The European Journal of Finance

Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713684998

Anyone for Tennis (Betting)?

David Forrest [a]; Ian Mchale [b]
[a] Centre for the Study of Gambling, University of Salford, Salford, UK
[b] Centre for Operational Research and Applied Statistics, University of Salford, Salford, UK

Online Publication Date: 01 December 2007
To cite this Article: Forrest, David and Mchale, Ian (2007) 'Anyone for Tennis (Betting)?', The European Journal of Finance, 13:8, 751 - 768
To link to this article: DOI: 10.1080/13518470701705736
URL: http://dx.doi.org/10.1080/13518470701705736

The European Journal of Finance
Vol. 13, No. 8, 751–768, December 2007

Routledge
Taylor & Francis Group

Anyone for Tennis (Betting)?

DAVID FORREST* & IAN McHALE**
*Centre for the Study of Gambling, University of Salford, Salford, Greater Manchester, UK, **Centre for Operational Research and Applied Statistics, University of Salford, Salford, Greater Manchester, UK

ABSTRACT *The most robust anomaly noted in the literature on wagering markets is (positive) longshot bias: over a period of 50 years, it has been well documented in horse betting that higher expected returns accrue to short- than to long-odds bets. However, a few examples of betting markets with zero or negative bias have been found, for example in certain American sports. The understanding of longshot bias is likely to be informed by comparing and contrasting conditions in markets displaying positive, zero, and negative bias but, to date, relatively few markets have been examined. This paper employs a large data set on professional men's tennis matches and a new econometric approach to the estimation of the relationship between returns and odds. It finds positive bias throughout the range of odds. It discusses this finding in the context of the debate on why biases exist and persist in wagering markets, focusing particularly on bettors' attitudes towards risk and skewness.*

KEY WORDS: Sports betting, longshot bias, risk preference

1. Introduction

Wagering markets are specialist financial markets where bettors purchase state-contingent assets. As in other financial markets, most studies (Sauer (1998) and Vaughan Williams (1999) provide excellent surveys) indicate that, broadly, prices are efficient but, again as in other financial markets, anomalies have been noted. The most celebrated and robust of these is termed *longshot bias*: in American, British, and Australian horse-betting markets, whether organized by bookmakers or on a pari mutuel (tote) basis, superior returns accrue to a strategy of placing short-odds bets relative to a strategy of placing long-odds bets. This stylized fact has been repeatedly confirmed by empirical studies dating as far back as Griffiths (1949), and a similar pattern has been shown to characterize dog-race-betting markets (Cain *et al.* (1996) for the UK; Terrell and Farmer (1996) for the USA). It represents an anomaly because in most conventional financial markets, expected return is higher in respect of assets associated with greater risk. In horse- and dog-betting markets, the opposite occurs.

A number of explanations have been proposed to account for the anomaly; for a survey, see Vaughan Williams (1999). Many start from the notion that the market is inefficient because of the presence of noise traders and high transactions costs that discourage arbitrage.[1] But, in the present paper, we are concerned principally with explanations based on bettors' risk preferences. Golec and Tamarkin (1998) suggest that bettors might be risk-averse but "skewness-loving".

Correspondence Address: David Forrest, Centre for the Study of Gambling, University of Salford, Salford, Greater Manchester M5 4WT, UK. Tel.: +44 161 295 3674; Fax: +44 161 295 2130; Email: d.k.forrest@salford.ac.uk

1351-847X Print/1466-4364 Online/07/080751–18
DOI: 10.1080/13518470701705736

In this explanation, the assumption made in most economic analysis, namely that investors are risk-averse, is retained; but it is hypothesized that at some level of wealth, the utility function would become risk-loving, and so betting at sufficiently long odds would be attractive. On this view, longshots provide evidently poor value (low expected return, high variance in return) in equilibrium, but bettors still accept the odds because they are compensated by high positive skewness. One of the implications of risk-aversion, combined with a taste for skewness, is that there should be a non-monotonic relationship between odds and returns. The present paper tests whether this implication of the Golec and Tamarkin approach is supported by the data from a betting market with a very wide range of odds.

One doubt that may be raised concerning this approach is that, in the case of sports betting, most wagers are observed to be small. Therefore, neither winning nor losing would take the bettor far from initial wealth. With a narrow range of wealth outcomes from the bet, the shape of the utility of wealth function is unlikely to be perceptibly different from linear, yet the explanations in the literature refer to concave and convex segments when accounting for the pattern in returns and indeed for the act of betting itself. However, if we accept that there are recreational benefits from betting, participants may have preferences over the amount they win and then a utility function defined over payoffs from a unit bet may conceivably be characterizable as having the shape of the Freidman–Savage utility function. In this paper, the focus is on whether such bettor preferences over the amounts won account for the presence of longshot bias in betting markets.

Any hypothesis needs to account not just for the dominant stylized fact of positive longshot bias in horse- and dog-betting markets, but also for any exceptions to the general rule that bets on short-odds events are financially superior to bets on long-odds events. And exceptions have been found. For example, no bias has been detected in Asian racetrack markets (Busche, 1994; Busche and Hall, 1988; Busche and Walls, 2000), and Vaughan Williams and Paton (1998) found bias absent also in the case of wagering on "higher-grade handicap" races in Britain. Beyond racing, Forrest *et al.* (2005) reported unbiased odds from examining bets available on nearly 10 000 English football matches, and Woodland and Woodland (1994, 2001, 2003) have even identified a reverse, i.e. negative, bias in respect of betting on the American team sports of baseball and (ice) hockey.[2] As noted by Vaughan Williams and Paton (1998), a successful general theory of wagering markets would explain any contrast in the pattern of returns as observed in different sub-sectors of betting.

This paper tests for the direction of bias in a hitherto unexplored betting market, that on men's professional tennis. This adds to the existing taxonomy of which markets display a positive and which a negative bias. Evidence to date relates to horse and dog betting and certain team sports. Attempts to account for different directions of bias have been based mainly on contrasting conditions in the horse- and baseball-betting markets. To be convincing, such explanations should be able to predict the direction of bias in other sub-sectors, such as tennis. Hence, evidence is needed from markets not previously studied in order to discipline the debate on longshot bias. We seek to provide such evidence for tennis and to examine the extent to which it is consistent with previous attempts (particularly, those linked to specific hypotheses about bettor risk preferences) to explain why bias is positive and negative in different markets. We argue below that tennis is an especially valuable source of evidence because of the very wide range of odds available on singles bets.

