on 240 matches for the 2000/2001 season.

For the first stage of the analysis, we need to be able to identify an outlying spread. This means that we require a minimum of three quotes to be available. Further, the quotes must not be symmetric about the average.[7] Applying these restrictions, we end up with a sample of 102 matches for the first stage. For the second stage, in which we compare returns on Quarb positions to others, we include all the matches for which we have data.

For each match we calculate the actual bookings score using data on the number of players booked and sent off taken from the data published by the specialist football magazine, 'Match'. We also have available for each game a second set of spread quotes taken once the market has settled down (denoted as 'Settled Prices'). These data are taken from the figures published by the daily sports newspaper, the Racing Post.

3.2 Empirical Approach

We consider two alternative predictors of the actual value of the asset: (i) the average mid-point of the market bid-offer spreads (MID); (ii) the mid-point of the outlying bid-offer spread (OUTLIER). Our empirical approach is in two stages. In the first stage we seek to establish whether either of these two measures is consistently superior as a predictor of the asset value. In the second stage, we investigate whether the traders can utilise the information from stage one to construct a profitable trading strategy.

A natural starting point for the first stage is to posit competing hypotheses about the determination of the actual value of the asset (in our case bookings points):

$$H_1\text{: } V_i = a + b.MID_i + u_i \quad (1)$$

$$H_2\text{: } V_i = b + c.OUTLIER_i + v_i \quad (2)$$

where V_i is the value of the asset, i (in our case the number of bookings points at the end of each game) and u_i and v_i are disturbance terms.

As H_2 cannot be written as a restriction of H_1, we are choosing between alternative, or non-nested, models. A wide literature exists on choosing between such models. The simplest approach is to choose the model with the lowest value of a particular model selection criterion, such as those of Schwarz, Bayes or Akaike. These criteria all impose a penalty on the log-likelihood in relation to the number of parameters being estimated. For

[7] For example, assume there are three quotes available: 32-35, 34-37 and 36-39. The two outlying spreads are symmetric in the sense that the mean of the outlying mid-points equals the market mean.

example, the Akaike Information Criterion (AIC) can be given by $AIC = -2(\log likelihood) + 2(k + 1)$ where k is the number of regressors. However, although such model selection criteria can rank the predictors in order of performance, they do not provide a formal test of whether the preferred predictor is significantly superior. Formal tests for choosing between non-nested models include the Cox test (Cox, 1961; 1962) and the J-test (Davidson and MacKinnon, 1981). A standard procedure in such tests is to estimate equations (1) and (2) and obtain the fitted values in each case ($V1_i$ and $V2_i$ respectively). The fitted values are then used to construct more general models as follows:

$$V_i = a + b.MID_i + d.V2_i + m_i \quad (1a)$$

$$V_i = b + c.OUTLIER_i + e.V1_i + n_i \quad (2a)$$

A standard t-test that $d = 0$ provides a valid test in favour of H_1, that is the market mid-point is the better predictor. Similarly, a test that $e = 0$ provides a test in favour of H_2, that is the outlier is the better predictor. If neither H_1 nor H_2 is rejected by these tests then there is evidence that neither predictor is significantly superior.

There is a complication in our case in that the dependent variable (the number of booking points accrued) cannot fall below zero and is, in effect, censored. Further, the values of the dependent variables are limited to particular values - 10, 20, 25, 30, 35 etc. Given the nature of the data, several estimation procedures suggest themselves. One possibility is to use the Tobit model to take account of the censored nature of the dependent variable. An alternative approach would be to treat the dependent variable as count data and to use count regression models to estimate the model. We report results using both approaches here.

A complication with using these approaches is that standard non-nested tests do not take into account the censored or count nature of the dependent variable and may be unreliable. Consequently, we apply the conditional moment test for omitted variables suggested by Pagan and Vella (1989) to our models, using the fitted values from the alternative model as the potentially erroneously omitted variable.

Following Greene (2000), the general form of the conditional moment test is given by:

$$C = I'M[M'M - M'G(G'G)^{-1}-G'M]^{-1}M'I \quad (3)$$

where M is a matrix in which the rows are the terms of the moment conditions; G is a matrix in which the rows are the terms in the gradient of the log-likelihood function; I is the identity matrix.

