3

The Efficiency of Australian Football Betting Markets

by

Timothy J. Brailsford †
Stephen A. Easton ‡
Philip K. Gray §
Stephen F. Gray ¶

Abstract:

This paper examines the efficiency of the two major Australian football betting markets: the Australian Rugby League (ARL) ***FootyTAB*** *market and the Australian Football League (AFL)* ***Footywin*** *market. Probit and ordered probit models are tailored to the unique structures of the markets. This circumvents some potential econometric problems, and also allows us to test betting strategies in which a bet is placed only when there is a high ex-ante probability of success. Our probit models are successful in predicting game outcomes in both the ARL and AFL. While several of our betting strategies generate significant profits, both in-sample and out-of-sample, we offer a number of reasons why we are cautious about interpreting these results as conclusive evidence of market inefficiency.*

Keywords:

MARKET EFFICIENCY; SPORTS BETTING; PROBIT.

† Department of Accounting and Finance, University of Melbourne, Parkville VIC 3052.
‡ Department of Accounting and Finance, Monash University, Clayton VIC 3168.
§ School of Economics and Finance, Queensland University of Technology, Brisbane QLD 4001.
¶ Fuqua School of Business, Duke University, Durham NC 27708, U.S.A.

We are grateful to Barry Oliver, Peter Whelan, and participants at the the 1995 AAANZ conference in Melbourne for valuable comments and suggestions. Funding from Coopers & Lybrand is gratefully acknowledged.

1. Introduction

The concept of market efficiency is now well established since the early work of Fama (1970) and others. Market efficiency is typically discussed and investigated in the context of financial markets due to the ready availability of large databases. Despite numerous tests of market efficiency on different samples, there is still no clear consensus of opinion [see Fama (1991) for a review]. Tests of efficiency tend to fall into two categories. First, market prices or returns are examined for their time-series properties. Researchers have looked for patterns in past returns, autocorrelations, mean reversion, the profitability of trading strategies based on filter rules, and trends [see Boudoukh, Richardson and Whitelaw (1994)]. This group of studies is generally recognised as testing weak form efficiency. Second, studies have examined the information content of key events such as accounting disclosures, release of economic data, contagion effects and security-specific events. This group of studies is generally recognised as testing semi-strong form efficiency. Other tests of market efficiency have examined the information content and rôle of private information. But these studies are not common.

In an efficient stock market, the share price is an unbiased view of the "true" value of the company. It is difficult, however, to formulate a direct test of the efficiency of a stock market—since the market is continually forming expectations of future cash flows, the true value of a particular investment is never revealed. Hence, efficiency tests have tended to focus on predictability of asset *returns*. Nevertheless, there is no reason why the principles of market efficiency cannot also be applied to any other competitive market.

Sports betting markets provide an ideal arena for testing efficiency since the range of possible asset payoffs is often known with certainty in advance. While the ex-ante distribution of these payoffs is not known, its form is simple and ex-post realisations are available soon after each bet is placed. This structure permits a direct test of market efficiency. In point spread betting markets, the spread should reflect the "true" outcome of the game. Unlike financial markets, we can observe the true outcome and assess whether spreads are unbiased. Further, just as share prices adjust to reflect the value of new information, spreads and odds in sports betting markets similarly adjust to reflect the value of new information. While the nature of prices is quite different in financial and sports betting markets, in both cases they represent information about expected outcomes. In this sense, we can examine the impact of publicly available information on spreads and odds (i.e. prices).

The financial markets literature has other strong analogies to sports betting markets. A recent strand of the literature has argued for the importance of psychological factors in financial markets [see Heisler (1994) and Thaler (1994) for reviews of this literature]. One example of a psychological factor is the attribute of overreaction. The early work of De Bondt and Thaler (1985, 1987) posited that the stock market overreacts to new information. Numerous papers have followed this work investigating the profitability of the resultant contrarian trading strategy. This literature is equally applicable to sports betting markets.

There is a well-documented favourite/long shot bias in racetrack betting [see Thaler and Ziemba (1988)]. That is, favourites tend to be under-backed and long shots tend to be over-backed. Hence, an obvious test of efficiency is whether a contrarian strategy yields excess profits. Moreover, the behavioural literature suggests that individuals underestimate the likelihood of extreme outcomes.[1] Combined with the favourite/long shot bias, this suggests that extreme wins by favourites will be underbacked, providing another natural test of efficiency.

There is also a growing literature on the effect of changes in CEO (and other senior executives) on equity returns.[2] This literature examines the importance of the CEO to the firm (as measured by market reaction). Similarly, we can examine the importance of key players to game outcomes. In some sports competitions, key players are missing due to representative commitments that occur during the regular season. While the event studies of CEO changes are unable to separate market perceptions from real effects (since the true value of the firm is never revealed), we can separately examine the market's view of the importance of key players (via spreads and dividends) and the actual importance (via the effect on actual winning percentages).

Finally, research has been conducted into the effect of "momentum" in stock returns [see Clarke and Statman (1994)]. It is argued that growth stocks that attain momentum continue to earn positive excess returns. In part, this phenomenon is related to the overreaction phenomenon. In sports betting markets, momentum is akin to teams that have exhibited a recent winning or losing streak. Tests of market efficiency can examine whether this momentum translates into predictable game outcomes and profitable betting strategies. In contrast, there is a strand of the applied psychology and economics literature that demonstrates that perceived "streaks" by sports teams really do not exist. Camerer (1989), for example, finds some evidence of a belief in the "hot hand" in spreads for professional basketball games, but finds no evidence of the existence of streaks in actual game outcomes. Teams with long winning streaks tend to perform worse than expected, reverting to the mean rather than continuing their streak.

This paper advances the literature that examines the efficiency of sports betting markets in a number of directions. First, the paper examines two Australian football betting markets that have unique and interesting structures that allow us to gain an increased understanding of the efficiency of markets. The Australian Rugby League (ARL) *FootyTAB* market offers a point spread on each game. While this market structure has many similarities to the standard point spread betting markets examined in previous research, there are some important differences (discussed in Section 3). In contrast, the Australian Football League

1. See, for example, Winkler (1967), Alpert and Raiffa (1969), Stael von Holstein (1971) and Kahneman, Slovic, and Tversky (1982).
2. See, for example, Denis and Denis (1995), Johnson, Magee, Nagarajan, and Newman (1985), Warner, Watts, and Wruck (1988), and Weisbach (1988).

(AFL) *Footywin* market does not offer point spreads. Rather, bettors are required to select game outcomes to within a 12-point range. This market structure is unique by world standards and complicates econometric tests of efficiency.

Second, we use econometric techniques that are tailored specifically to the unique market structures. In the ARL *FootyTAB* point spread market, we use a discrete-choice probit model rather than the OLS regression methodology commonly used in prior literature. The sensitivity of OLS to extreme outliers can make OLS estimates difficult to interpret. This is because the return on a bet is the same whether a team beats the spread by one point or by thirty points. Whereas OLS effectively gives more weight to the more extreme result, the probit model sees both results as "wins" and treats them equally. A discrete-choice probit model, however, is not suited to the AFL *Footywin* market where bettors are required to forecast winning/losing margins to within a 12-point range rather than a simple binary win/loss. Accordingly, we develop an ordered probit model that predicts the range in which the game outcome is likely to fall.

Third, we examine the efficiency of the football betting markets from the perspective of profitable trading rules. Market inefficiency requires that trading strategies earn consistent profits. This paper examines the profitability of trading strategies based on biases documented both in this paper and in previous research. The probit model is ideally suited to examine issues such as the probability of Team A beating the spread against Team B in the ARL, conditional on available information, or the probability of an AFL result falling in a certain 12-point range. Consequently, we can examine the profitability of trading rules that involve, for example, placing a bet only if the conditional probability of success exceeds X% (some filter).

Fourth, we consider a number of explanatory variables that have strong analogies to financial markets and behavioural research. Does the market overreact to new information so that a contrarian betting strategy is profitable? Is the absence of key players efficiently incorporated into spreads? Can we condition on recent winning or losing streaks to say something about the (conditional) distribution of future outcomes?