2. Width of Odds

One reason for contrasting patterns between different areas of betting may be that different widths of odds are observed in different markets (Woodland and Woodland, 1999). Consider some relevant facts, first from racing. Vaughan Williams and Paton (1998) found no bias in English horse-racing odds when their sample was confined to handicap races. By design, runners in a handicap race are intended to have fairly equal chances of winning, and the odds range may be expected therefore to be relatively narrow; this is indeed the case. We recently reviewed odds data collected on 31 037 runners in British (jumps) races. The proportion of runners who were strong favourites, defined as having odds of evens or shorter, was much lower in handicap than in non-handicap races (0.7% compared with 2.8%). The reason for the failure of an odds bias to emerge in handicap races might be that one observes only a compressed range of odds compared with other racing.

In sports-betting markets, one does not generally find extreme favourites or rank outsiders. In English soccer and in the American team sports, all contests are prospectively quite close because the sport is organized to promote competitive balance. So, true probabilities of a particular team winning will almost never justify odds such as 10/1, which are available routinely in, say, the British racing programme. And whereas, in our large sample of British (jumps) races, much less than 3% of runners had starting odds at evens or shorter, it is not unusual in a sports fixture to find *both* teams quoted at odds-on. In this setting of team sport where all bets are in a short range either side of evens, positive longshot bias is *not* found.

In contrast to most sports leagues, Singapore's S-League (soccer) is spectacularly unbalanced,[3] and extreme odds *are* often available. For example, one team was recently 1/20 to win, notwithstanding that it was playing away. Forrest and Simmons (2001) found that odds on S-League games did display a marked positive bias in contrast to the difficulty in establishing any bias in the odds on English football and the clear negative bias in American team sports. Moreover, within English football, betting on exact scores (rather than on the home/draw/away result) does offer the possibility to wager on highly unlikely events (such as a scoreline of 6–5), and in this sector, a positive bias has again been reported (Cain *et al.*, 2000).

A pattern appears to emerge. Where odds cover a wide range, as in betting on a typical horse race or the exact score in a soccer game, then positive longshot bias characterizes the odds. Where the odds range is narrow, as in betting on the winner in most team sports, findings are of zero or negative bias. This picture is consistent with the role accorded to skewness in returns in Golec and Tamarkin (1998). In their model, bettor utility is related positively to expected return, negatively to variance, and positively to skewness. If only odds in a range around evens are available, as in markets on the winner of team sports events, skewness is close to zero on all bets, and risk-aversion then yields a positive relationship between the mean and variance of returns, i.e. "negative" or reverse longshot bias occurs. Where some outcomes that are very unlikely become available for betting, as in many horse races and all exact-score soccer betting, such wagers attract bettors with a preference for high skewness. Bets where low returns coexist with high risk still sell because they offer high skewness: "positive" or traditional bias is then observed in the general pattern of returns.

The "love of skewness" attributed to bettors by Golec and Tamarkin may be interpreted in the context of a function relating utility to the payoff, X, from a successful unit bet. A taste for skewness implies that such a utility function will display less risk-aversion, as higher values of X are considered (i.e. the third derivative of utility with respect to X will be positive). Thus, the utility function could display risk-aversion for low values of X but become risk-loving in a range corresponding to higher values of X. In this case, the utility function would resemble that of Friedman and Savage (1948).[4] Such a utility function could permit the reconciliation of the positive longshot bias found in horse-betting markets with the negative bias reported for baseball. Because baseball is relatively competitively balanced, all bets carry a (low) return, with X quite close to one. If the utility function is concave for the relevant range of X, bettors will act as if risk-averse in choosing between bets and require a risk premium for backing the underdog: there will be a negative longshot bias. In contrast, bets on genuine outsiders are regularly available in horse races. The bettor can choose between bets offering low and high values of X. If X for the longshot lies beyond the value at which the point of inflection of the utility function occurs (it changes from concave to convex), the longshot bet may be preferred by the representative bettor: odds will then evolve so that in equilibrium, the longshot bet is associated with lower expected returns, i.e. there is a standard (positive) longshot bias.

The evolution of odds to reflect bettor preferences is consistent with a model in which bookmakers seek to maximize their expected winnings from bettors. Such bookmakers will exploit bettor preferences when they determine the odds. For example, if bettors prefer longshots over favourites at odds that deliver the same expected return, bookmakers will shorten the odds, that is, make them more unfavourable (Levitt, 2004). Profit-maximizing behaviour in odds setting therefore permits the relationship between the odds and expected bettor returns to reveal bettor preferences.

Consistent with this approach, Cain and Peel (2004) offered empirical evidence that findings for racing and baseball are not necessarily contradictory. They noticed that one large data set employed in a racing study had a range of implicit probabilities (derived from the odds) of 0.003 to 0.5, whereas the range of *favourite* odds in the baseball data analysed by Woodland and Woodland (1994) was roughly 0.52 to 0.7. The overlap of odds ranges was therefore limited. Utility functions may be such as to generate a negative return–risk (mean–variance) relationship in one range of odds but a positive relationship in another. That horse and baseball betting exhibited opposite biases was not necessarily an anomaly given the different ranges used in the estimation. Cain and Peel therefore advocated merging data from different betting markets to cover as wide a range of odds as possible and then estimating the relationship between return and probability (i.e. odds), allowing for the possibility of a turning point.

In their illustration, Cain and Peel created composite data sets from US and Hong Kong horse racing and US baseball markets and found some evidence of a turning point in the relationship between odds and expected return, consistent with a Friedman–Savage shaped utility function. But such a meta-study has drawbacks. Some approaches to longshot bias, such as taken by Vaughan Williams and Paton (1998), account for differences in results as between racing and baseball by differences in transactions costs and differences in the proportions of informed and uninformed bettors. If these factors are important, aggregation of the data may not permit legitimate inferences to be drawn concerning underlying bettor utility functions. Indeed, if Hong Kong and American racetracks and sports books on American baseball attract different types of bettor, there might not be a common representative utility function at all.

In the present paper, we analyse the relationship between odds and returns in the betting market on a single sport, men's professional tennis. Without the need to aggregate across sports, this market allows the analysis of wagering opportunities across almost the complete odds range from zero to one probability-odds. This arises from the way in which tennis tournaments are structured. In early rounds, the seeding system ensures a large proportion of matches that pitch the best players against those with much lower rankings and sometimes against players who have barely qualified to participate at all. In such matches, odds for favourites can be shorter than any found in racing and odds against underdogs very long. In later rounds, the best players typically play each other and odds on offer resemble those found in most team sports. The opportunity to observe return–risk patterns over a wide range of odds in the betting market on a single sport motivates our study of tennis. Given the range of odds on offer, it will be possible to test whether there is a central maximum in the relationship between odds and yield, as Cain and Peel found, without the danger that the result is spurious because of aggregation across heterogeneous markets.

