
The Cost of Insider Trading: Evidence from Defined Markets

Leighton Vaughan Williams¹

Abstract

In this paper, two complementary markets for bets on UK horse races are examined for evidence of information inefficiency in the form of unexploited price differentials between these markets for the same product. It is shown that bettors could on average improve their expected return by shifting their betting patterns in a defined way. Evidence is adduced which supports the hypothesis that the apparent inefficiency is a consequence of perceived insider activity. In particular, in situations where there are likely to be limited opportunities for insider trading, the 'inefficiency' disappears. Although the evidence produced here applies to a relatively unsophisticated market, it poses legitimate questions about the operation of more complex financial markets.

1. Introduction

A standard definition of information efficiency is provided in Fama (1991, p.1575), that '... prices fully reflect all available information.'

In conventional financial markets, the existence of information efficiency as defined is closely related to the absence of opportunities for traders to earn abnormal returns. In betting markets, however, the odds offered to bettors are designed so that the odds-setters earn a profit overall, i.e. odds-takers secure a loss overall. Information

efficiency in this context is, therefore, the idea that no one better is able to earn above-average returns or that no given trading strategy can out-perform another, except by chance.

In this paper, a particular test of information efficiency in racetrack betting markets is proposed and employed, i.e. whether bettors can expect to earn systematically different returns by placing different types of bet about events with identical probabilities of success. The test is based on the premise that in an informationally efficient market, in the sense here employed, these returns should tend to converge, at least net of differential costs of implementation. Otherwise, bettors should be able to adjust their betting patterns so as to increase their expected return over time. It is proposed in this paper to ascertain whether there is evidence of such an inefficiency, and if so the reason for its existence and persistence.

A racetrack betting market, it should be noted, is a relatively simple, unsophisticated form of financial market - characterized by a sequence of relatively short-lived and usually independent markets, and a well-defined end-point at which the value of the outcome becomes certain. This defined termination point is indeed particularly appealing in that it helps avoid many of the analytical problems associated with indefinite future outcomes. This contrasts with more sophisticated

financial markets, where the value of an asset in the present is dependent both on the present value of future cash flows and on the uncertain price at which it can be sold at some future point in time.

In seeking to draw more general conclusions for financial markets from a specific study of betting markets it is, however, very important to be aware of the specific nature and context of the market/s being investigated. This study in any case offers a valuable insight into the operation of bettors and betting markets in their own right.

2. Tests of market inefficiency

An early test for the existence of differential returns to different types of bet about events with comparable probabilities of success was formulated and investigated by Ali (1979).

Ali (1979) tested the hypothesis of differential returns to two forms of 'exotic bet', known as the 'daily double' and the 'parlay.' In a 'daily double' the better selects the winner of two consecutive races before the first race is run, securing a return only if both horses win. In the 'parlay', the better selects a series of horses, betting the total proceeds of each win on the next, until a pre-determined number of wins, or until one loses. The usefulness of comparing these two types of bet is that one can be constructed from the other. Ali's test of market efficiency is based on the idea that in an efficient market bets will be valued according to their probability distributions alone, and so the return to a daily double will be the same as that to the corresponding parlay. Using data from 34 racetracks in the U.S. and Canada between September and December 1975, his results were consistent with the hypothesis that both bets were identically priced and therefore, that the efficient market hypothesis could not be rejected. Although later work by Asch and Quandt (1987), using U.S. data, appeared to

contradict this conclusion, Dolbear (1993) cast doubt on their results. For Hong Kong data, however, Lo and Busche (1994) found clear evidence of differential returns, and therefore of inefficiency.