The remainder of the paper is organised as follows. Section 2 reviews previous research on the efficiency of sports betting markets. Section 3 discusses the different structures of the ARL and AFL betting markets. We outline our econometric methods and report the results in Section 4. Section 5 concludes.

2. Previous Tests of Sports Betting Market Efficiency

In the market efficiency literature, especially as it relates to sports betting, two definitions of efficiency have been used. The narrow view, which is primarily of academic interest, posits that the return from any betting strategy should be negative and equal in magnitude to the commission of the betting house. The broad view, which is also of practical interest, posits that no betting strategy should yield significantly positive returns (after commissions) on average. In this paper we focus on the latter definition. We test for efficiency by searching for betting strategies that yield significantly positive returns, on average.[3] If the market is efficient, there exists no such strategy, as the spread or odds capture all relevant information. Previous research has also distinguished between statistical and economic efficiency. The former examines whether game outcomes are statistically predictable, while the latter focuses on trading profits from implementable betting strategies. This paper examines both of these issues.

The question of whether organised sports betting markets are efficient has recently received much attention in the literature. For example, Golec and Tamarkin (1991) document that spreads set in the U.S. National Football League (NFL) betting market are systematically biased predictors of actual results. This finding is sometimes offered as evidence of inefficiency in the sports betting market, although it is not clear that this bias can be exploited via a profitable trading strategy. In particular, Golec and Tamarkin document that betting on home team underdogs was consistently profitable over their sample period. Gray and Gray (1995) show that this bias has attenuated over time and no longer exists. Gandar, Zuber, O'Brien, and Russo (1988) are unable to find any *statistical* evidence of inefficiency in the NFL betting market using a relatively simple econometric model. They establish evidence of *economic* inefficiency, however, by demonstrating the profitability of various betting strategies. Thaler and Ziemba (1988) review the literature on racetrack betting markets, noting the consistent favourite/long shot bias [see for example, Ali (1977)]. The expected return from betting on favourites is significantly greater than the expected return from betting on long shots. Swidler and Shaw (1995) fail to find evidence of this bias in a small racetrack betting market. Woodland and Woodland (1994) document that the favourite/long shot bias in racetrack betting exists in reverse for baseball bettors, but that no betting strategy earns profits in excess of commissions. Brown and Sauer (1993) show that observable variables can be used to predict outcomes of professional basketball games, beyond the information contained in the spread.

The Australian evidence with respect to sports betting markets is limited. Bird and McCrae (1987) and Bird, McCrae, and Beggs (1987) examine the efficiency of bookmakers' markets at Melbourne horse racetracks. Consistent with overseas evidence, they find that average returns from betting on favourites are significantly greater than average returns from betting on long shots, but that no strategy yields a positive return. Bird and McCrae note that prior knowledge of movements in odds during the course of betting could lead to significant returns. Such knowledge, however, is privileged information and in the football markets studied in this paper, is impossible to obtain. In contrast, Tuckwell (1983) finds that changes in horse racing odds do not move in a random fashion and concludes

3. Since the range of betting strategies is limitless, our tests do not provide a complete test of efficiency. Nevertheless, we do test a diverse set of strategies.

that racetrack betting in Australia is inefficient. While this result is prime facie evidence of statistical inefficiency, there is no direct link to economic inefficiency.

3. Market Structure

3.1 ARL FootyTAB Betting

The ARL competition is effectively a national league, comprising either sixteen or twenty teams during our sample period.[4] There are eight games per week (ten in 1995) over a regular season of 22 rounds.[5] Betting on ARL games is facilitated by the State-run Totalizator Agency Board (TAB), and trades under the name *FootyTAB*. Each Monday, a panel of "experts" sets the spreads for the following round of games. The consensus spread is published on Tuesday and becomes the official *FootyTAB* spread, remaining fixed over the course of the week. The relative volume of betting on each side is not captured by a changing spread, but rather by a changing dividend. The dividend is calculated in parimutuel fashion after the outcome of the match, such that 75% of the total amount wagered is returned to successful bettors. Individual bettors, therefore, do not know what the dividend will be when placing their bets.[6]

There are several betting systems offered: *Pick the Score* and *Pick the Result* (select the winning team and margin) are offered on several selected games each week, *Pick the Winners* requires bettors to pick the winning team, relative to the spread, in all games for the week, and *Pick the Margins* requires bettors to select the winning team and margin (to within a 12-point range) in each game. *Pick the Winners*, however, is the only system that requires bettors to take the point spreads into account. Accordingly, this paper's analysis of point spread betting focuses on *Pick the Winners*.

There are three major differences between ARL *FootyTAB Pick the Winners* betting and standard point spread betting markets. First, the *FootyTAB* spreads do not capture changing conditions leading up to match day (for example, weather conditions, injuries to players, and ground conditions). While this may introduce biases into the spread for a single game, such biases are unlikely to be systematic. Moreover, Gander et al. (1988) report that even when spreads are truly market-determined, less than 5% of all "line movements" are due to public information about injuries and weather. Second, *FootyTAB* bettors do not know what the winning dividend will be until after the games are completed. The implications of this are discussed further in Section 4. Third, *FootyTAB* bettors cannot place a bet

4. In 1995, the ARL competition was expanded from sixteen to twenty teams.
5. Most games are played on Sunday afternoon. One game per week is played on Friday night and several games per week are played on Saturday.
6. In a standard point spread betting market, a fixed dividend is known prior to placing the bet. The relative volume of betting on each side causes the spread to adjust.

on a single game—*Pick the Winners* requires correctly selecting the result in *all* games. While other betting establishments in Australia offer point spread betting on single games, data are not readily available.[7]

3.2 AFL Footywin Market

The AFL is also effectively a national league, similar in structure to the ARL. The AFL competition comprised between fourteen and sixteen teams during our sample period, with between seven and eight games per round. The number of rounds per season varied from 22 to 24.

Betting on AFL games is conducted by *TabCorp* under the name of *Footybet*.[8] *Footybet* is the generic term given to a variety of betting schemes requiring bettors to correctly predict the range in which game outcomes will fall: (*a*) *Footywin* takes bets on single games, (*b*) *Footydouble*, *Footyxtradouble*, *Footyquad* and *Footy Pick 6* require bettors to correctly predict the range in which the game outcome will fall for two, three, four and six specified games, respectively. Our focus in this paper is on *Footywin*.

Footywin is a unique form of odds scheme that does not require the posting of a spread. Rather, bettors are required to predict the outcome of a game to within a 12-point range. There are fifteen of these 12-point "bins".[9] Bin 1 covers a 1 to 12 point win by the home team. Bin 2 covers a 13 to 24 point win by the home team, and so on. Bin 7 covers a home team win in excess of 72 points. Bin 8 covers a drawn game. Bin 9 covers a 1 to 12 point win by the away team and so on up to Bin 15 which covers a 73 point plus win by the away team. The dividend is calculated, in parimutuel fashion, after the outcome of the match, such that a certain percentage of the total amount wagered is returned to successful bettors.[10]

---

7. For example, point spread betting is offered in some casinos, at racetracks, and by legal betting houses operating out of the Northern Territory, in addition to illegal gambling. It is implicitly assumed in this paper that the spreads and support rates for the *Pick the Winners* system would be applicable for single game betting and that we can therefore treat each game as if it were available for individual betting. In the Las Vegas NFL betting market, the spread on a particular game is the same for a single-game bet as for a multiple-game bet. While we cannot confirm that this would be the case for our ARL data set, we see no reason to conclude that the *Pick the Winners* system causes any bias in the spreads or support rates. Moreover, it is clear that as spreads become more biased, game outcomes become more predictable and the total amount bet decreases. In the extreme case, a spread may be so biased that the outcome (relative to that spread) is virtually certain. In this case, all bets would be placed on the same team and all bettors would lose the 25% *FootyTAB* commission. Therefore, no money will be bet when the outcome is certain. Since the objective of the TAB is to maximise the total amount bet, it seems likely that the *FootyTAB* spread would not be dramatically biased for either a single-game or a multi-game bet.
8. *TabCorp* was formed in August 1994 from the privatisation of the Victorian State-run Totalizator Agency Board.
9. *Footywin* was originally established with 23 bins of variable point ranges on nominated games. Our sample covers the population since 12-point bins were introduced in 1987.
10. The commission retained by *TabCorp* and its predecessor has varied over the sample period, but averages around 20%.