We now consider the implied moment condition for model (1). If the variable $V2_i$ has been correctly omitted from model (1), this implies the following expectation for uncensored (non-limit) observations:

$$E[V2.(V - a - b.MID)] = 0 \quad (4)$$

The sample counterpart to this moment restriction is

$$m_i = 1/n\ [S(V2_i.u_i) \text{ where } u_i = V_i - a - b.MID_i] \quad (5)$$

and (in the case of the Tobit model), the sample counterpart for limit observations is

$$m_i = V2_i(- s\,.\,l_i) \quad (6)$$

where s is the standard error of the Tobit regression and $l_i = f(b.MID_i/s)/[1 - j(b.MID_i/s)]$ and where f is the standard normal probability density function and j is the standard normal cumulative density function.

As we are considering only the omitted variable moment condition, M is a single column vector, the elements of which are given by m_i. On the null hypothesis that the variable has been validly omitted, the test statistic, C, follows a chi-squared distribution with one degree of freedom. A similar argument holds for testing model (2).

Assuming that one predictor is found to be systematically superior, the second stage of our analysis aims to determine whether or not this information can be used to construct a profitable trading strategy. Our approach is to calculate the returns that would have been made by a trading strategy based on the superior predictor as revealed by stage one. We then use standard t-tests to test for the existence of abnormal returns and for returns that are superior to those from noise trading. In the first case, we test the null hypothesis that the returns based on this strategy are equal to zero against the alternative that they are positive. Second, we test the null hypothesis that the returns are equal to the returns from all other bets against the alternative that they are greater. To guard against accusations of data mining, we repeat the tests out of sample using a set of data reserved from the Stage 1 analysis.

4. Results

Stage 1

In Table 1, we report the Tobit regression estimates of models (1) and (2) using the mean mid-point as our measure of average. For both models, the predicting variable is positive and strongly significant, but the lower value of the AIC suggest that the market mid-point (model (1)) performs better than the outlier (model (2)). This is confirmed by the estimates of equations (1a) and (2a). The fitted values from model (2) are only weakly significant in equation (1a) and attract a negative sign. The fitted values from model (1), on the other hand, are strongly significant in equation (2a) and render the main predictor insignificant. In other words, applying the standard J-test to our models provides significant evidence in favour of model 1 and against model 2. Lastly, using the conditional moments test to take account of the truncated nature of our dependent variable, we are unable to reject H_1 (that model 1 is the correct specification) at any conventional significance level, whereas we reject H_2 (that model 2 is correct), albeit only at the 10% level.

Using the median value as the market average (reported in Table 2), the results are even clearer. The Tobit regression results are fairly close to those using the mean value, but, in this case, the conditional moments tests suggest rejection of the H_2 at the 5% level, whilst still suggesting that H_1 cannot be rejected.

The count regression results are not reported here, but have exactly the same pattern as the Tobit regressions. Taken together, the evidence appears to be conclusive in suggesting that the market mid-point is systematically superior to the outlier in predicting actual asset values. An obvious question that arises here is whether there exist systematic differences in the performance of outliers between the firms. For example, it may be that one firm has superior information in this market and that outlying prices quoted by this firm will tend to be better predictors than outlying prices quoted by others. To test this hypothesis, we modify model (1) and (2) by allowing the slope variables to take different values for each of the five companies in our sample (not reported here). For each of the regression models, we find no significant differences between the companies, implying that there is no evidence that the performance of any of the firms was systematically superior to the others in our sample.

Stage 2

We now examine whether the information gained from Stage 1, can be used to devise a profitable trading strategy. Specifically, we examine the returns to trading against the subset of outliers in which a Quarb exists, that is, where the market mid-point lies entirely outside the outlying spread. The logic behind this strategy is that, in the absence of other information, the mid-point of all spreads provides us with an obvious point estimate of the expected value of the asset. On this basis, we can expect positive returns as long as this value is greater (less) than the price at which we buy (sell). For example, if the mean mid-point of all spreads in the market is 35.5 and the outlying spread is 36-39, the strategy would be to sell bookings with the outlying company at 36. If the outlying spread is 32-35, the strategy would be to buy bookings at 35. An outlying spread of 35-38 would imply no trade.

Note that when a true arbitrage position exists, the Quarb strategy may suggest a buy and a sell bet at the same time. Our strategy is based on opposing the outliers so, in these cases, only one trade is allowed, using the quote furthest from the mid-point. For example, if there are three quotes of 24-28, 26-30 and 29-33, the mean mid-point is 28.3. The outlying quote (furthest from the mean) is 29-33 and the strategy would be to sell at 29.