A more direct test of information inefficiency is to identify cases of unexploited arbitrage opportunities. Attempts to apply this sort of approach have particular relevance in the U.K., where two co-existing markets are available to bettors - the 'Tote' and Bookmakers, each offering independent odds about identical events. The Tote (Horserace Totalisator Board) offers a system (sometimes referred to as a 'parimutuel' system) in which bettors compete for shares of a fixed proportion of the pool of all bets, the size of the fixed proportion being determined in advance outside of the control of the bettors. The proportion of the pool available for distribution is always less than one, and hence the aggregate return to bettors must always be less than the aggregate stakes of bettors, i.e. the expected average profit is negative. The odds set in the pool are determined by the relative weight of money placed on each outcome, and are therefore demand-determined. Bookmakers on the other hand set odds as active suppliers, although these odds are likely to be influenced by the activities of the demand side of the market, i.e. of bettors. Unlike in the Tote system, the proportion of aggregate stakes returned to bettors is not pre-determined and fixed. The so-called 'Starting Price' is the price (odds) at which a sizeable bet could be placed with bookmakers on the course at the start of the race. This is determined by independent assessors at the racetrack. Gabriel and Marsden (1990, 1991) compared the returns at the odds implied ex-post in the starting price and tote returns, observing that, 'Since the differing bets are two options for purchasing exactly the same

item (a bet to win on a specific horse), we would expect the odds to converge.' (p.877).

In fact, an examination of the difference in the mean tote and starting price payouts revealed a significantly greater expected return to bets on the tote. Blackburn and Peirson (1995), Cain, Law and Peel (1997), and Vaughan Williams and Paton (1997a) repeat this analysis for different data sets, and in each case confirm significant differences in the returns to each type of bet, in favour of tote bets, at least at higher odds.

This paper employs a data set composed of races run prior to the introduction of direct access terminals to tote pools in a wide range of off-course bookmakers ('Tote Direct'), in order to test for the existence of a differential between tote and starting price returns. In this set of races, there is clear scope for the existence of unexploited arbitrage opportunities, although not exclusively so. Moreover, one might expect that the arbitrage opportunities would persist in any subset of these races. To test this, the set of all races is divided into four subsets, based on the likely potential for insider activity in these races.

3. Insiders and information efficiency

The motivation of this analysis is to test a hypothesis that insider trading can produce an information inefficiency in complementary betting markets, in the form of a differential expected return to outcomes characterized by identical probabilities of success. In particular, does the existence of a clear potential for insider activity in certain types of race lead to a differential expected return to bets in Tote and bookmaker-based markets, and does the absence of this potential eliminate this differential.

The reasoning behind the hypothesis that insider activity may produce this form of

inefficiency is that bookmakers may be reluctant to set large odds, even on horses whose estimated objective probability of winning is low, because of an adverse selection problem facing them in the context of insiders who may possess superior information to the bookmakers, or who can even (in the stylized models of Shin, 1991, 1992, 1993) 'observe' the outcome with certainty. In such an environment, bookmakers may 'squeeze' prices, particularly about a 'longshot', even when no bets have been placed on it, because the 'insider' may pounce at any time. Because there is often a 'limit to the size of the wager accepted at the quoted price' (Shin, 1992, p.428), bookmakers are thus able to control and limit their liabilities. The 'insider' may, however, still prefer to use the bookmakers to the Tote, since a substantial bet with the Tote may depress the odds considerably, especially in a small pool. This problem is compounded if the consequent price signal leads to imitative behaviour by outsiders, or else if informed bettors feel the need to take an early fixed price with bookmakers for fear that '...they do not have a monopoly on the insider information.' (Bird and McCrae, 1994, p.578).

Schnytzer and Shilony (1995) are also clear about the advantage to insiders of betting with bookmakers:

Given that with bookmakers, the payout contingent on a win is known when the bet is placed and that inside information is likely to be more accurate as race time draws near, we should expect most 'insiders' to bet with bookmakers. (p.964).

If this is what is happening, a test which distinguishes the data by the likely presence of insider trading may be useful.

The data sample is thus sub-divided into a sample composed of races likely to be characterized by the absence (or a low incidence) of useful private information. If the observed inferiority of returns at bookmakers' odds is caused by the threat of insider trading, then the differential should be less (or disappear even) in samples characterized by the likelihood of low levels of insider activity.