Individual bettors, however, do not know what the dividend will be when placing their bets.

4. Econometric Methods and Results

4.1 Data

Our ARL data consists of *FootyTAB* spreads, game outcomes, dates and locations for all teams from 1989 to 1995, a total of 1,276 games.[11] These data were obtained directly from the New South Wales Totalizator Agency Board. The ARL data set also includes the total amount bet, the total pool, and the total number of winning bets, from which the dividend can be calculated, as well as the proportion of bets placed on each team (the "support rate").

Our AFL data consists of game outcomes, dates, locations, and the dividend[12] for all games from 1987 to 1995, a total of 1,437 games.[13] These data were obtained from a combination of the Victorian Totalizator Agency Board (prior to TabCorp), various TAB agencies, and the Australian Football League.

4.2 The Efficiency of the ARL FootyTAB Market

4.2.1 Summary Statistics For each ARL game, the points spread and the result can be defined relative to either team. If Team X is given 10.5 points start against Team Y, then the start is +10.5 from the perspective of Team X and –10.5 from the perspective of Team Y. We denote the team from whose perspective the start and result is defined as the *team of record*. Since there is no single correct way of choosing the team of record, three different methods are used, following Golec and Tamarkin (1991). First, the team of record is defined to be the favourite. Second, the team of record is defined to be the home team. Third, the team of record is chosen randomly to avoid any systematic effects. Assigning the team of record on the basis of a specific attribute may uncover a conditional bias, such as favourites and home teams receiving too much start, for example. Randomly assigning the team of record is unlikely to uncover any bias, since there is nothing systematic about the teams of record.

Summary statistics of point spreads and outcomes for ARL games are reported in Table 1. The team of record is defined in the three different ways described above. There appears to be some overconfidence in the favourite, as the favourite wins by less than expected (according to the spread). On average, the favourite gives up 8.18 points start and wins by only 7.25 points. In contrast, the

11. No spread was offered on one game in 1993, Round 5 reducing our sample to 1,275 games.
12. The minimum quoted unit was changed in 1992 from $0.50 to $1. All data prior to 1992 were rescaled.
13. There were two missing observations in 1989, Round 17 and one in 1993, Round 5, leaving a sample of 1,434 games.

Table 1

Summary Statistics: ARL[1]

	Point Spread			Outcome		
Parameter	Favourite	Home	Random	Favourite	Home	Random
Mean	8.18	0.75	0.20	7.25	4.18	–0.35
Std Dev.	5.13	9.63	9.66	16.45	17.48	17.97

Note: 1. The data consist of 1,275 ARL games from 1989 to 1995. The *Point Spread* columns represent expected outcomes, based on the spread. The *Outcome* columns refer to actual results. The *Favourite* columns refer to the average points scored by the favourite in excess of the underdog. The *Home* columns refer to the average points scored by the home team in excess of the away team. The *Random* columns refer to the average points scored by a randomly selected team in excess of the opponent.

value of the home field advantage appears to be discounted. On average, the home team gives up 0.75 points start and wins by 4.18 points. Obviously, there are no systematic biases when the team of record is defined randomly.

These preliminary results are suggestive of potential systematic biases, or inefficiencies, in the ARL *FootyTAB* market. In the following section, we test for statistical evidence of the apparent inefficiencies.

4.2.2 Predictability of Game Outcomes For a point spread betting market to be statistically efficient, the spread must be an unbiased predictor of the game outcome—the game-time spread should reflect the aggregate information in the market. Therefore, knowing the magnitude of the spread, or whether the team of record is the home team or the favourite, should not help in consistently predicting whether or not the team of record will beat the spread. To examine whether the market efficiently incorporates all publicly available information into the price, it is simply a matter of testing whether any publicly available information is useful in predicting game outcomes, beyond the information contained in the spread. Previous tests of efficiency have used an OLS framework to test this [see Golec and Tamarkin (1991) for a review].

$$R_i = b_0 + b_1 S_i + b'_2 X_i + \varepsilon_i \tag{1}$$

where R_i is the game outcome, S_i is the spread, and X_i represents all other conditioning information. If the spread captures all relevant information, $b_1 = 1$, $b_0 = 0$ and all of the elements of b_2 are zero. Market efficiency is typically examined by testing this joint hypothesis. In particular, if any of the elements of

b_2 are significantly different from zero, the information in X_i can be used to predict whether or not the spread will be beaten.

In this paper, we formally test for statistical efficiency using a probit model rather than the standard OLS methodology. Since OLS has the potential to overweight outliers, the OLS slope coefficient is influenced not only by the identity of the winning team, but also by the winning margin. Games in which the winning margin is large have a greater impact on the slope coefficient. In the ARL *FootyTAB* market, the only relevant question is *whether* a certain team beat the spread; the *amount* by which the team beat the spread is irrelevant. Therefore, rather than the winning or losing margin itself, we use as our dependent variable a discrete-choice variable which is defined as follows:

$$Y_i = \begin{cases} 1 & \text{if the team of record beat the spread} \\ 0 & \text{otherwise.} \end{cases}$$

We begin by examing the home team and underdog biases documented in the NFL by Golec and Tamarkin (1991) and suggested by Table 1 using the following probit model:

$$Y_i^* = b_0 + b_1\, SPREAD_i + b_2\, HOME_i + b_3\, FAV_i + \varepsilon_i \tag{2}$$

where the team of record is defined on the basis of random selection. In a probit model, $Y_i = 1$ if $Y_i^* > 0$ and $Y_i = 0$ otherwise, where Y_i and Y_i^* are defined above and

$$HOME_i = \begin{cases} 1 & \text{if the team of record is playing at home} \\ 0 & \text{otherwise} \end{cases}$$

$$FAV_i = \begin{cases} 1 & \text{if the team of record is the favourite} \\ 0 & \text{otherwise} \end{cases}$$

The spread is included as an explanatory variable to capture any bias associated with the degree of favouritism. For example, very strong favourites may be more likely to beat the spread than mild favourites. The results of estimation of the probit model appear in Table 2.

Table 2 indicates that home teams and underdogs are more likely to beat the spread than are away teams and favourites. The coefficients on our home team dummy variable (b_2) and the favourite dummy variable (b_3) are statistically significant. The coefficient on the spread is insignificantly different from zero, indicating that strong and mild favourites are equally likely to beat the spread. The probit model correctly predicts whether a team will beat the spread for 59.64% of ARL games.

Next, we expand the set of conditioning information to capture two additional effects. First, we examine whether winning or losing “streaks” are incorporated into the spread in an unbiased way. This has the analogy of momentum investing in equity markets. We define $HL4_i$ to be the number of

Table 2

Probit Results: ARL[1]

Parameter	Estimate	Std Error
b_0	-0.158^*	0.083
b_1	0.008	0.007
b_2	0.533^*	0.074
b_3	-0.399^*	0.135
% Correct Predictions:	59.64	

Note: 1. The data consist of 1,275 ARL games from 1989 to 1995. * = significant at 5%. The model is $Y_i^* = b_0 + b_1\, SPREAD_i + b_2\, HOME_i + b_3\, FAV_i + \varepsilon_i$, where $Y_i = 1$, indicating that the team of record beat the spread, if $Y_i^* > 0$ and $Y_i = 0$ otherwise.

times the home team has beaten the spread in the last four games played (a number between zero and four), $AL4_i$ to be the number of times the away team has beaten the spread in the last four games played (a number between zero and four), and FAV_i takes a value of one when the home team is the favourite and zero otherwise.