This betting market has other advantages to the analyst seeking illumination as to the sources of longshot bias in betting markets generally. In contrast to team sports, individual players seldom attract committed fans whose allegiance may be reflected in betting volumes; such "sentiment" has the potential to distort odds in the betting market (Avery and Chevalier, 1999; Forrest and Simmons, 2002). Indeed, sentiment bias could conceivably have created a misdiagnosis of reverse bias in American team sports, because the stronger teams serve larger markets and their greater resources make them favourites disproportionately often; the relatively unfair odds, which are, on an average, posted with respect to favourites, could then be designed to exploit preferences for backing particular teams rather than to cater for return–risk preferences.

Compared with horse racing, tennis betting is also less characterized by factors, other than bettors' risk preferences, that have been linked to longshot bias in some of the previous literature. First, it is a highly specialized betting field, and prices will not have the potential to be distorted by noise traders to the same extent as at the racetrack where money from casual, relatively uninformed visitors can be influential: casual betting appears to be rare in tennis (for example, all British mainstream newspapers list racing and football odds, but one has to purchase the specialist betting daily to find listings for tennis odds). Second, any biases that do emerge from the presence of noise traders are more likely to be arbitraged away because transactions costs in the form of bookmaker commissions were less than half of those in racing (in our data set, introduced in the next section, the mean over-round across all matches was approximately 1.075).

An additional point of contrast between horse racing and tennis is that there is less complex information to process in the latter field. Tennis has only two competitors who compete relatively frequently, whereas a horse race has numerous runners and jockeys.

There is no prior evidence, based on large data sets, regarding longshot bias in betting markets for individual as opposed to team sports. This is despite the fact that it would be likely to discipline the debate on sources of the longshot anomaly by adding to the taxonomy of sectors according to whether they display a positive, zero, or a negative bias. Cain *et al.* (2003) tabulated returns on bets in different odds ranges in boxing, snooker, and tennis. In each case, very poor returns were associated with betting on rank outsiders but, otherwise, no clear pattern emerged in respect of the relationship between odds and returns. This is unsurprising because the samples, particularly for boxing and tennis, were very small. For tennis, they considered only bets available on 91 matches in the 1996 Wimbledon tournament. The standard deviation on the return to a unit bet is very high since its value is always either -1 or a positive number. Finding statistically significant relationships between odds and return therefore requires large data sets. Shmanske (2005) studied the golf-betting market. Again, the sample was limited in size, but evidence was found for an apparent positive bias in the sense that short odds outcomes (such as any unspecified outsider winning the tournament) yielded the highest returns to bettors; but the small sample size makes the conclusion tentative.

The availability of a large data set for tennis permits us to offer a potentially more reliable analysis of the return–risk relationship in the context of an individual sport and one where odds on offer cover a wider range than in team sports or even horse racing. The motivation is to further understanding of longshot bias and width of odds. Specifically, we test for whether Cain and Peel's finding that there is a central maximum in the relationship between odds and return is replicated when one employs a data set from the cleaner environment represented by a single rather than an array of sports.

3. Data

We obtained odds for each competitor in each of over 8500 men's singles tennis matches (i.e. in respect of over 17 000 possible bets) played in international tournaments between June 2002 and August 2005. The source was the archive section of the website www.tennis-data.co.uk.[5] Odds are listed for eight bookmakers[6] (though the set of matches covered differed slightly across the firms). All eight enjoy a high reputation and are known as tennis specialists, willing to accept wagers into the thousands of euros (some other bookmakers may offer tennis betting only as a service to their clients and will have low betting limits). Of those listed, we choose to report analysis based on the odds offered by Bet365, a firm with high recognition in all sectors, not just tennis. However, we checked and found that findings were robust regardless of whichever bookmaker's set of odds was selected for the study. At the end of the paper, we explore the extent to which betting returns can be improved by wagering at the best odds available on a particular bet across the eight firms. For all firms, odds recorded on the website are generally closing odds. As Woodland and Woodland (1999) pointed out, closing odds have the advantage of having been modified by betting volumes. They therefore represent more genuine market prices, influenced by bettor preferences, than odds available earlier in the betting period. It should be noted that odds movements tend to be smaller than in horse-race markets. For example, if decimal odds (defined below) were to move from 1.50 to 1.55 or 1.45, this would be regarded as normal. However, a move from 1.50 to 1.75 or 1.30 would be wholly exceptional and would indicate significant news such as of a player injury.

For Bet365, odds were available for 8233 matches.[7] Odds in the archive are quoted in decimal format. If odds are d, a successful bet with a stake of one unit would give a payout to the bettor of d, yielding him a profit of $(d - 1)$. For our purposes, we converted the decimal odds to the alternative format of *probability-odds*, which would in this case be $(1/d)$. For example, if decimal odds are quoted as 3.50, a successful one unit bet would deliver a payout of 3.50 units and a profit to the bettor of 2.50 units. Corresponding probability-odds are 0.286.

It is, of course, the case that the sum of the probability-odds in a match exceeds 1 to allow for bookmaker profit. The margin by which it exceeds 1 is known as the "over-round" and is a measure of transactions costs for bettors. In our data set, the mean over-round is 0.075, which implies that, if the bookmaking firm held the same liabilities on each outcome, it would earn a return of $0.075/1.075$, or 7.0% of stakes. This is lower than in horse-race-betting markets worldwide. There is very little variation in over-round across the matches in our data set; the standard deviation is 0.008.

4. Odds and Return

In this and in the next section, we explore the relationship between odds and return. The general pattern is summarized in Table 1, which shows the returns to placing a unit stake on every bet available within each of 10 odds ranges. There is a close to monotonically increasing relationship between probability-odds and return. In the 56 cases where the odds indicated that the player was most unlikely to win (odds 0 to 0.1), not a single bet was successful, and the return was -1 to a portfolio of all such wagers. As odds increase beyond 0.8, a negative return was still incurred, but it became less than 0.02 of stakes.

The information in Table 1 is useful for indicating the presence of positive longshot bias in this betting market. However, it is inadequate for estimating the precise relationship between odds and return to the extent that groupings of odds have to be imposed on the data, and this process of

Table 1. Mean returns by odds range for Bet365

Odds range	Number of bets	Mean return	SD	t statistic
0.0–0.1	56	−1.000	0.000	
0.1–0.2	834	−0.326	2.034	−4.622
0.2–0.3	1674	−0.257	1.567	−6.706
0.3–0.4	2060	−0.160	1.306	−5.557
0.4–0.5	2506	−0.065	1.124	−2.889
0.5–0.6	2601	−0.072	0.911	−4.038
0.6–0.7	2658	−0.061	0.750	−4.197
0.7–0.8	1850	−0.025	0.603	−1.814
0.8–0.9	1539	−0.018	0.451	−1.555
0.9–1.0	682	−0.017	0.296	−1.482
Total	16 460	−0.102	1.077	−12.194

aggregation risks introducing measurement bias. We therefore estimate the relationship between odds and yield using data on individual bets.