Insiders have a particular potential for advantage in races where the form of the horses is less clearly established in the public domain. In these cases, private information available to a limited number of people with access to the off-course form and ability of a horse is especially valuable. In 'handicap' races, horses are allotted weights so as to equalize as far as possible their chances of winning. To enter these races, horses must usually have run at least three times, or have won a race. Thus, form in handicap races is relatively well-established, and there are likely to be limited opportunities for those with access to private information to secure any useful advantage. In non-handicaps, the public form of the entrants need not be established as a pre-condition of entry, and so opportunities for advantageous use of private information are greater.

For this reason, Crafts (1985) proposed separating handicap races from non-handicaps, as a means of distinguishing races by the likely incidence of insider trading. It is possible, however, that insiders may be able to use public information to improve their private information, as suggested in Vaughan Williams and Paton (1997b). Because of this it may be preferable to consider only higher grade handicaps as indicative of the absence of useful private information. The reason is that these race types (excluding both non-handicaps and handicaps rated below 100)² are the subject of particular media attention

and might be expected to offer very little opportunity for non-disclosure of useful private information. It is assumed that the informational content of any private information available about these race types will be close to zero.

One might argue that lower grade handicaps should be examined as a separate group, i.e. as a subset of the group of races which are not 'higher grade handicaps.' There is a clear case, however, for regarding lower grade handicaps and non-handicaps as one group for purposes of distinguishing access to inside information. This is because both these types of race are characterized by relatively restricted access to public information on the current form of race entrants, and by particular opportunities for the profitable use of private information.

4. Data and estimation

In this section, the mean tote and starting price payouts are compared at an aggregated level and at various levels of sub-aggregation, in order to examine whether any systematic differences can be identified, and if so whether these are systematically exploitable.

The data set collected for this study covers the period from Thursday, March 19, 1992 to Saturday, May 16, 1992 inclusive, a period prior to the introduction of Tote Direct terminals (which allow direct access to tote pools) into the offices of non-Tote bookmakers. This set of races may well offer scope for the existence of unexploited arbitrage opportunities between tote and fixed-odds returns, although potential opportunities are by no means restricted to this period (see, for example, Vaughan Williams and Paton, 1997a).

In all 510 races are recorded, with 509 containing data on the differences between the starting price and tote dividend.³ Since the

Table 1: Testing the significance of differences in the mean payout to a level unit stake on winning horses at tote prices and at starting prices

Races	Number of Observations	Mean Tote Payout	Mean Starting Price Payout	Difference	
All	509	7.86 (10.65)	6.74 (6.93)	1.12* ⁺⁺⁺	t = 4.19 z = 1.82
SP # 8	371	3.65 (2.76)	3.55 (2.20)	0.10	t = 0.87 z = 0.87
SP > 8	138	19.18 (14.93)	15.31 (7.95)	3.87* ⁺	t = 4.14 z = 3.96

Standard deviations are in parentheses. * Significant at the 1 per cent level (using t-tests). ⁺ Significant at the 1 per cent level (using Wilcoxon matched-pairs signed rank tests). ⁺⁺⁺ Significant at the 10 per cent level (using Wilcoxon matched-pairs signed rank tests).

Table 2: Testing the significance of differences in the mean payout to a level stake on winning horses at tote prices and at starting prices - for all non-handicap, handicap, non-higher grade handicap and higher grade handicap races

Races	Number of Observations	Mean Tote Payout	Mean Starting Price Payout	Difference	
Handicaps	224	9.73 (11.38)	8.44 (7.85)	1.29*	t = 2.97 z = 1.34
Non-Handicaps	285	6.39 (9.81)	5.40 (5.77)	0.99*	t = 2.95 z = 1.34
Higher Grade Handicaps	64	9.62 (12.49)	8.36 (7.66)	1.26	t = 1.67 z = 0.75
Non-Higher Grade Handicaps	445	7.61 (10.35)	6.50 (6.79)	1.11* ⁺⁺⁺	t = 3.84 z = 1.72

Standard deviations are in parentheses. * Significant at the 1 per cent level (using t-tests). ⁺⁺⁺ Significant at the 10 per cent level (using Wilcoxon matched-pairs signed rank tests).

tote dividend is declared inclusive of a unit stake, the tote odds are examined net of this, i.e. by subtracting 1 from the published dividend. Standard t-tests for paired data are employed to compare differences in payout. Following Gabriel and Marsden (1990, 1991), the Wilcoxon matched-pairs signed rank test is also used, as this makes no assumption about the distributions from which the data was drawn.