In Table 3, we report the mean return to all bets and returns to trades based on the Quarb strategy for the ‘within sample’ data on which we conducted our Stage 1 tests - the 1999/2000 season. To avoid charges of data mining in a search for profitable trading opportunities, in Table 4, we report returns using the reserved sample, namely matches with available data for the 2000/2001 season.

For each game there are two possible bets: buy bookings and sell bookings. In the 1999/2000 season (‘within sample’) this means there are 414 bets in total. The mean return to a unit £1 stake on every possible bet was £0.077 (standard error 1.349). Of these bets, the Quarb strategy suggests 60 trades. The mean return to the Quarb trades is £9.817 (standard error 3.660). Given the nature of spread betting it is difficult to present this as a rate of return. However, on the basis of £1 being placed on each Quarb trade, a bettor would have won £973 on 36 winning bets and lost £384 on 22 losing bets, a net profit of £589 over the season. The remaining two bets yielded a return of zero.

We conduct two different t-tests on these returns. The first is a test of the null hypothesis that the mean return to Quarb trades is zero against the one-sided alternative that it is positive. The second test is of the null hypothesis that the mean return to Quarb trades is equal to the mean return to all other bets against the one-sided alternative that the Quarb mean return is greater. For the 1999/2000 season, both the t-tests suggest rejection of the null at the 1% significance level. In other words, there is evidence of significantly abnormal and superior returns in this sample of data.

Using the 2000/2001 season (reported in Table 4), the mean return to a unit stake on all bets is -£0.202 (standard error 1.024). With this sample, there are 80 trades suggested by the Quarb strategy. The mean return to these bets is lower than in 1999/2000 at £4.825 (standard error 2.152). However, the mean is significantly greater than zero at the 5% level and greater than the mean return to other bets at the 1% level, suggesting the Quarb strategy would have led to superior returns even in the reserved sample. In this season, a £1 stake on each Quarb bet would have yielded winnings of £836 on 50 winning bets and losses of £450 on 28 losing bets, a net profit of £386.

Robustness Check 1: are these prices available?

We have calculated returns to Quarb trades using the opening prices - the initial prices quoted by bookmakers prior to trading. These include some prices that represent one side of a true arbitrage position. A natural question is whether the quoted prices are actually available to bettors. The authors' experience in these markets suggests that, in general, the published prices are available, but, particularly when they are part of an arbitrage position, prices often move within a few minutes of the opening of the market. Further, bookmakers sometimes limit stakes in the case of prices that are one side an arbitrage position. Consequently, in Tables 3 and 4, we also report returns excluding all games in which the opening prices represent a true arbitrage. This restricts the sample size considerably, to just 19 trades in the 'within sample' (Table 3). The small sample size, along with the fact that all returns are not normally distributed, means that we cannot appeal to the Central Limit Theorem for our t-tests. However, the mean returns to Quarb bets are still positive and superior to returns to other bets. For the reserved sample (Table 4), 32 bets remain, suggesting that the t-tests are valid. The mean return to Quarb bets in this sample is 5.625 (standard error 3.418) and the t-tests suggest evidence in favour of abnormal returns at the 10% significance level and in favour of superior returns at the 5% level.

An alternative approach is to examine the returns to the Quarb bets, based on the settled prices. As explained above, these are published some time after the market has settled down and almost always represent prices that are available in practice. Returns to the Quarb trades suggested by the opening prices, but evaluated at these settled prices, are given in the third column of Tables 3 and 4. For the 'within sample', returns to Quarb bets are lower than at opening prices, but still significantly positive and significantly superior to other bets, both at the 5% level. For the 'reserved sample' the returns are positive and superior in each case, although significance levels are lower than for the 'within sample'.

Robustness Check 2: controlling for risk

The role of risk in these markets is not the primary focus of this paper, but it might reasonably be asked whether the positive returns to Quarb trades can be explained by risk. The two types of bets contained within this spread betting market (buy and sell) represent directly contrasting risk positions. For example, given a spread of 36-40, a bettor who sells bookings at 36 with a stake of £1 knows with certainty that winnings cannot exceed £36. The maximum loss, is unknown, however, and (effectively) open-ended. On the other hand a bettor who buys at 40 knows with certainty that the maximum loss will be £40 but the winnings are open ended. In this limited sense, we can say that sell bets are riskier propositions to bettors than buy bets. Consequently, sell bets may attract a risk premium. If sell bets are over-represented in the sample of Quarb bets, then this may go some way to explaining our finding of abnormal returns.