Analysis here encounters the problem that within the set of 16 460 bets observed, returns are not independent as between each of the two bets available in a single match. If one player wins, the other loses. If the return on one bet in a match is positive, that on the other is necessarily −1. Therefore, regression of yield on odds would generate downwardly biased standard errors. We resolve the issue by randomly selecting one player in each match so that there is only one observed bet for each event. This leaves us still with a sample size in excess of 8000.

If the tennis-betting market were not subject to longshot bias, bets at any odds would be associated with the same (negative) expected return. It would not be possible to improve returns by adopting a strategy of betting, in particular odds ranges (and therefore there would be no violation of strong efficiency, as defined by Thaler and Ziemba (1988), which rules out that any betting plan is superior to any other). The hypothesis that there is no bias, and there is strong efficiency, implies therefore that the relationship between odds and yield is linear with a slope of zero. Because linearity is part of the null hypothesis, we first fitted a linear model, $return = \alpha + \beta * odds$, to the data. In the absence of longshot bias, the constant term would be negative (reflecting bookmaker win) and the coefficient on probability-odds would be zero. This is akin to the widely employed test for efficiency in betting markets organized on a spread basis, where the equation estimated is: $actual\ points\ difference = \alpha + \beta * spread$. However, in tests for spread markets, efficiency implies $\alpha = 0$ and $\beta = 1$, rather than $\alpha < 0$ and $\beta = 0$ as in the present setting of an odds-based market.

In our exercise for the odds-based, tennis-betting market, the intercept term was negative as expected (bettors, on an average, lose money), but the estimated coefficient on probability-odds exceeded zero (t statistic: 7.14), indicating superior returns to backing more heavily favoured players, i.e. positive longshot bias. Absence of bias is therefore rejected even if, in statistical analysis, one imposes the restriction of linearity implicit in the null hypothesis (imposing linearity, the relationship between probability-odds and return was estimated by the equation, return $= -0.280 + 0.351 *$ probability-odds; the standard errors were 0.032 and 0.057, respectively). Given bias as demonstrated, the case for imposing linearity collapses, and we therefore base subsequent analysis on less-restrictive specification.

In adopting a more flexible approach than imposing linearity, we also recognized that the estimation exercise can, in fact, be reduced to relating the probability of losing/winning to quoted probability-odds. This is because the relationship between yield and probability-odds is arithmetic: conditional on a bet being successful, the return is determined only by the odds.

Let

$$X_i = 0, \quad \text{if bet } i \text{ loses,}$$
$$X_i = 1, \quad \text{if bet } i \text{ wins.}$$

Now, let r_i represent the return to a unit bet. For any given probability-odds, p_i, r_i is deterministic in the sense that if X_i is known, r_i is known:

$$r_i\big|_{X_i=1} = \frac{1}{p_i} - 1, \qquad r_i\big|_{X_i=0} = -1. \tag{1}$$

The expected value of the return, r_i, is then given by:

$$E\{r_i\} = \theta_i\left(\frac{1}{p_i} - 1\right) + (1-\theta_i)(-1) = \frac{\theta_i}{p_i} - 1, \tag{2}$$

where θ_i is $\Pr(X = 1)$. We use an estimate of θ_i, $\hat{\theta}_i$, to be obtained from the estimation of the relationship between the true probability of winning and probability-odds. Note that we do not scale the probability in each match to sum to 1. The purpose of the analysis is to test whether it is possible, merely from a study of past prices, to formulate trading strategies that earn superior returns in the future. This would be a violation of strong efficiency according to the definition of Thaler and Ziemba (1988). It would be a violation of even weak efficiency, according to their terminology, if a strategy yielded positive returns. Examination of efficiency does not require the scaling of odds found in some betting market studies.

The first step is to estimate the relationship between true win probability and bookmaker probability-odds. To do this, we fit a semi-parametric regression model. This provides flexibility, not present in previous studies, that implicitly or explicitly impose functional forms on the relationship. It means that we do not make any assumptions regarding the shape of the link function used, as would be the case if a logit or probit link function were adopted. Indeed, we have found that the familiar "S-shaped" curves of the logit and probit link functions are not appropriate for the data. We fit a semi-parametric regression model given by:

$$\theta = m(p) + \varepsilon, \tag{3}$$

where θ is, as before, the true probability of a win and p the covariate, i.e. the bookmaker probability-odds. In the semi-parametric regression, the function m effectively weights the average values of the response variable locally. Just how locally this averaging occurs is governed by a smoothing parameter, h. Since we are dealing with a binary variable, the model is adjusted and we transform θ using the logit transformation, and hence $m(p)$ is a function of the form:

$$\hat{m}(p) = \frac{\exp(\hat{\alpha} + \hat{\beta}p)}{1 + \exp(\hat{\alpha} + \hat{\beta}p)}, \tag{4}$$

so as to ensure that the estimated probabilities lie between 0 and 1. $\hat{m}(p)$ is estimated using a *local likelihood* procedure. At each value of the predictor variable p, $\hat{\alpha}$ and $\hat{\beta}$ are calculated, resulting in the estimated function $\hat{m}(p)$. From this function, we can evaluate $\hat{\theta}$. This semi-parametric model and its estimation using a local likelihood approach are described in more detail by Bowman and Azzalini (1997). The model was fitted and estimated in R (R Development Core Team, 2005). The smoothing parameter, h, has been chosen so as to minimize the cross-validation function, the optimal value of h being 0.04.

As in the linear estimation, we randomly selected one player from each match and then fitted a semi-parametric regression model, relating the probability of winning a bet to the quoted probability-odds, to the data on over 8000 wagers. The result is shown in Figure 1 where the dashed lines represent the upper and lower 95% confidence interval for the fitted regression line.

It is however true that a different relationship will be fitted according to which set of players has been included in the sample following a random selection of one bet per match. To satisfy ourselves as to the robustness of the result, we repeated the exercise in 10 further trials. In each case, the outcome was very similar to that reported here, with identical implications as described below.