Table 1 compares the mean payout to a level stake on winning horses at starting prices for all races. Following evidence in Blackburn and Peirson (1995), Cain, Law and Peel (1997) and Vaughan Williams and Paton (1997a) that any bias is concentrated in the higher odds categories, a comparison is also made of returns at sub-groupings of odds centred on 8 to 1. This is an arbitrary but not unconventional dividing line between more and less favoured horses - see, for example, Asch, Malkiel and Quandt (1982). All payouts are calculated exclusive of the stake.

Thus, for the set of all races, using t-tests for paired data the null hypothesis of no difference between the tote and starting price payouts about identical winning horses is rejected at all conventional levels of significance. The difference favours the tote. It is, however, only rejected at the ten per cent level, using the Wilcoxon matched-pairs signed rank test suggested by Gabriel and Marsden (1990). For the set of all horses with starting prices of less than or equal to 8 to 1, the null hypothesis of no difference cannot be rejected using either test at any conventional level of significance. There is, however, a clear bias in favour of the tote for the sub-group of horses with starting prices above 8 to 1, these results indicating that the bias in favour of the tote is greater for horses less favoured in the market. Table 2 compares the mean payout to a level unit stake on winning horses at tote prices and at

starting prices - for all non-handicap, handicap, non-higher grade handicap, and higher grade handicap races.

All payouts are calculated exclusive of the stake.

Thus, using a t-test for paired data, it is possible to reject a null hypothesis of no difference in tote and starting price payouts about identical winning horses in all races except higher grade handicaps. The Wilcoxon matched-pairs signed rank test is unable to distinguish these samples as clearly, the test indicating rejection of this null hypothesis only for non-higher grade handicaps, and then only at the ten per cent level of significance. However, the z-value for higher grade handicaps was markedly lower than for any of the other samples. In assessing the legitimacy of these conclusions it should, of course, be noted that the number of observations in the 'higher grade handicap' category is relatively small.

It might be argued that different types of race are characterized by different proportions of horses winning in given odds ranges. In particular, it may be that handicaps, designed as they are to equalize the chances of race entrants, are characterized by a difference in the frequency with which higher-odds horses win. Indeed, Table 3 suggests that this is the case, the incidence of longer-odds horses (as partitioned here) winning being somewhat lower in handicaps (19.6 per cent) than in non-handicaps (36.6 per cent). The essential issue, however, is whether the expected return at Tote and bookmaker odds is different for handicaps and non-handicaps considered separately.

Table 3: Testing the significance of differences in the mean payout, at lower and higher odds, to a level stake on winning horses at tote and at starting prices - for handicaps and non-handicaps.

Races	Number of Observations	Mean Tote Payout	Mean Starting Price Payout	Difference	
Handicaps SP>8	56	19.21 (15.98)	14.73 (6.41)	4.48* +	t = 2.78 z = 2.63
Handicaps SP≤8	229	3.26 (2.82)	3.11 (2.22)	0.15	t = 1.65 z = 0.08
Non- Handicaps SP>8	82	19.15 (14.26)	15.71 (8.87)	3.45* +	t = 3.06 z = 2.91
Non- Handicaps SP≤8	142	4.28 (2.54)	4.25 (1.97)	0.03	t = 0.27 z = 1.13

* Significant at the one per cent level (using t-tests). + Significant at the ten per cent level (using Wilcoxon matched-pairs signed rank tests).