Second, our ARL data also allows us to examine the market's assessment of the importance of key players. In the ARL competition, several representative matches are played each year *during* the regular season.[14] This results in some teams being without several key players for some regular season games. To examine whether the spread unbiasedly captures the effect of these missing key players, we define the dummy variable:

$$REP_i = \begin{cases} 1 & \text{if the team of record has more players missing than opponent} \\ -1 & \text{if the team of record has less players missing than opponent} \\ 0 & \text{otherwise} \end{cases}$$

The marginal impact of both of these effects for ARL games is estimated using the following augmented probit model. To aid the interpretation of results and economise on parameters, we define the team of record to be the home team. Therefore Y_i takes a value of one when the home team beats the spread. This

14. Representative matches consist of the three-match Queensland versus New South Wales *State of Origin* series and a three-match *Test* series against a visiting international team.

model can, therefore, be thought of as a tool to predict the probability that the home team will beat the spread:

$$Y_i^* = b_0 + b_1\, HL\,4_i + b_2\, AL\,4_i + b_3\, FAV_i + b_4\, REP_i + \varepsilon_i. \tag{3}$$

Table 3 contains results of the augmented probit model.

Table 3

Augmented Probit Results: ARL[1]

Parameter	Estimate	Std Error
b_0	0.645*	0.107
b_1	–0.112*	0.031
b_2	0.018	0.024
b_3	–0.355*	0.088
b_4	–0.045	0.046
% Correct Predictions:	59.77	

Note: 1. The data consist of 1,034 ARL games from 1989 to 1995. * = significant at 5%. The augmented probit model is $Y_i^* = b_0 + b_1\, HL\,4_i + b_2\, AL\,4_i + b_3\, FAV_i + b_4 REP_i + \varepsilon_i$, where $Y_i = 1$, indicating that the home team beat the spread, if $Y_i^* > 0$ and $Y_i = 0$ otherwise. Some early season observations are lost due to the construction of our recent-record variables. We have reëstimated the model using data from the end of the previous season to construct these recent-form variables. The results are virtually identical and are available on request.

The home team and underdog biases are again obvious. b_0 is positive and statistically significant, indicating that home teams are more likely to beat the spread than are away teams. Moreover, b_3 is significantly negative indicating that home favourites are less likely to beat the spread. These results confirm the home-underdog bias in the setting of spreads. Table 3 does not support the hypothesis that teams on winning or losing streaks gain momentum. The negative coefficient on b_1 indicates that home teams on a winning streak are less likely to beat the spread. The positive coefficient on b_2 suggests the home team is more likely to beat the spread when their visiting opponent is on a winning streak, although this coefficient is not statistically significant. These results are consistent with Camerer (1989) who found that teams on winning streaks tend to perform worse than expected (i.e. spreads overreact to recent form).

The representative game dummy variable enters with a negative coefficient possibly indicating that a team missing several key players is less likely to beat the spread. That is, the spread tends to underestimate the importance of losing key players. But since b_4 is not statistically significant, we draw no conclusions. Overall, the augmented probit model correctly predicts whether the home team will beat the spread in 59.77% of the ARL games.

4.2.3 Specification Tests To further document the performance of our probit model, Table 4 compares the model probability of success with the actual success rate.

Table 4

ARL Probit Probabilities and Actual Success Rates[1]

Probit Probability	Proportion of Games (%)	Success Rate (%)
< 45	3.00	51.61
45 – 50	8.61	46.07
50 – 55	14.89	54.55
55 – 60	30.85	55.80
60 – 65	24.47	63.24
65 – 70	9.09	65.96
> 70	9.09	75.53

Note: 1. The data consist of 1,034 ARL games from 1989 to 1995. *Probit Probability* refers to the probability that the home team will beat the spread, generated from the augmented probit model, *Proportion of Games* refers to the proportion of games in the sample that have a probit probability in the particular range, and *Success Rate* refers to actual success rates in the sample.

For example, we take all games for which the probit model predicts that the probability that the team of record will beat the spread is within a certain range (e.g., 50–55%) and calculate the actual proportion of successes for these games.

Table 4 reports that the probabilities generated by the probit model are remarkably accurate. Apart from the lower-extreme category, where there are very few observations, the model probabilities correspond closely to the actual frequency of success. Since the model is in terms of the probability that the home team will beat the spread, and because home teams are relatively more likely to beat the spread, the range of probabilities is not symmetric around 0.5.

4.2.4 Tests of Actual Betting Strategies The results in Table 4 suggest a novel betting strategy: place bets only on those games where the probability of success is greater than X%. Under this rule, *X* is used as a filter, so that bets are placed only when there is a high probability of success. Table 5 reports the success rate of this strategy, and some other, simple betting rules. The success rate refers to the prediction of game outcomes and at this stage ignores the parimutuel dividend.

Table 5 also reports two *p*-values that indicate the likelihood that a random betting strategy would generate more successes than a particular strategy. The bootstrap *p*-value is constructed by randomly sampling *N* games, with replacement, from our data set, where *N* is the number of games on which a bet was placed under the particular strategy.[15] For each game, we randomly select a team and place a notional bet on that team. This whole procedure is repeated 1,000 times. The bootstrap *p*-value is the proportion of the 1,000 simulations that have success rates greater than the success rate of our particular strategy.

Parametric *p*-values, which are based on the binomial distribution, are also reported in Table 5. Suppose a particular strategy generates *P* successes from *N* bets. If we were to randomly select a team on which to bet for a particular game, the probability that we would have chosen the correct team is 0.5. To test whether the success rate from a particular strategy is significant, or merely within the bounds of chance, we calculate, from the binomial distribution, the probability of obtaining more than *P* "successes" from *N* "draws" where the probability of success is 0.5.

Table 5 shows that the success rate of betting on underdogs and home teams is 54.46% and 59.43%, respectively. Combining these biases in a home-underdog betting strategy has a success rate of 65.38%, indicating that ARL game outcomes are quite predictable even with very simple betting rules.

Betting strategies based on the probit model can generate success rates of over 70% out-of-sample.[16] Betting on all teams predicted by the probit model has

---

15. While strategies such as betting on home teams can be implemented on all games, other strategies depend on match-specific circumstances and hence do not involve placing a bet on every game. For example, betting on home team underdogs can not be implemented when the home team is the favourite. Also recall from Footnote 6 that although *FootyTAB* bets cannot be placed on individual games, we assume that there is no reason that spreads for multi-game bets should differ from spreads for single-game bets.
16. For all of our model-based strategies, we estimate the model using data from 1989 to 1993 and report out-of-sample success rates using data from 1994 and 1995. Thus, the parameter estimates and probit probabilities would have been available at the time the bets were placed. In contrast to Zuber, Gandar, and Bowers (1985), our model performs equally well over the in-sample and out-of-sample periods. The high success rates generated out-of-sample, and the relatively simple and intuitive nature of our model, indicate that our in-sample performance is not due to over-fitting.

Table 5

ARL Betting Strategy Success Rates

Betting Strategy	Proportion of Games	Success Rate	Bootstrap *P*-Value	Binomial *P*-Value
Bet on Underdog	100.00	54.46	0.000	0.001
Bet on Home Team	100.00	59.43	0.000	0.000
Bet on Home Team Underdog	45.15	65.38	0.000	0.000
Bet on Team Predicted By Probit Model				
In-Sample	100.00	59.77	0.000	0.000
Out-of-Sample	100.00	58.10	0.001	0.002
Bet Only When Probit Probability of Success Exceeds 55%				
In-Sample	76.50	61.44	0.000	0.000
Out-of-Sample	83.17	59.54	0.001	0.001
Bet Only When Probit Probability of Success Exceeds 60%				
In-Sample	42.65	66.44	0.001	0.000
Out-of-Sample	47.94	64.24	0.000	0.000
Bet Only When Probit Probability of Success Exceeds 65%				
In-Sample	18.18	70.75	0.000	0.000
Out-of-Sample	19.37	63.93	0.014	0.015
Bet Only When Probit Probability of Success Exceeds 70%				
In-Sample	9.09	75.53	0.000	0.000
Out-of-Sample	10.79	73.53	0.001	0.003