The relationship between expected return and probability-odds, derived from the semi-parametric regression according to equation (2), is illustrated by Figure 2, which also shows the corresponding 95% confidence band for the particular randomly selected sample of bets. According to Cain and Peel (2004), a utility function such as that associated with Friedman and Savage (1948), i.e. having both risk-averse and risk-loving sections, would be expected to generate a relationship with a central maximum. However, although the relationship between odds and yield, as shown in Figure 2, is not perfectly monotonic, there is a clear and consistent positive

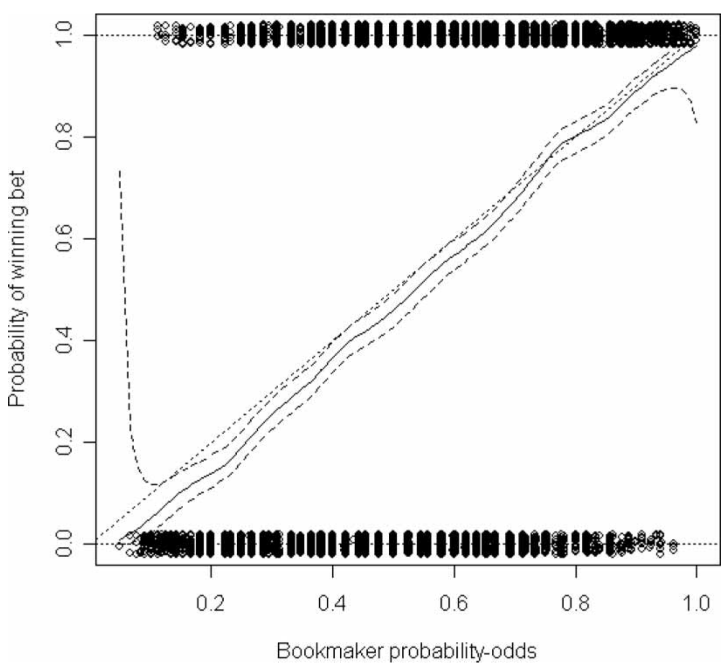

Figure 1. Fitted semi-parametric model. *Note*: The observations are represented by circles that have been spread a little so as to help the reader view the distribution of points at zero and one. A 45° line has been added to facilitate interpretation

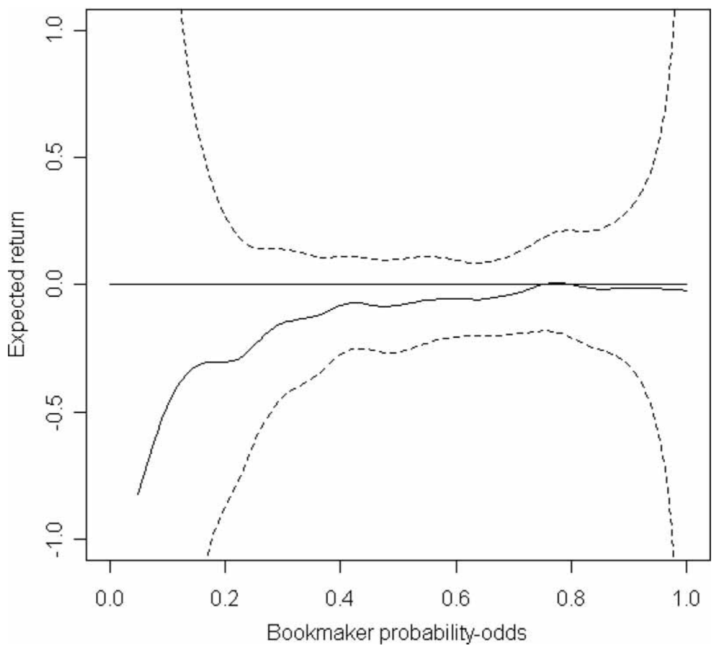

Figure 2. Expected yield and confidence intervals from the fitted semi-parametric model

longshot bias throughout the odds range, with no central maximum. Heavy losses are associated with bets on rank outsiders, whereas expected yield is close to zero for the strongest favourites.

The pattern of returns was stable over time so that bettors could have used their observations from one period to improve their returns in the next period. We illustrate this in Table 2, which shows yields to betting in different odds ranges in the sub-periods 2002–2003 and 2004–2005. For example, the smallest losses in 2002–2003 were for bettors focussing on the 0.7–1.0 odds range, and in 2004–2005, these odds groupings again yielded the best returns.

Table 2. Returns by odds range for Bet365, 2002–2003 and 2004–2005

	2002–2003				2004–2005			
Odds range	Number of bets	Mean return	SD	*t* statistic	Number of bets	Mean return	SD	*t* statistic
0.0–0.1	4	−1.000	0.000		52	−1.000	0.000	
0.1–0.2	314	−0.272	2.124	−2.271	520	−0.358	1.979	−4.122
0.2–0.3	718	−0.295	1.526	−5.185	956	−0.228	1.597	−4.412
0.3–0.4	954	−0.129	1.321	−3.028	1106	−0.186	1.292	−4.787
0.4–0.5	1153	−0.089	1.122	−2.702	1353	−0.044	1.126	−1.439
0.5–0.6	1286	−0.068	0.906	−2.696	1315	−0.076	0.916	−3.010
0.6–0.7	1178	−0.052	0.745	−2.397	1480	−0.068	0.754	−3.480
0.7–0.8	839	−0.032	0.609	−1.529	1011	−0.020	0.599	−1.055
0.8–0.9	656	−0.006	0.441	−0.337	883	−0.027	0.458	−1.741
0.9–1.0	234	−0.030	0.329	−1.415	448	−0.010	0.276	−0.736
Total	7336	−0.097	1.071	−7.788	9124	−0.106	1.082	−9.391

Our data pertain to a single sport with an even wider odds range than that available to Cain and Peel; considerable numbers of wagers were available even towards the extremities (for example, 682 had odds in the 0.9–1 range). The pattern that is shown has similarities with those found by Cain and Peel to the extent that the relationship between yield and odds is particularly steep in the long odds range (where skewness in the distribution of returns changes rapidly) and relatively flat around the odds of evens. However, there is no evidence of a turning point, and the findings are therefore consistent with bettors behaving as if risk-loving over the whole range of possible payoffs. As risk becomes very low, expected loss virtually disappears, as one would anticipate if predicting on the basis of simply assuming a risk-loving utility function.

Our findings (which appear to imply that tennis bettors love both risk and skewness) cast doubt on whether bettors in all markets exhibit similar risk preferences capable of being represented by a utility function shaped as in Friedman and Savage (1948). They are unsupportive of the notion that focus on such a utility function can reconcile the differences in the direction of longshot bias found in studies of wagering markets on different activities.

5. Microanalysis

It could be argued, if only speculatively, that the aggregation across the whole data set is in fact inappropriate. Because of the way tennis tournaments are structured, matches where long odds are available may take place on different days from those where odds are compressed. Hence, in one sense, bets with very short or very long odds may not be made in the same market as those with odds close to evens. We therefore review evidence on the return–risk relationships revealed by considering, as Woodland and Woodland (1999) did for baseball, groups of observations defined by how short odds the favourite was in an individual match.