In fact, Table 3 shows that sub-dividing the races into handicaps and non-handicaps does not alter the conclusions drawn from Table 1. As before, the expected tote payout is significantly greater than that at starting prices in the higher odds category. For the set of all horses in the lower odds category, on the other hand, the null hypothesis of no difference cannot be rejected using either test at any conventional level of significance.

5. Discussion

A number of studies have demonstrated a tendency for the expected return to a unit bet in U.K. Tote betting markets to exceed that in U.K. bookmaker markets, at higher odds. One possible explanation, explored here, is that at higher odds bookmakers contract their odds in the face of an adverse selection problem posed by the threat of insiders who possess superior information to themselves. Tote

payouts, on the other hand, being driven by the demand-side of the market, are not prey to this. If this explanation is correct, it would also explain at last in part why bookmaker-based markets are characterized by the favourite-longshot bias identified by, for example, Dowie (1976), Henery (1985), Thaler and Ziemba (1988), i.e. a particular bias against the expected return to bets at higher odds. This would, of course, not explain the bias observed in parimutuel markets. Even if bettors are characterized as pure investors, this bias can still be reconciled with an hypothesis of market efficiency if bettors love risk (e.g. Quandt, 1986; Hamid, Prakash and Smyser, 1996) or skew (e.g. Bird, McCrac and Beggs, 1987; Golec and Tamarkin, 1998). In fact, bettors may be acting at least in part as consumers rather than investors, and as such be responding to a wider set of motivations, perhaps more akin

to a leisure activity (see, for example, Bruce and Johnson, 1992, 1995), or to a broader concept of rationality traceable to the psychological theories of Kahneman and Tversky (1984) and Thaler (1985). Thus, Thaler and Ziemba (1988, p.171) propose that economic behaviour may be context-specific, in which '...people adopt mental accounts and act as if the money is not fungible.' This might explain why a person may be risk-seeking at the racecourse, but risk-averse with respect to pension provisions.

This paper is not designed to disentangle these explanations, but rather to distinguish the influence of one factor, i.e. the threat of insider trading, on bookmaker-based betting markets. To do this a set of races likely to be characterized by a particularly low incidence of insider activity is identified. In these races, it appears that the tendency displayed in other races for bookmakers to offer lower expected returns than the Tote at higher odds disappears. This is at least contributory evidence toward a partial explanation of the favourite-longshot bias in bookmaker-based betting markets. Moreover, this confirms evidence in Vaughan Williams and Paton (1997b) that in races likely to be characterized by a low incidence of insider trading, the bookmakers' margin is uniquely unrelated to the number of runners (and average odds level) in a race.

If this explanation is correct, it points to the fact that insider trading (or even the threat of insider trading) has a part to play in the artificial constriction of returns by the supply side to the demand side of the market here examined. Since this constriction is not arbitrated away in the available complementary (Tote) market, it also raises the issue of whether the market is inefficiently responding to market preferences, or whether the market preferences are themselves somewhat idiosyncratic.

In seeking to apply these results to financial markets, however, it is important to acknowledge the individual and unique features of betting markets which distinguish them from more conventional financial markets. In particular, a betting market is a relatively unsophisticated form of financial market, characterized in the bookmaker-based sector by comparatively high transactions costs, the influence of non-financial factors as a motivating force, and a relatively powerful supply side. Bruce and Johnson (1996, p.8) contrast these '...large, concentrated, commercially-focused, well-resourced, well-informed sellers, protected from their liability by their ability to engage in market influencing strategy' with a demand-side made up of '... a fragmented set of individually very small operators, informationally-disadvantaged, with mixed objectives and limited assets and without effective opportunity to control liabilities via secondary market activity.'

Even so, there are extant in the literature a number of theoretical models of how a conventional financial market might respond to insider activity, for example by increasing the margin (or bid-ask spread) of the market-maker (see, for example, Copeland and Galai, 1983; Glosten and Milgrom, 1985; Kyle, 1985). Interestingly, the empirical evidence produced in this study of betting markets is broadly in line with the theoretical outcomes predicted by these models.