Note: 1. The data consist of 1,275 ARL games from 1989 to 1995. In fitting the augmented probit model, games from early in the season are lost due to the construction of our "recent form" variables. In these cases, the available sample consists of 1,034 games. The *Success Rate* is the proportion of times that the particular strategy correctly predicted the outcome. *In-Sample* refers to cases where the model is both estimated and tested over the whole sample. *Out-of-Sample* refers to cases where the model is estimated using data from 1989 to 1993 and tested using data from 1994 and 1995. Both *P*-Values are the probability that a strategy of randomly selecting teams will yield a success rate greater than that of the particular strategy being tested. The *Bootstrap P*-Value is based on 1,000 simulations. The *Binomial P*-Value is based on a success probability of 0.5.

an out-of-sample success rate of 58.10%. While this success rate is less than the naïve home-underdog strategy, a bet can be placed on every game. Some probit filter betting strategies yield substantially higher success rates. For example, the strategy of placing bets only when the probit probability of success exceeds 70% performs extremely well yielding an out-of-sample success rate of 73.53%. That is, when the probit model predicts that the probability that the home team will beat the spread is greater than 70%, we bet on the home team, and when the predicted probability that the home team will beat the spread is less than 30%, we bet on the away team. This strategy, however, cannot be implemented on every game since approximately only one game in ten (or one game per week) triggers the probit filter.

4.2.5 Profitability of Betting Strategies In a standard point spread betting market, it is a simple task to evaluate the profitability of betting strategies since the dividend is fixed and known with certainty prior to placing the bet.[17] Recall, however, that the ARL *FootyTAB* market does not pay a fixed dividend, but rather a portion of the total amount bet. Since the eventual *FootyTAB* dividend depends on the support rate for each team and the total prize pool, the profitability of betting strategies depends on the uncertain dividend. We cannot, therefore, conclude that high success rates in predicting game outcomes will necessarily lead to significant profits. Although outcomes appear to be quite predictable, this may attract many winning bets, resulting in a dividend that is just commensurate with the probability of success.

To measure the profitability of ARL betting strategies, we begin by examining the actual support rate for various ranges of probit probabilities. The support rate is the actual proportion of the total amount bet that is placed on the home team. The distribution of support rates is graphed in Figure 1. Clearly, bettors recognise the home team bias in the "expert" spread as evidenced by the high support rates for home teams. In Table 6 we examine the relationship between support rates and actual success rates. For example, for games in which the probit probability of the home team beating the spread is between 65% and 70%, we calculate the average support rate that the home team received. This support rate and knowledge of the 75% payout rate are sufficient to estimate the dividend and the corresponding break-even success rate.

The average support rates in Table 6 suggest a favourite/long shot bias. Teams that are predicted by the probit model to be highly likely to win do not receive commensurately higher support rates. In contrast, teams that are predicted to be relatively less likely to win receive higher support rates. This bias leads to a profitable strategy—bet on strong favourites. The estimated break-even success

17. For example, Gandar, Zuber, O'Brien and Russo (1988) show that the "eleven for ten" payout in NFL betting leads to a 52.4% break-even success rate.

Figure 1

Proportion of Bets Placed on Home Teams

rate for betting on the strongest favourites is 71.86%, while the actual success rate is 75.53%. Recall, however, Table 5 reported that there is approximately only one strong favourite each week.

We conclude, therefore, that there is little evidence of profitable betting strategies in the ARL *FootyTAB* market given the relatively large (25%) commission retained by the TAB. There is an apparent oversupport of low probability teams and undersupport of high probability teams which is consistent with the favourite/longshot bias in racetrack betting markets. The most frequently cited explanation for this phenomenon is that recreational bettors derive utility from capturing the bragging rights associated with backing longshots. While this bias results in an apparently successful strategy of betting on strong favourites, the opportunities for implementing this strategy are uncommon. Further, the high commmission in *FootyTAB* betting probably forces serious gamblers out of this market leaving only recreational bettors.[18]

---

18. Consistent with this conclusion, the total amount bet on the *FootyTAB* pick-the-winners system amounts to a little over $7 million per season (or just over $320,000 per week) which is considerably less than the amount bet on horse-racing, for example.

Table 6

Profitability of ARL Probit Betting Strategies

Model Probability	Average Support Rate	Average Dividend	Break-even Success Rate	Actual Success Rate
< 45	58.19	1.29	77.59	51.61
45 – 50	56.78	1.32	75.73	46.07
50 – 55	54.81	1.37	73.07	54.55
55 – 60	53.96	1.39	71.94	55.80
60 – 65	54.56	1.37	72.74	63.24
65 – 70	52.18	1.44	69.57	65.96
> 70	53.89	1.39	71.86	75.53

Note: 1. The data consist of 1,034 ARL games from 1989 to 1995. *Model Probability* refers to the probability that the home team will beat the spread according to the augmented probit model in Equation (3). The *Average Support Rate* is the actual average proportion of bettors who favoured the home team when the probit probability was within the specified range. The *Average Dividend* is the expected total return on a successful bet, given the support rate for that team and a 75% payout rate. The *Break-even Success Rate* is the estimated proportion of bets that must be successful to break-even, given the dividend for that type of bet. The *Actual Success Rate* is the actual proportion of bets that are successfully classified by the model.

4.3 The Efficiency of the AFL Footywin Market

4.3.1 Predictability of Game Outcomes Recall that AFL *Footywin* betting requires bettors to predict the 12-point bin in which the result will fall. According to the broad view of market efficiency, no betting strategy should consistently generate positive returns. As a preliminary examination of this issue, we calculate the historical probabilities of results falling within each bin and compare them with the average dividend for each bin. That is, we are examining the strategies of always betting on Bin 1 or always betting on Bin 2 and so on. The returns from these naïve strategies are reported in Table 7.

Table 7 reports measures of the significance of the success rate and average return of the various trading strategies. To test the significance of the success rate, the parametric (binomial) p-value is the probability that a random betting strategy would generate more successes than the particular strategy. To test the significance of the average return, the bootstrap p-value is the proportion of simulations in which a strategy of randomly selecting a bin generated a greater return than the particular strategy.

Table 7

Outcomes and Dividends: AFL *Footybet*[1]

Bin	Frequency (Binomial *P*-Value)	Average Dividend (Standard Deviation) [Minimum-Maximum]	Average Return (%) (Bootstrap *P*-Value)
Home 1–12	0.1059 (0.000)	10.73 (3.76) [5.30–25.70]	13.64 (0.049)
Home 13–24	0.0934 (0.001)	9.12 (4.03) [5.20–27.90]	–14.87 (0.388)
Home 25–36	0.0990 (0.000)	7.91 (3.41) [3.90–30.80]	–21.76 (0.612)
Home 37–48	0.0878 (0.008)	10.55 (6.00) [4.30–42.20]	–7.40 (0.223)
Home 49–60	0.0669 (0.747)	14.17 (11.38) [4.70–60.80]	–5.22 (0.186)
Home 60–72	0.0411 (1.000)	21.23 (19.94) [5.80–99.90]	–12.71 (0.336)
Home 73+	0.0878 (0.008)	10.93 (18.82) [1.20–188.30]	–4.03 (0.170)
Draw	0.0091 (1.000)	67.12 (37.25) [21.60–153.50]	–39.19 (0.975)
Away 1–12	0.0843 (0.029)	11.55 (5.33) [5.10–56.40]	–2.63 (0.155)
Away 13–24	0.0815 (0.068)	9.37 (3.93) [5.00–32.80]	–23.57 (0.678)
Away 25–36	0.0857 (0.018)	9.67 (6.10) [4.60–43.30]	–17.15 (0.449)

Table 7 cont'd.