The theoretical model in Woodland and Woodland is essentially similar to that underpinning Cain and Peel's search for a turning point in the overall relationship between odds and yield. It is that utility functions are characterized by a distaste for risk and a preference for skewness. Thus, a positive return–risk relationship should be observed in a market where neither outcome was priced very far from the evens. This is because skewness, in such cases, is not a major issue since it is quite close to zero and not very different between the alternative bets on the two teams. In the betting markets for such games, negative longshot bias should be observed (such that the yield on underdog bets exceeds that on favourite bets). However, as one moves towards considering games with a very strong favourite and a very weak underdog, bettors will be attracted by the increasingly positive skewness in the distribution of returns from the long-odds underdog proposition. The degree of negative longshot bias should therefore fall as increasingly strong favourites are considered and may eventually be reversed to become positive bias at sufficiently extreme odds. Woodland and Woodland's evidence was broadly consistent with these propositions, though the lack of very strong favourites precluded complete testing (which was possibly why no reversal of bias was in fact observed): only 3% of games offered odds over two. An advantage of tennis is that over half of the matches featured such a bet so that there are very adequate numbers of observations in odds ranges where there is large difference in skewness of returns as between favourite and underdog. If bettors are risk-averse but skewness-loving, this should be manifest in a tendency for longshot bias to be reversed when considering bets within matches where each player has either very short or very long odds. This is what we test in this section.

Table 3 shows the return achieved by betting one unit on all wagers available in odds classes defined by favourite odds.[8] In order to examine more closely the relationships between odds ranges, return, variance, and skewness, we also include columns for average variance and average

Table 3. Returns to favourites and underdogs by favourite odds range

	Favourite						Underdog				
Favourite odds range	Number of matches	Mean return	SD	*t* statistic	Average variance	Average skewness	Mean return	SD	*t* statistic	Average variance	Average skewness
0.55–0.60	1173	−0.053	0.857	−2.111	0.735	−0.185	−0.097	1.002	−3.300	1.004	0.199
0.60–0.65	1167	−0.051	0.787	−2.234	0.620	−0.354	−0.099	1.088	−3.107	1.187	0.353
0.65–0.70	1491	−0.069	0.720	−3.677	0.514	−0.568	−0.069	1.208	−2.209	1.450	0.566
0.70–0.75	1005	−0.051	0.642	−2.494	0.399	−0.891	−0.110	1.320	−2.636	1.701	0.860
0.75–0.80	1189	0.002	0.539	0.119	0.300	−1.256	−0.237	1.427	−5.721	2.076	1.269
0.80–0.85	448	−0.021	0.478	−0.941	0.222	−1.622	−0.229	1.585	−3.061	2.484	1.614
0.85–0.90	747	−0.022	0.407	−1.483	0.159	−2.079	−0.280	1.741	−4.395	2.982	2.019
0.90–0.95	496	−0.022	0.325	−1.479	0.102	−2.785	−0.314	2.087	−3.346	4.068	2.783
0.95–1.0	186	−0.004	0.198	−0.276	0.052	−4.201	−0.664	1.704	−5.314	4.263	4.365
Total	7902	−0.039	0.656	−5.318	0.429	−1.051	−0.166	1.381	−10.703	1.905	1.049

skewness for bets in each odds class. The latter two quantities are calculated as follows. If, as before, we let the true probability of winning bet i, be θ_i, estimated from our semi-parametric regression as $\hat{\theta}_i$, and the bookmaker probability-odds of bet i, be p_i, then for a one unit bet the expected return of bet i is given by equation (2). The variance of the return on each individual bet, i, is given by

$$\mathrm{Var}(r_i) = \frac{\hat{\theta}_i}{p_i^2}(1 - \hat{\theta}_i) = \sigma_{r_i}^2 \tag{5}$$

and the skewness of each return (the standardized third moment) is thus given by

$$\mathrm{Skew}(r_i) = \frac{\hat{\theta}_i(2\hat{\theta}_i - 1)(\hat{\theta}_i - 1)}{p_i^3 \sigma_{r_i}^{3/2}}. \tag{6}$$

Table 2 shows the average variance and average skewness of bets for favourites and underdogs in odds classes defined by favourite odds.

Consistent with Woodland and Woodland, underdogs tend to become increasingly unfair betting propositions, as more extreme odds are considered. However, in no odds range, even where the two players' odds are very close, does the expected return on the underdog exceed that on the favourite. Indeed, in every group except one, where returns are equal, the data indicate positive longshot bias. This suggests that risk-loving rather than risk-averse bettors dominate the market. Again the pattern is consistent with typical bettors being risk- and skewness-loving (i.e. they are risk-loving and become more risk-loving as they consider bets with higher and higher payoffs).

6. Betting in High-Profile Tournaments

We have focussed on implications of evidence from tennis for theories of longshot bias based on hypotheses concerning the utility functions of representative bettors. However, a separate strand in the literature (initiated by Shin (1992)) represents the positive longshot bias as a defensive reaction by bookmakers to the possible holding of private information by privileged bettors. In support of this approach, Vaughan Williams and Paton (1998) claimed that there were differences in whether a positive longshot bias existed according to how high profile the (racing) event was. Here, therefore, we consider separately patterns of returns as between Grand Slam and non-Grand Slam tournaments. Grand Slams are the highest profile events. Upsets should be fewer for two reasons. First, matches are longer, played over five rather than three sets, making it less likely that outcomes will be determined by random events such as poor umpire calls. Second, prize money and prestige is much higher than in other tournaments. With incentives to effort greater, there is less likelihood that a star player will be beaten because he is, for example, treating the event as a practice or underperforming because he does not want to risk aggravating an injury. Therefore, relevant private information concerning the event becomes much less likely to be a factor in the betting market. One might then predict that the positive longshot bias we have found for tennis might be mitigated in the case of Grand Slam tournaments.

In fact, Table 4, which displays mean returns by odds ranges, shows no difference between the patterns for non-Grand Slam and Grand Slam sub-samples. There is no evidence here that there is a role for inside information in accounting for the positive longshot bias we have documented.