6. Summary and conclusions

This paper has investigated the expected return to level unit stake bets placed on horses at starting price and Tote odds. Evidence is produced for the whole sample of races which demonstrates the superiority of expected tote returns, at least at higher odds levels. Restricting the samples, however, on the basis of the likely potential for insider trading

yields some interesting results. In particular, the differential between tote and starting price odds is less (or disappears) in samples which might be expected to contain lower levels of insider activity, i.e. higher grade handicaps. Without stating the case too strongly, it may be reasonable to infer from this that the findings are at least indicative of some effect on starting prices resulting from bookmakers' response to potential insiders, and that this effect is not fully reflected in tote prices. Whatever conclusions one draws from this evidence, the findings taken in aggregate are in any case consistent with the existence of an apparent information inefficiency with respect to the set of public information in these markets. Given that the opportunities for identifying and exploiting inefficiencies in these relatively unsophisticated markets are relatively clear and straightforward, the results presented here pose questions about the efficiency of more complex and sophisticated financial markets. Even so, the relevance and legitimacy of these questions must be viewed within the context of the unique and distinct nature of betting markets.

An avenue for future research might be to investigate the interface between conventional betting markets and financial markets, in the form of the rapidly growing availability of spread betting markets. In these markets, a 'commodity' (which may range from the number of goals in a football match to the price of gold) may be bought at one figure, determined by the market-maker, and sold at another. The difference between these numbers (essentially the bid-ask spread) is a measure of the market-maker's margin. Returns or losses to the better are equal to the difference between the traded figure and the actual outcome. The advantage offered by these markets to the empirical researcher is that they generate a set of well-defined

outcomes, but at the same time are much closer in their nature and operation to conventional financial markets than are traditional betting markets. In particular, spread betting is treated in law in the same way as trading on financial markets. As such, it is subject to the 1986 Financial Services Act, Schedule 1, sections 9 and 12, and is regulated by the Securities and Futures Authority Ltd. Because of the tax system, these two markets are also characterized by rather more comparable transactions costs. The amounts that bettors stand to win or lose from a single bet are also much higher, and so the market is likely to be populated by traders whose motivations are more similar to that of financial traders. Finally, the direct quote of a spread between a buying and selling price also makes them more amenable to a theoretical modelling normally applied to more conventional financial markets.

In the meantime, the evidence presented here might usefully add to the debate about the costs of insider trading in a defined market.

Endnotes

1. The Nottingham Trent University
2. Flat race handicaps are normally rated between 0 and up to 115, a higher rating indicating a higher grade race. Thus a typical low grade handicap, for example, would be restricted to horses rated from 0 to 60.
3. The missing item is the first race at Newcastle on 20 April 1992, which failed to produce a tote dividend owing to a technical fault.