Bin	Frequency (Binomial *P*-Value)	Average Dividend (Standard Deviation) [Minimum-Maximum]	Average Return (%) (Bootstrap *P*-Value)
Away 37–48	0.0481	11.55	–44.45
	(0.999)	(7.78)	(0.996)
		[4.90–62.80]	
Away 49–60	0.0334	17.44	–41.68
	(1.000)	(13.37)	(0.993)
		[5.60–58.50]	
Away 60–72	0.0300	36.49	9.35
	(1.000)	(80.41)	(0.064)
		[7.00–536.20]	
Away 73+	0.0460	14.96	–31.21
	(0.999)	(20.76)	(0.863)
		[2.00–104.00]	

Note: 1. The data consist of 1,434 AFL games from 1987 to 1995. The *Average Dividend* is the payout on a $1 bet. The *Average Return* is the average percentage return over 1,434 bets. A dividend of $536.20 in the Away 60–72 bin occurred in Round 5 of 1994 when Hawthorn upset the previous season's premiers, West Coast Eagles, at Subiaco. Without this result, the average return from betting on this bin is –27.96%, which is not statistically significantly different from the return of a random betting strategy.

In many respects, the results in Table 7 are consistent with evidence relating to other sports betting markets. Consistent with the existence of a home-field advantage, home teams win 58.19% of games (Bins 1–7), while away teams have an unconditional probability of success of only 40.90% (Bins 9–15).

There also appears to be a favourite/long shot bias. Long shots, or bins that have little probability of occurring, tend to be overbet, and favourites, or bins that are relatively more likely to occur, tend to be underbet.[19] For example, whereas drawn games occur 0.91% of the time, the average dividend that is paid on a drawn game is only $67.12. The strategy of betting on draws consistently loses, yielding a gross return of $0.61 for every $1 bet. Conversely, the home 1–12 point bin has a probability of occurring of 10.59%, making it relatively quite likely. The

19. While we have identified these biases ex-post, it seems clear from the history of AFL games *before* our sample period, that draws are much less likely than home team 1–12 point wins.

average gross return from placing a $1 bet on the home team to win by 1–12 points is $1.14. Therefore, it appears that a simple strategy of betting on the home team by 1–12 points yields consistent profits, even after commissions. This provides some preliminary evidence of inefficiency in the AFL *Footywin* market.

We now examine the efficiency of the AFL *Footywin* market in more detail by considering whether available information can be used to predict which bin will result. This is done via an ordered probit model. We begin by defining an ordered categorical variable R_i to represent game outcomes. $R_i = 1$ if the home team wins by 73 points or more, $R_i = 2$ if the home team wins by 60 to 72 points, $R_i = 8$ for a draw, and so on until $R_i = 15$ if the away team wins by 73 points or more. Next, we define a continuous variable Y_i^* to be a linear function of our explanatory variables:

$$\begin{aligned} Y_i^* &= X_i \beta + \varepsilon_i \qquad (4) \\ &= \beta_0 + \beta_1 HP_i + \beta_2 AP_i + \beta_3 HWP_i + \beta_4 AWP_i + \beta_5 HA_i + \varepsilon_i \end{aligned}$$

where ε_i is assumed, without loss of generality, to be distributed standard normal.

In our set of explanatory variables, *HP* represents the percentage of home games won by the home team in the current season, *AP* represents the percentage of away games won by the away team in the current season, *HWP* represents the weighted performance of the home team over the last four games, and *AWP* represents the weighted performance of the away team over the last four games. To calculate the weighted performance, we give a team a score between –7 and +7 for each game according to the bin in which the result falls. A loss by more than 72 points scores –7, a draw scores zero, and a win by more than 72 points scores +7 and so on. These scores are accumulated over the four most recent games, giving a total score between –28 and +28. These variables capture the strength of the recent performance of each team and provide a way of testing whether the market over- or under-compensates for streaks or momentum. Finally, we consider the possibility of a home-ground advantage. The home-ground advantage is now largely removed for games between two Victorian-based teams, with ground rationalisation forcing most teams to share their home ground. We argue that it is generally only when a team travels interstate that it is at a disadvantage to the home team. In this case, the visiting team faces the problems of travel and acclimatisation and must play in front of a particularly parochial crowd. Accordingly, *HA* is an indicator variable that takes a value of zero if the game is between two Victorian-based teams and one otherwise.

The ordered probit technique divides the standard normal density (for ε_i) into fifteen segments, one for each possible bin, as shown in Figure 2 for the case where $X_i \beta = 2$. Next, the ordered probit model posits that $R_i = 1$ if $Y_i^* \le 0$, $R_i = 2$ if $0 < Y_i^* \le a_1$, $R_i = 3$ if $a_1 < Y_i^* \le a_2$, and so on until $R_i = 15$ if $a_{13} < Y_i^*$. Now $Y_i^* \le 0$ if $\varepsilon_i \le -X_i \beta$ and $0 < Y_i^* \le a_1$ if $X_i \beta < \varepsilon_i \le -X_i \beta + a_1$ and so on.

For $\varepsilon_i \sim N(0, 1)$, $Pr(\varepsilon_i \le -X_i \beta) = \Phi(-X_i \beta)$ where $\Phi(\cdot)$ is the normal cumulative distribution function. Similarly, $Pr(-X_i \beta < \varepsilon_i \le -X_i \beta + a_1) = \Phi(-X_i \beta + a_1) - \Phi(-X_i \beta)$ and so on. Hence, the contribution to the log-likelihood function from a particular game *i* is

$$l_i = \begin{cases} \ln[\Phi(-X_i\,\beta)] & if\, R_i = 1 \\ \ln[\Phi(-X_i\,\beta + a_1) - \Phi(-X_i\,\beta)] & if\, R_i = 2 \\ \vdots & \\ \ln[1 - \Phi(-X_i\,\beta + a_{13})] & if\, R_i = 15. \end{cases} \qquad (5)$$

Consequently, the log-likelihood function to be maximised is

$$\sum_{i=1}^{N} l_i. \qquad (6)$$

We use maximum likelihood estimation to choose parameters $\beta_j;\ j = 0, \ldots, 5$ which maximise the probability that our model produces the game outcomes R_i; $i = 1, \ldots, N$ that we observe. The optimization was performed using the BFGS algorithm in GAUSS. Table 8 reports the estimates for this model and Figure 2 is drawn according to the estimates of $a_{1,} \ldots, a_{13}$.

Statistically, the ordered probit model performs well, with all of the coefficients reaching statistical significance and 14.26% of bins being correctly predicted by the model. On average, a strategy of betting on randomly selected bins yields a success rate of only 6.67% (that is, one bin out of fifteen). In 1,000 simulations of this randomising strategy, none yielded success rates above 14.26% meaning that the bootstrap *p*-value for the hypothesis that our model performs significantly better than random chance is less than 0.001. Furthermore, the model’s prediction is within one bin either side of the actual result 26.37% of the time, whereas the corresponding success rate for the randomising strategy is only 20.00% (three bins out of fifteen). Here the bootstrap *p*-value is also less than 0.001, again indicating that our model’s success is statistically significant.

Moreover, the signs of the coefficients are as expected. The coefficients on *HP* and *AP* are positive and negative, respectively. This indicates that a larger home team win is expected when the home team has a good home record and the away team has a poor record on the road over the current season. A similar explanation pertains to the coefficients of *HWP* and *AWP* which are positive and negative, respectively. The coefficient of *HA* is positive, which indicates that a larger home-team win is expected when there is an interstate home-ground advantage. The significance of β_0 indicates the additional success enjoyed by home teams.

4.3.2 Profitability of Betting Strategies While our model appears to exhibit statistical significance, high success rates in predicting bins does not guarantee a positive return. Since dividends are determined in parimutuel fashion, we must evaluate the economic significance by examining the profitability of betting strategies. In addition to simple betting strategies (for example, always backing the home team), we examine whether our ordered probit model can be employed in a betting strategy that generates significant profits. The results of various betting

Table 8

Ordered Probit Results: AFL[1]

Parameter	Estimate	Std Error
β_0	1.7487*	0.1254
β_1	0.0050*	0.0015
β_2	–0.0059*	0.0013
β_3	0.0190*	0.0039
β_4	–0.0213*	0.0038
β_5	0.1857*	0.0628
% Correct Predictions:	14.26	
% Predictions Correct To Within ± 1 Bin:	26.37	

Note: 1. The data consist of 1,164 AFL games from 1987 to 1995. * = significant at 5%. The augmented probit model is $Y_i^* = \beta_0 + \beta_1 HP_i + \beta_2 AP_i + \beta_3 HWP_i + \beta_4 AWP_i + \beta_5 HA_i + \varepsilon_i$. Some early season observations are lost due to the construction of our recent-record variables. We have reëstimated the model using data from the end of the previous season to construct these recent-form variables. The results are virtually identical and are available on request.

strategies are reported in Table 9. Bootstrap and parametric *p*-values are also reported. The sample sizes used in calculating *p*-values correspond to the number of bets placed. For example, a filter-type strategy may only be triggered 10% of the time so that we place only 140 bets over our sample period. Our *p*-values are then calculated on the basis of a sample size of 140.