We control for differential risk by separating out returns to sell and buy bets and then testing whether the returns to sell (buy) bets involving Quarbs are significantly greater than the returns to all other sell (buy) bets. The separate returns 'within sample' are reported in Table 5 and for the 'reserved sample' in Table 6. There is indeed some evidence of a risk premium to sell bets. The mean return to all Quarb sell bets in the reserved sample is 2.879 points (standard error 2.034), whereas the mean return to buy bets is -6.285 points (standard error 1.965). There is also a premium in the 'reserved sample', although of a lower magnitude. However, in both samples, the returns to Quarb bets are superior than to others for both buy and sell bets. In the 'within sample', the Quarb sell bets are significantly superior to the other sell bets at the 1% level, whilst the Quarb buy bets are significantly superior at the 10% level. In the 'reserved sample' both types of Quarb bets are superior at the 10% level. When we estimate the returns using the settled prices, the Quarb returns are still higher for both types of bets, but the level of significance is reduced. Eliminating all games in which the opening prices suggest a true arbitrage leads to very small sample sizes for each of the buy and sell bets. However, even in this case, the Quarb returns (not reported here) are superior for both types of bet and in both samples.

In summary, both within and out of sample, there is evidence that bettors can take advantage in practice of the superior performance of the market mid-point as a predictor to earn positive and superior returns based on published prices. Even having controlled for differential risk, there is still considerable evidence that trading on the basis of Quarbs allows bettors to make positive returns.

5. Conclusions

In this study, we have examined the scope for earning superior returns on the basis of simply defined trading rules in a rapidly growing sector of the betting market, known as index (or spread) betting. In this form of betting, bookmakers quote a bid-offer spread about the result of some future event, and bettors are invited to buy (sell) at the top (bottom) end of the spread. The particular market examined, chosen for its relative popularity and high profile, is based on the number of disciplinary points awarded in identified football matches in the UK (the ‘bookings market’).

The results of our study suggest that the mid-point of all quotes is a better forecast of the actual outcome in the bookings market than is the mid-point of the spread offered by the market outlier. This casts doubt on a hypothesis that market-makers who set quotes out of line with the prevailing view do so because they possess better (even privileged) information, or that they are able to process a given set of information more effectively than the market as a whole.

Using the notion of quasi-arbitrages or Quarbs, we find that it is possible to devise a trading strategy on the basis of the outlying spread that yields returns, both within and out of sample, that are consistently positive and superior to those that might be expected from noise trading. Further, this result is robust to a variety of checks to control for the possibility that published prices might not be available, and also for the impact of differential risk.

References

Avery, C. and Chevalier, J. (1999), 'Identifying Investor Sentiment from Price Paths: the case of football betting', *Journal of Business*, 72 (4), 493-521.

Brown, W.O. and Sauer, R.D. (1993), 'Does the Basketball Market Believe in the Hot Hand? Comment', *American Economic Review*, 83 (Dec.), 1377-86.

Cain, M., Law, D. and Peel, D. (2000), 'The Favourite-Longshot Bias and Market Efficiency in UK Football Betting', *Scottish Journal of Political Economy*, 47 (1), 25-36.

Camerer, C. (1989), 'Does the Basketball Market Believe in the "Hot Hand"?', *American Economic Review*, 79, 1257-61.

Copeland, T. and Galai, D. (1983), 'Information Effects on the Bid Ask Spread', *Journal of Finance*, Vol. 38, No. 5, pp. 1457-1469.

Cox, D. (1961), 'Tests of Separate Families of Hypotheses', *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, 1, Berkeley: University of California Press.

Cox, D. (1962) 'Further Results on Tests of Separate Families of Hypotheses', *Journal of the Royal Statistical Society, Series B*, 24, 406-24.

Davidson, R. and MacKinnon, J. (1981), 'Several Tests for Model Specification in the Presence of Alternative Hypotheses', *Econometrica*, 49, 781-93.

De Long, J.B., Shleifer, A., Summers, L.H. and Waldmann, R.J. (1990), 'Noise Trader Risk in Financial Markets', *Journal of Political Economy*, 98 (August), 703-38.

De Long, J.B., Shleifer, A., Summers, L.H. and Waldmann, R.J. (1991), 'Survival of Noise Traders in Financial Markets', *Journal of Business*, 64 (Jan.), 1-20.