Table 4. Mean returns by odds range, Grandslams only

Odds range	Number of bets	Mean return	SD	*t* statistic
0.0–0.1	36	−1.000	0.000	
0.1–0.2	299	−0.351	2.053	−2.958
0.2–0.3	365	−0.410	1.429	−5.478
0.3–0.4	328	−0.073	1.349	−0.984
0.4–0.5	348	−0.081	1.125	−1.345
0.5–0.6	335	−0.066	0.909	−1.320
0.6–0.7	364	−0.066	0.749	−1.691
0.7–0.8	337	−0.024	0.600	−0.721
0.8–0.9	345	0.018	0.412	0.805
0.9–1.0	265	−0.022	0.296	−1.195
Total	3022	−0.131	1.122	−6.426

7. Betting at "Best Odds"

Odds in tennis appear, then, to be biased such that a strategy of accepting only shorter odds wagers yields a superior mean return. But in no odds group was there evidence that expected return was consistently positive. This appears to imply that the market is weak, but not strong, efficient in the Thaler–Ziemba sense. However, returns were examined using only odds offered by Bet365. In the analysis of betting markets, it has sometimes been claimed that a profit could have been made by identifying a bias and combining this knowledge with a strategy of shopping around, so as to place bets at different bookmakers according to which had the best price quote for that particular bet (Simmons *et al.*, 2003). The eight bookmakers in our sample all enjoy a strong reputation and operate internet accounts. Once accounts had been established, a serious bettor would therefore have faced low transactions costs: current odds could be compared on a single site and the wager then placed where prospectively there was the highest return. By definition, given heterogeneity in odds, this must improve bettor return. The interest is whether a strategy with positive expected returns emerges.

Table 5. Best odds return by odds range defined by Bet365

	2002–2003				2004–2005			
Odds range	Number of bets	Mean	SD	*t* statistic	Number of bets	Mean	SD	*t* statistic
0.0–0.1	4	−1.000	0.000		52	−1.000	0.000	
0.1–0.2	314	−0.249	2.191	−2.010	520	−0.216	2.490	−1.977
0.2–0.3	718	−0.265	1.597	−4.442	956	−0.145	1.785	−2.518
0.3–0.4	954	−0.092	1.381	−2.050	1106	−0.116	1.408	−2.743
0.4–0.5	1153	−0.058	1.162	−1.695	1353	0.023	1.210	0.700
0.5–0.6	1286	−0.037	0.938	−1.425	1315	−0.018	0.976	−0.684
0.6–0.7	1178	−0.025	0.767	−1.125	1480	−0.012	0.806	−0.555
0.7–0.8	839	−0.010	0.623	−0.472	1011	0.029	0.633	1.474
0.8–0.9	656	0.022	0.455	1.255	883	0.024	0.490	1.471
0.9–1.0	234	−0.011	0.337	−0.493	448	0.058	0.723	1.695
Total	7336	−0.068	1.112	−5.254	9124	−0.040	1.221	−3.129

Table 5 shows returns to placing unit bets on all wagers in a given Bet365 odds range but using the supplier offering the most generous price. The sample is for matches where Bet365 odds were available: the best price may not necessarily have been the best from eight bookmakers since there were missing odds in some cases (we used whatever odds were available). There is some evidence of profitable trading opportunities that repeated from one sub-period to the next, i.e. of a violation of weak efficiency. In 2003–2004, bets in the 0.8–0.9 odds range yielded a positive return in excess of 2%. A bettor noting this could have accepted all such bets in 2004–2005 and again enjoyed a positive return of more than 2%. However, in neither case was the return statistically significant except at generous significance levels.

8. Closing Remarks

We have employed a new econometric approach to the estimation of the relationship between odds and return in a betting market, using it to test for longshot bias in the context of wagering on men's professional tennis. The tennis-betting market we study appears to be conducted between well-informed traders, and transactions costs are lower than in horse racing. The market we analyse therefore lacks some features claimed by previous authors to play a significant role in generating a bias. We find nevertheless that there is a positive longshot bias in the tennis-betting market.

Vaughan Williams and Paton (1998) offered a general theory of bias in betting markets, designed to be consistent with the findings in American baseball and different categories of horse racing. The danger in basing explanations of longshot bias on evidence from a limited range of betting activities is underlined by our findings for tennis. First, Vaughan Williams and Paton predicted that the positive bias would be eroded as the potential for insiders to hold the relevant information diminished; but we found a positive bias even in Grand Slam events where effort is likely to be maximized, and therefore inside knowledge concerning a player's commitment to the upcoming match likely to be less relevant. Second, they predicted that bias would be lessened or reversed as the proportion of informed bettors in a market increased because informed bettors preferred backing favourites; but, though tennis betting is a specialist rather than a mass market activity, strong positive bias remains in the odds. Third, they related the negative bias in baseball odds to the very low transactions costs compared with horse racing; but tennis betting, with its positive bias, is also cheaper than racing in terms of bookmaker over-round, so low transactions costs in themselves do not appear to lead to diminished or zero or reverse bias.

Some prior literature finds a negative longshot bias in particular markets but this has always been in settings where only a narrow range of odds was available. Some authors have attempted to reconcile the findings of positive and negative bias in different areas of betting by pointing to the fact that different odds ranges comprise the sample in different studies. Across the whole possible odds range, there might be positive or negative bias according to where in the range one looks. Contrasting evidence in different sectors might then be explained not by the classic analysis of market inefficiency (many noise traders, high transaction costs) but by reference to bettors' risk preferences, in particular the possibility that they are similar to those exhibited by a Friedman–Savage utility function (which has been redefined here to relate to the size of the potential amount won for a unit stake rather than total wealth). If it had this shape for the representative bettor, it would be capable of generating either a negative or positive longshot bias according to where in the odds range one looked. This was the point emphasized in Cain and Peel (2004).

The way tennis is organized offers an opportunity to study a complete range of odds on players winning matches in the context of a single sport. Our evidence pointed to positive longshot bias with no evidence of negative bias in any part of the odds range. This calls into question whether a focus on a Friedman–Savage-shaped utility function is the most appropriate approach in trying to explain why a longshot bias is positive in some markets and negative in others. The evidence for tennis is consistent with a simple assumption that bettors have risk-loving utility functions rather than a Friedman–Savage-shaped function (it would also be consistent with the alternative hypothesis that bettors prefer longshots because there is greater expected utility when bragging rights could be won).

Risk-loving utility functions constitute the simplest story to account for why people bet at all. But the story does not explain why a negative bias occurs in certain team sports betting markets. One possibility is that in the absence of genuine longshots in such betting markets, bragging rights accrue to those with the highest proportion of winning bets. In this case, bettors would prefer to back favourites and therefore accept a lower rate of return than that on underdogs, hence a negative bias. Another area for investigation might be the extent to which odds in such markets are distorted by sentiment.