The naïve strategies of betting on home teams only (all bins) and away teams only (all bins) yield negative returns.[20] Since probit-based betting strategies require knowledge of the model’s coefficients, we parameterise the model on data from 1987 to 1992 and then apply the estimates to out-of-sample data from 1993 to 1995. The results in Table 9 indicate that some strategies generate significant positive returns. If bets are placed on the predicted bin, the strategy yields an

20. Other naïve strategies not reported in the paper (including backing a particular team consistently, backing interstate teams consistently, and backing a particular team in Friday night games only) also yield negative returns.

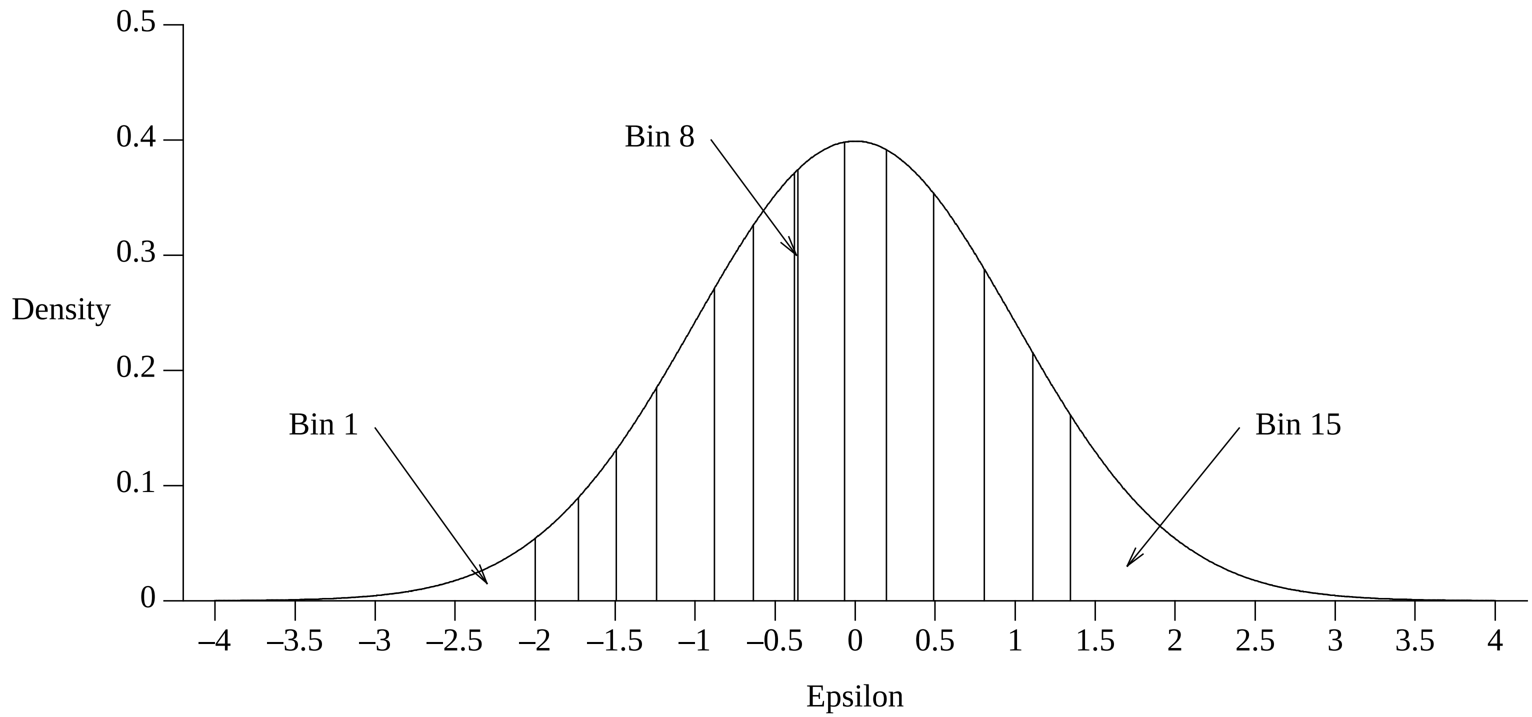

Figure 2

Ordered Probit Model

Table 9

Profitability of AFL Betting Strategies[1]

Betting Strategy	Proportion of Games	Cost of Bet	Success Rate (Binomial *P*-Value)	Average Dividend	Average Return (%) (Bootstrap *P*-Value)
Bet on Home Team (Bins 1–7)	100.00	7.00	58.19 (0.000)	11.13	–7.48 (0.223)
Bet on Away Team (Bins 9–15)	100.00	7.00	40.86 (1.000)	12.52	–26.91 (0.769)
Bet on Predicted Bin					
In-Sample	100.00	1.00	14.26 (0.000)	7.94	13.26 (0.060)
Out-of-Sample	100.00	1.00	14.44 (0.000)	8.50	22.74 (0.130)
Bet on Two Most Likely Bins					
In-Sample	100.00	2.00	25.26 (0.000)	8.00	2.12 (0.114)
Out-of-Sample	100.00	2.00	24.55 (0.000)	8.14	–0.11 (0.246)
Bet on Predicted Bin if Probability is Greater Than 25%					
In-Sample	3.87	1.00	31.11 (0.000)	2.61	–18.89 (0.371)
Out-of-Sample	6.13	1.00	17.65 (0.000)	4.50	–20.58 (0.382)
Bet on Predicted Bin and Bin on Either Side[2]					
In-Sample	100.00	2.70	29.30 (0.000)	8.43	–8.24 (0.212)
Out-of-Sample	100.00	2.67	29.96 (0.002)	9.07	1.82 (0.231)

average return of 22.74% out-of-sample. The bootstrap *p*-value (0.130) indicates that it is unlikely that a random betting strategy could produce a return greater than 22.74%.[21] Some betting strategies utilising probit filters also generate positive

21. That is, 13% of 1,000 simulations, where bets were placed on randomly selected teams, generated average returns greater than 22.74%. The sample size for each simulation matched the number of games in the out-of-sample period. While this cannot be interpreted as conclusive evidence of predictability, it does provide some indication that the profitability of this strategy is not due solely to chance.

Table 9 cont'd.

Betting Strategy	Proportion of Games	Cost of Bet	Success Rate (Binomial *P*-Value)	Average Dividend	Average Return (%) (Bootstrap *P*-Value)
Bet $3 on Predicted Bin and $1 on Bin on Either Side[2]					
In-Sample	100.00	4.70	29.30	16.17	0.92
			(0.000)		(0.117)
Out-of-Sample	100.00	4.67	29.96	17.26	10.78
			(0.002)		(0.187)
Bet on Extreme Bin if Win Probability Exceeds 70%					
In-Sample	31.36	1.00	19.18	6.34	21.62
			(0.000)		(0.057)
Out-of-Sample	36.10	1.00	19.00	8.08	53.50
			(0.000)		(0.100)

Notes: 1. The data consist of 1,164 AFL games from 1987 to 1995. Some early season observations are lost due to the construction of our recent-record variables. We have reëstimated the model using data from the end of the previous season to construct these recent-form variables. The results are virtually identical and are available on request. Where a strategy requires knowledge of model coefficients, we calculate success rates and returns over an out-of-sample period, 1993–1995, where the model is parameterised using data from 1987–1992.