Gandar, J.M., Zuber, R.A., O'Brien, T and Russo, B. (1988), 'Testing Rationality in the Point Spread Betting Market', *Journal of Finance*, 43 (Sept.), 995-1008.

Gandar, J.M., Dare, W.H., Brown, C.R. and Zuber, R.A. (1998), 'Informed Traders and Price Variations in the Betting Market for Professional Basketball Games', J*ournal of Finance*, 53 (Feb.), 385-401.

Glosten, L. and Milgrom, P. (1985), 'Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders', *Journal of Financial Economics*, 14, 71-100.

Golec, J. and Tamarkin, M. (1998), 'Bettors Love Skewness, Not Risk, at the Horse Track', *Journal of Political Economy*, 106, 1 (March), 205-25.

Gray, P.K. and Gray, S.F. (1997), 'Testing Market Efficiency. Evidence from the NFL Sports Betting Market', *Journal of Finance,* 52 (Sept.), 1725-37.

Greene, W.H. (2000), *Econometric Analysis,* New Jersey: Prentice-Hall.

Haigh, J. (1999), '(Performance) Index Betting and Fixed Odds', *The Statistician*, 48, 3, 425-434.

Hurley, W. and McDonough, L. (1995), 'A Note on the Hayek Hypothesis and the Favorite-Longshot Bias in Parimutuel Betting', *American Economic Review*, 85, 949-55.

Jackson, D.A. (1994), 'Focus on Sport: Index Betting on Sports', *The Statistician*, 43, 2, 309-315.

Kyle, A. (1985), 'Continuous Auctions and Insider Trading', *Econometrica*, 53, 1315-1335.

Pagan, A. and Vella, A.F. (1989), 'Diagnostic Tests for Models Based on Individual Data: a survey', *Journal of Applied Econometrics*, 4 (supplement), S29-S59.

Paton, D, Vaughan Williams, L and Siegel, D (2002), 'A Policy Response to the E-Commerce Revolution: the case of betting taxation in the UK', *Economic Journal*, forthcoming.

Sauer R.D. (1998), 'The Economics of Wagering Markets', *Journal of Economic Literature*, 36**,** 2021-64.

Schnytzer, A. and Y. Shilony (1995), 'Inside Information in a Betting Market', *Economic Journal*, 105, 963-71.

Shin H.S. (1993), 'Measuring the Incidence of Insider Trading in a Market for State-Contingent Claims', *Economic Journal*, 103, 1141-1153.

Shleifer, A. and Vishny, R.W. (1997), 'The Limits of Arbitrage', *Journal of Finance*, 52 (March), 35-55.

Thaler R. and Ziemba W. (1988), 'Parimutuel Betting Markets: Racetracks and Lotteries' *Journal of Economic Perspectives*, 2, 161-74

Vaughan Williams, L. and Paton, D. (1997), 'Why is There a Favourite-Longshot Bias in British Racetrack Betting Markets?', *Economic Journal*, 107, 150-158.

Vaughan Williams, L. (2001), 'Can Bettors Win? A Perspective on the Economics of Betting', *World Economics*, January, 2 (1), 31- 48.

Woodland, L.M. and Woodland, B.M. (1994), 'Market Efficiency and the Favorite-Longshot Bias: the baseball betting market', *Journal of Finance,* 49 (March), 269-79.

Table 1: Tobit Regression Estimates of Bookings Points: mid-points against outliers (means)

	1 Model 1	2 Model 1a	3 Model 2	4 Model 2a
MID_i	1.861*** (0.334)	3.162*** (0.812)	-	-
$V2_i$	-	-0.964* (0.549)	-	-
$OUTLIER_i$	-	-	1.428*** (0.342)	-1.377 (0.784)
$V1_i$	-	-	-	1.699*** (0.436)
Constant	-43.271*** (14.417)	-64.269*** (18.717)	-	32.528 (19.715)
N: uncensored	90	90	90	90
N: censored	12	12	12	12
Log-likelihood	-446.88	-445.36	-452.42	-445.36
AIC	899.6	898.72	910.84	898.72
Conditional Moments Test	0.453		2.787*	
Conclusion	Do not reject H_1		Reject H_2 at 10% level	

Notes:
(i) The dependent variable in each case is the number of bookings points as described in the text.
(ii) Figures in brackets are standard errors.
(iii) *** indicates significance at the 1% level; ** at the 5% level; * at the 10% level.
(iv) The conditional moments test is based on that of Pagan-Vella for omitted variables and distinguishes between the two, non-nested models. See the text for more details.