The longshot anomaly is unlikely to be explained fully from analysing only one betting sector. Resolution will be more likely if additional sectors are examined for positive or negative bias since a successful theory must account for heterogeneity of findings across racing and other sports. While evidence has built on team sports, ours is the first study employing a large data set to examine a betting market for an individual sport. It proved to be more like the market for horse racing than those for team sports, with a positive relationship between odds and return throughout an odds range that extended virtually from 0 to 1. This casts doubt on whether an assumption that bettor utility functions display convex and concave segments can account for heterogeneity in findings on longshot bias across betting sectors. It will be of interest whether the finding is repeated in the future for other individual sports. If it is, the literature might wish to explore alternative explanations than risk preference in accounting for the existence of positive and negative longshot bias in different betting markets.

Acknowledgements

We are grateful to our student, Graham Pigott, who first drew our attention to the potential of the betting market on men's tennis to inform the debate on longshot bias in wagering markets, for his stimulating ideas and for assembling, and carrying out very valuable preliminary analysis of, the data. The odds data on British jumps racing to which reference was made in the commentary were collected by Kristian Strangeway.

Notes

1. See, for example, Terrell and Farmer (1996), Hurley and McDonough (1995), and Vaughan Williams and Paton (1998).
2. Longshot bias has not been subject to the same analysis in the betting markets on the other principal American team sports, football, and basketball, since, in each of these cases, the primary offering of bookmakers is handicap betting, i.e. bettors are asked to judge whether a team will "beat the spread" and bets on all teams carry the same odds, so there is, in this case, no range of odds across which returns could vary.
3. The standard deviation (across clubs) of the proportion of matches won in the season (with draws counted as half a win) is an often-used measure of the degree of balance in a sports league. In the S-League, it was 0.189 in 1999 and 0.184 in 2000. By way of comparison, the equivalent measure for the English Premier League in 1999–2000 was 0.141 and that for the supposedly uncompetitive Scottish Premier League was 0.148.
4. Friedman and Savage considered a utility of wealth function. Their function had risk-averse and risk-loving segments, and this could "explain" why many individuals buy both insurance and lottery tickets. The focus on the marginal utility of wealth was plausible, given that a house fire or a jackpot win would be associated with a move to a very different level of wealth for most individuals. However, it is less clear that the analysis illuminates betting behaviour. Most bets

are for a small stake and relate to a relatively small possible payoff. For most bettors, losing or winning would make little difference to their level of wealth. The analysis of the decision to bet or not would then focus on a very narrow segment of the utility of wealth function, which therefore may be taken as imperceptibly different from linear (i.e. risk-neutral). Given that track or bookmaker commissions are positive, the bet would not be made. That people do bet we therefore view as revealing a recreational value to betting. When we refer to the utility function, we have in mind the utility attached to winning different amounts from a unit bet. This *may* have the shape of the Friedman–Savage function but does not relate to total wealth.

[5] Access to the site requires the payment of a subscription.

[6] The eight were: Bet365, Bet & Win, Centrebet, Expekt, Gamebookers, Interwetten, Pinnacles, and Sportingbet.

[7] There were six matches in which betting was not available on the "favourite".

[8] We excluded the 325 matches with no favourite (in each of these matches, odds on either player were quoted as 1.833 (11/12 in UK format)) and six matches where no betting was available on the favourite.

References

Avery, C. and Chevalier, J. (1999) Identifying investor sentiment from price paths: The case of football betting, *Journal of Business*, 72, pp. 493–521.

Bowman, A. W. and Azzalini, A. (1997) *Applied Smoothing Techniques for Data Analysis: The Kernel Approach with S-Plus Illustrations* (Oxford: Oxford University Press).

Busche, K. (1994) Efficient market results in an Asian setting, in: D. Hausche, S. Y. Lo and W. T. Ziemba (Eds) *Efficiency in Racetrack Betting Markets*, pp. 615–616 (London: Academic Press).

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

Busche, K. and Walls, W. D. (2000) Decision costs and betting market efficiency, *Rationality and Society*, 12, pp. 477–492.

Cain, M. and Peel, D. (2004) The utility of gambling and the favourite-longshot bias, *European Journal of Finance*, 10, pp. 379–390.

Cain, M., Law, D. and Peel, D. (1996) Insider trading in the greyhound betting market, Discussion Paper, University of Salford.

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, pp. 25–36.

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

Forrest, D. and Simmons, R. (2001) Efficiency in an Asian sports betting market, Working Paper, University of Salford.

Forrest, D. and Simmons, R. (2002) Outcome uncertainty and attendance demand: The case of English soccer, *Journal of the Royal Statistical Society (Series D) The Statistician*, 51, pp. 229–241.

Forrest, D., Goddard, J. and Simmons, R. (2005) Odds-setters as forecasters: The case of English football, *International Journal of Forecasting*, 21, pp. 551–564.

Friedman, M. and Savage, L. (1948) The utility analysis of choices involving risk, *Journal of Political Economy*, 56, pp. 279–304.

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

Griffiths, M. (1949) Odds adjustment by American horse race bettors, *American Journal of Psychology*, 62, pp. 290–294.

Hurley, W. and McDonough, L. (1995) A note on the Hayek hypothesis and the favourite-longshot bias in pari mutuel betting, *American Economic Review*, 85, pp. 949–955.

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

R Development Core Team (2005) *R: A Language and Environment for Statistical Computing* (Vienna, Austria: R Foundation for Statistical Computing). ISBN 3-900051-07-0, available at http://www.R-project.org.

Sauer, R. (1998) The economics of wagering markets, *Journal of Economic Literature*, 36, pp. 2021–2064.

Shin, H. (1992) Prices of state contingent claims with insider traders, and favourite-longshot bias, *Economic Journal*, 102, pp. 426–435.

Shmanske, S. (2005) Odds-setting efficiency in gambling markets: Evidence from the PGA Tour, *Journal of Economics and Finance*, 29, pp. 391–402.

Simmons, R., Forrest, D. and Curran, A. (2003) Efficiency in the handicap and index betting markets for English rugby league, in: L. Vaughan Williams (Ed.) *The Economics of Betting*, pp. 114–134 (London: Routledge).

Terrell, D. and Farmer, A. (1996) Optimal betting and efficiency in parimutuel betting markets, *Economic Journal*, 106, pp. 846–868.

Thaler, R. and Ziemba, W. (1988) Anomalies – Pari mutuel betting markets: racetracks and lotteries, *Journal of Economic Perspectives*, 2, pp. 161–174.

Vaughan Williams, L. (1999) Information efficiency in betting markets: A survey, *Bulletin of Economic Research*, 51, pp. 1–30.

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

Woodland, L. and Woodland, B. (1994) Market efficiency and the favourite-longshot bias: The baseball betting market, *Journal of Finance*, 49, pp. 269–279.

Woodland, B. and Woodland, L. (1999) Expected utility, skewness and the baseball betting market, *Applied Economics*, 31, pp. 337–345.

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

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