2. Draws are ignored draws when considering neighbouring bins. That is, if the home team is predicted to win by 1-12 points, the neighbouring bins are considered to be a home 13–24 point win and an away 1–12 point win.

returns. For example, betting $3 on the predicted bin and $1 on the neighbouring bins yields an average return of 10.78% out-of-sample (bootstrap *p*-value = 0.187).

Having already noted a possible favourite/long shot bias, Table 9 reports that a betting strategy based on this bias is profitable. Betting on the extreme bin when the probit model estimates the probability of a win (by any margin) exceeds 70% generates a 53.50% return out-of-sample (bootstrap *p*-value = 0.100). Note, however, that probit filters restrict the bettor from betting on *every* game. For example, the favourite/long shot strategy is only available for 36.10% of the out-of-sample games.

We are cautious, however, in interpreting these results as clear evidence of market inefficiency. While the results in Table 9 indicate that significant positive returns are sometimes available, we have three caveats. First, the results may be specific to our sample. Second, we have reported favourable results and other permuatations of strategies and model design were tried. Hence, we are to some extent guilty of data-mining, even though the apparently profitable strategies are straightforward, intuitive and analagous to results reported in other sports betting markets. Third, as the dividend is determined by reference to the size of the betting pool, any substantial bet would influence the final dividend. Therefore, exploitation of the apparent inefficiency may cause the disappearence of the profitable strategies.[22] The analogy to financial markets is strong: arbitrage opportunities are exploitable for price takers, but the actions of price makers may change the price such that the arbitrage is no longer available.

5. Conclusions

From our analysis, it is clear that game outcomes are predictable in the sense that predictive models perform significantly better than random selection, and that some betting strategies generate significant positive returns in both the ARL and AFL. In the ARL, there is a bias in the posted spreads towards home teams and underdogs. Due to the large transaction costs, however, it is difficult to consistently earn economic profits by exploiting this bias. One strategy, however, appears to be profitable. A probit-based betting strategy backing strong favourites has success rates of over 75%, which we estimate to be sufficient to generate significant profits. Although this strategy also earns significant positive returns out-of-sample, we are cautious about concluding that the ARL *FootyTAB* is inefficient based upon this single result.

In the AFL, a naïve strategy of backing the home team to win by 1–12 points generates an average return of 13.63%. Our ordered probit model is quite successful in predicting the range in which the result of a game will fall. When tested out-of-sample, betting strategies using probit-based filters have success rates high enough to generate significant positive profits. Betting on the predicted bin yields a return over 22%, and betting on the extreme bin when the probit-based win probability exceeds 70% yields an average return of 53%. This evidence is prime facie indicative of market inefficiency. We have, however, raised concerns as to whether this apparent inefficiency is exploitable.

Finally, we do not claim that the probit model captures all relevant information. We have excluded many variables that could have information content (such as weather conditions, history between the two teams, position on the ladder, crowd size, identity of coaches, teams' records at the ground). Rather, our purpose has been to examine market efficiency in the context of prior systematic biases documented in the literature and drawing analogies to financial markets.

22. Indeed, it is likely that the dividend would be quite sensitive to a large bet. For years in which we have data on the size of the betting pool, the average pool per game is $5,835, $5,415, $5,660 and $7,549 for 1985, 1986, 1987 and 1990, respectively.

Extending this type of model to examine potential new biases along these lines may be a fruitful area for further research.

(Date of receipt of final typescript: November 1995.)

References

Ali, M.M., 1977, Probability and utility estimates for racetrack bettors, *Journal of Political Economy*, 85, 803–815.

Alpert, W. and H. Raiffa, 1982, A progress report on the training of probability assessors, in: D. Kahneman, P. Slovic, and A. Tversky, *Judgement Under Uncertainty: Heuristics and Biases* (Cambridge, Cambridge University Press).

Bird, R. and M. McCrae, 1987, Tests of the efficiency of racetrack betting using bookmaker odds, *Management Science*, 33, 1552–1562.

Bird, R., M. McCrae, and J. Beggs, 1987, Are gamblers really risk takers? *Australian Economic Papers*, December, 237–253.

Boudoukh, J., M.P. Richardson, and R.F. Whitelaw, 1994, A tale of three schools: insights on autocorrelaion of short-horizon stock returns, *Review of Financial Studies*, 7, 539–573.

Brown, W.O. and R.D. Sauer, 1993, Fundamentals or noise? Evidence from the professional basketball betting market, *Journal of Finance*, 48, 1193–1209.

Camerer, C.F., 1989, Does the basketball market believe in the “hot hand”? *The American Economic Review*, 79, 1257–1261.

Clarke, R. and M. Statman, 1994, Growth, value, good and bad, *Financial Analysts Journal*, Nov-Dec, 82–86.

De Bondt, W. and R.H. Thaler, 1985, Does the stock market overreact? *Journal of Finance*, 40, 793–805.

De Bondt, W. and R.H. Thaler, 1987, Some further evidence on investor overreaction and stock market seasonality, *Journal of Finance*, 42, 557–581.

Denis, D.J. and D.K. Denis, 1995, Performance changes following top management dismissals, *Journal of Finance*, 50, 1029–1057.

Fama, E.F., 1970, Efficient capital markets: a review of theory and empirical work, *Journal of Finance*, 25, 383–417.

Fama, E.F., 1991, Efficient capital markets II, *Journal of Finance*, 46, 1575–1617.

Gandar, J., R. Zuber, T. O’Brien, and B. Russo, 1988, Testing rationality in the point spread betting market, *Journal of Finance*, 43, 995-1008.

Gray, P.K. and S.F. Gray, 1995, Testing market efficiency: evidence from the NFL sports betting market, Working Paper, Fuqua School of Business, Duke University.

Golec, E.K. and J.A. Tamarkin, 1991, The efficiency of sports betting markets, *Journal of Financial Economics*, 4, 653–665.

Heisler, J., 1994, Recent research in behavioural finance, *Financial Markets, Institutions and Instruments*, 3, 76–104.

Johnson, W.B., R.P. Magee, N.J. Nagarajan, and H.A. Newman, 1985, An analysis of the stock price reaction to sudden executive deaths, *Journal of Accounting and Economics*, 7, 151–174.

Kahneman, D., P. Slovic, and A. Tversky, 1982, *Judgement Under Uncertainty: Heuristics and Biases* (Cambridge, Cambridge University Press).

Stael von Holstein, C-A., 1971, Two techniques for assessment of subjective probability distributions—an experimental study, *Acta Psychologica*, 35, 478–494.

Swidler, S. and R. Shaw, 1995, Racetrack wagering and the uniformed bettor: a study of market efficiency, *The Quarterly Review of Economics and Finance*, 35, 305–314.

Thaler, R.H., ed., 1994, *Advances in Behavioural Finance* (New York, Sage Foundation).

Thaler, R.H. and W.T. Ziemba, 1988, Parimutuel betting markets: racetracks and lotteries, *Journal of Economic Perspectives*, 2, 161–174.

Tuckwell, R.H., 1983, The thoroughbred gambling market: efficiency, equity and related issues, *Australian Economic Papers*, 22, 106–118.

Warner, J., R. Watts, and K. Wruck, 1988, Stock prices and top management changes, *Journal of Financial Economics*, 20, 461–492.

Weisbach, M., 1988, Outside directors and CEO turnover, *Journal of Financial Economics*, 20, 431–460.

Winkler, R., 1967, The assessment of prior distributions in Baysian analysis, *Journal of the American Statistical Association*, 62, 776–800.

Woodland, L.M. and B.M. Woodland, 1991, The effects of risk aversion on wagering: point spread versus odds, *Journal of Political Economy*, 99, 638–653.

Woodland, L.M. and B.M. Woodland, 1994, Market efficiency and the favorite-longshot bias: the baseball betting market, *Journal of Finance*, 49, 269–279.

Zuber, R.A., J.M. Gandar, and B.D. Bowers, 1985, Beating the spread: testing the efficiency of the gambling market for national football league games, *Journal of Political Economy*, 93, 800–806.