References

- Ali M M (1979) 'Some Evidence on the Efficiency of a Speculative Market', *Econometrica*, 47, 387-392.
- Asch P, Malkiel B G and Quandt R E (1982) 'Racetrack Betting and Informed Behavior', *Journal of Financial Economics*, 10, 187-194.
- Asch P and Quandt R E (1987) 'Efficiency and Profitability in Exotic Bets', *Economica*, 59, 278-298.
- Bird R and McCrae M (1994) 'Efficiency of Racetrack Betting Markets: Australian Evidence', in *Efficiency of Racetrack Betting Markets*, Donald B Hausch, Victor S Y Lo and William T Ziemba (eds.) London: Academic Press, 575-582.
- Bird R, McCrae M and Beggs J (1987) 'Are Gamblers Really Risk Takers?', *Australian Economic Papers*, December, 237-253.
- Blackburn P and Peirson J (1995) 'Betting at British Racecourses: An Analysis of Semi-Strong Efficiency Between Bookmaker and Tote Odds', *Studies in Economics*, University of Kent, 95/4.
- Bruce A C and Johnson J E V (1992) 'Toward an Explanation of Betting as a Leisure Pursuit', *Leisure Studies*, 11, 201-218.
- Bruce A C and Johnson J E V (1995) 'Costing Excitement in Leisure Betting', *Leisure Studies*, 14, 48-63.
- Bruce A C and Johnson J E V (1996) 'Market Power in British Pari-Mutuel and Bookmaker-Based Horserace Betting Markets: An Empirical Analysis', *University of Nottingham School of Management and Finance Discussion Papers*, 1996.11
- Cain M, Law D and Peel, D (1997), 'Insider Trading and Market Efficiency in British Racetrack Betting', *Salford Papers in Gambling Studies*, 97-01.
- Copeland T and Galai D (1983) 'Information Effects on the Bid Ask Spread,' *Journal of Finance*, 38, 1457-69.
- Crafts N F R (1985), 'Some Evidence of Insider Knowledge in Horse Race Betting in Britain', *Economica*, 52, 295-304.
- Dolbear F T, Jr (1993), 'Is Racetrack Betting on Exactas Efficient?', *Economica*, 60, 105-111.
- Dowie J A (1976) 'On the Efficiency and Equity of Betting Markets', *Economica*, 43, 139-150.
- Fama E (1991) 'Efficient Capital Markets, II', *Journal of Finance*, 46 (5), 1575-1617.
- Gabriel P E and Marsden J R (1990) 'An Examination of Market Efficiency in British Racetrack Betting', *Journal of Political Economy*, 98, 874-885.
- Gabriel P E and Marsden J R (1991) 'An Examination of Market Efficiency in British Racetrack Betting: Errata and Corrections', *Journal of Political Economy*, 99, 657-659.
- 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

- Racetrack', *Journal of Political Economy*, 106 (1), 205-225.
- Hamid S S, Prakash A J and Smyser M W (1996) 'Marginal Risk Aversion and Preferences in a Betting Market', *Applied Economics*, 28, 371-376.
- Henery R J (1985) 'On the Average Probability of Losing Bets on Horses with Given Starting Price Odds', *Journal of the Royal Statistical Society*, 148 (4), 342-349.
- Kahneman D and Tversky A (1984) 'Choices, Values and Frames', *American Psychologist*, 39, 341-350.
- Kyle A (1985) 'Continuous Auctions and Insider Trading', *Econometrica*, 53, 1315-1335.
- Lo V S Y and Busche K (1994) 'How Accurately Do Bettors Bet in Doubles?', in *Efficiency of Racetrack Betting Markets*, Donald B Hausch, Victor S Y Lo and William T Ziemba (eds.), London: Academic Press, 465-468.
- Quandt R E (1986) 'Betting and Equilibrium', *Quarterly Journal of Economics*, 101, 201-207.
- Schnytzer A and Shilony Y (1995) 'Inside Information in a Betting Market', *Economic Journal*, 105, 963-971.
- Shin H S (1991) 'Optimal Betting Odds Against Insider Traders', *Economic Journal*, 101, 1179-1185.
- Shin H S (1992) 'Prices of State Contingent Claims with Insider Traders, and the Favourite-Longshot Bias', *Economic Journal*, 102, 426-435.
- Shin H S (1993) 'Measuring the Incidence of Insider Trading in a Market for State-Contingent Claims', *Economic Journal*, 103, 1141-1153.
- Thaler R (1985) 'Mental Accounting and Consumer Choice', *Marketing Science*, 4, 199-214.
- Thaler R and Ziemba W (1988) 'Anomalies - Parimutuel Betting Markets: Racetracks and Lotteries', *Journal of Economic Perspectives*, 2, 161-174.
- Vaughan Williams L and Paton D (1997a), 'Does Information Efficiency Require a Perception of Information Inefficiency?', *Applied Economics Letters*, 4, 615-617.
- Vaughan Williams L and Paton D (1997b), 'Why is There a Favourite-Longshot Bias in British Racetrack Betting Markets?', *Economic Journal*, 107, 150-158.