Table 2: Tobit Regression Estimates of Bookings Points: mid-points against outliers (medians)

	1 Model 1	2 Model 1a	3 Model 2	4 Model 2a
MID_i	1.860*** (0.321)	2.414*** (0.623)	-	-
$V2_i$	-	-0.450 (0.433)	-	-
$OUTLIER_i$	-	-	1.428*** (0.342)	-0.642 (0.619)
$V1_i$	-	-	-	1.298*** (0.335)
Constant	-43.597*** (13.963)	-51.305*** (15.770)	-	16.133 (16.927)
N: uncensored	90	90	90	90
N: censored	12	12	12	12
Log-likelihood	-445.98	-445.44	-452.42	-445.44
AIC	897.96	898.88	910.84	898.88
Conditional Moments Test	0.273		4.605**	
Conclusion	Do not reject H_1		Reject H_2 at 5% level	

Notes:
See Table 1.

Table 3: Tests for Abnormal/Superior Returns 1999/2000 Season (Within Sample)

	1 All Quarbs	2 Excluding Arb Positions	3 Settled Prices
Number of bets	414	282	408
Mean return to all bets	0.077	-0.915	-0.346
	(1.349)	(1.706)	(1.325)
Number of Quarbs	60	19	59
Mean return to Quarbs	9.817	14.158	6.677
	(3.660)	(7.890)	(3.661)
t-tests:			
Abnormal Returns	2.682***	-	1.824**
Superior Returns	2.897***	-	2.189**

Notes:
(i) Opening prices are those announced at the start of the market. Settled prices are those published in the Racing Post on the day of the game in question. In each case, returns are calculated to the most favourable price.
(ii) Figures in brackets are standard errors.
(iii) The t-tests for abnormal returns are 1-tailed tests of the hypothesis that the mean return to Quarb bets = 0 against the alternative that the mean > 0. The tests for superior returns are 1-tailed tests of the hypothesis that the mean return to Quarb bets = the mean return to other bets against the alternative that the Quarb mean is greater. The t-tests for superior returns allow for unequal variances across samples. Tests are not performed if the sample size is less than 20.
(iv) *** indicates the null hypothesis is rejected at the 1% level; ** at the 5% level; * at the 10% level.

Table 4: Tests for Abnormal/Superior Returns 2000/2001 Season (Out of Sample)

	1 All Quarbs	2 Excluding Arb Positions	3 Settled Prices
Number of bets	480	362	470
Mean return to other bets	-0.202	-0.801	-0.225
	(1.024)	(1.215)	(1.033)
Number of Quarbs	80	32	80
Mean return to Quarbs	4.825	5.625	2.413
	(2.152)	(3.418)	(2.132)
t-tests:			
Abnormal Returns	2.242**	1.646*	1.131
Superior Returns	2.474***	1.930**	1.157

Notes:
See Table 3.

Table 5: Tests for Risk-Adjusted Abnormal/Superior Returns 1999/2000 Season (Within Sample)

	1	2	3	4
	Opening Prices		**Settled Prices**	
	Sell Bets	**Buy Bets**	**Sell Bets**	**Buy Bets**
Number of bets	207	207	204	204
Mean return to all bets	4.913	-4.758	4.578	-5.270
	(1.866)	(1.895)	(1.836)	(5.462)
Number of Quarbs	25	35	24	35
Mean return to Quarbs	19.72	2.743	17.125	-0.486
	(3.171)	(5.590)	(3.282)	(5.462)
t-tests				
Abnormal Returns	6.219***	0.491	3.670***	-0.089
Superior Returns	4.471***	1.523*	4.319***	0.997

Notes:
See Table 3.

Table 6: Tests for Risk-Adjusted Abnormal/Superior Returns 2000/2001 Season (Out of Sample)

	1	2	3	4
	Opening Prices		**Settled Prices**	
	Sell Bets	**Buy Bets**	**Sell Bets**	**Buy Bets**
Number of bets	240	240	235	235
Mean return to all bets	2.829	-3.233	2.451	-2.901
	(1.439)	(1.514)	(1.454)	(1.450)
Number of Quarbs	53	27	53	27
Mean return to Quarbs	6.113	2.296	3.811	-0.333
	(2.377)	(4.377)	(2.355)	(4.327)
t-tests				
Abnormal Returns	2.572**	0.525	1.618*	-0.077
Superior Returns	1.438*	1.345*	0.599	0.632

Notes:
See Table 3.