



Testing Market Efficiency: Evidence from the NFL Sports Betting Market

Philip K. Gray; Stephen F. Gray

The Journal of Finance, Vol. 52, No. 4 (Sep., 1997), 1725-1737.

Stable URL:

<http://links.jstor.org/sici?sici=0022-1082%28199709%2952%3A4%3C1725%3ATMEEFT%3E2.0.CO%3B2-M>

The Journal of Finance is currently published by American Finance Association.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/afina.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Testing Market Efficiency: Evidence From The NFL Sports Betting Market

PHILIP K. GRAY and STEPHEN F. GRAY*

ABSTRACT

This article examines the efficiency of the National Football League (NFL) betting market. The standard ordinary least squares (OLS) regression methodology is replaced by a probit model. This circumvents potential econometric problems, and allows us to implement more sophisticated betting strategies where bets are placed only when there is a relatively high probability of success. In-sample tests indicate that probit-based betting strategies generate statistically significant profits. Whereas the profitability of a number of these betting strategies is confirmed by out-of-sample testing, there is some inconsistency among the remaining out-of-sample predictions. Our results also suggest that widely documented inefficiencies in this market tend to dissipate over time.

SPORTS BETTING MARKETS PROVIDE an ideal arena for testing market efficiency. In many financial markets, it is difficult to formulate a direct test of market efficiency since the true value of a particular investment is never revealed. Therefore, efficiency tests have tended to focus on the predictability of asset *returns*. In sports betting markets, however, 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* realizations are available soon after each bet is placed. This structure generates a relatively simple direct test of market efficiency.

The question of whether organized sports betting markets are efficient has received recent attention in the market efficiency literature. Golec and Tarmarkin (1991), for example, document that spreads set in the National Football League (NFL) betting market are systematically biased predictors of actual results. Findings such as this are sometimes offered as evidence of inefficiency in the sports betting market, although it is not always clear that this bias can be exploited via a profitable trading strategy. 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 do, however, establish evidence of *economic* inefficiency by demonstrating the profitability of various betting strategies. Thaler and Ziemba

* Philip K. Gray is from Queensland University of Technology and the Australian Graduate School of Management. Stephen F. Gray is from the University of Queensland and Duke University. We thank Barry Oliver, Peter Whelan, an anonymous referee, and participants of the 1995 AAANZ conference for valuable comments and suggestions. Funding from Coopers and Lybrand is gratefully acknowledged.

(1988) review the early literature on racetrack betting markets, noting the consistent favorite/longshot bias. The expected return from betting on favorites is significantly greater than the expected return from betting on longshots. Woodland and Woodland (1994) document that the favorite/longshot bias in racetrack betting exists in reverse for baseball bettors, but that no betting strategy admits 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.

This article advances the literature that examines the efficiency of sports betting markets in a number of directions. First, we reexamine the *statistical efficiency* of NFL betting markets using a discrete-choice probit model rather than the ordinary least squares (OLS) regression methodology used previously in the literature. The sensitivity of OLS to extreme outliers may make OLS results difficult to interpret. This is because the return on a bet is the same whether a team beats the spread by 1 point or by 30 points. Whereas OLS effectively gives more weight to the more extreme result, the probit model sees both results as “wins” and treats them equally.

Second, we examine the *economic efficiency* of the market from the perspective of trading strategies. The existence of consistent statistical biases in point spreads is not, in itself, evidence of inefficiency. In the strict sense, market inefficiency requires that trading strategies can exploit biases to earn consistent profits. This article examines the profitability of trading strategies based on biases documented both in this article and in previous research. Also, our probit model is ideally suited to examine issues such as the probability of team A beating the spread against team B, conditional on available information. Consequently, we can examine trading rules that involve, for example, placing a bet only when the conditional probability of success exceeds X percent (some filter). A series of in-sample tests demonstrates that betting strategies based on such probit filters can generate statistically significant profits. Whereas the profitability of some of these betting strategies is confirmed by out-of-sample testing, there is some inconsistency among the remaining out-of-sample predictions. This leads us to question whether the apparent inefficiencies will hold up over the long run. Moreover, our results suggest that widely documented inefficiencies in this market (such as the home-underdog bias) tend to dissipate over time.

Third, we consider a number of explanatory variables that have strong analogies to financial markets and behavioral research. For example, a team's recent record is analogous to recent stock market performance. The growing literature on momentum/contrarian investment is concerned with whether it is possible to earn abnormal returns by looking at the recent performance of a stock. Can we condition on “runs” in stock prices to say something about the (conditional) distribution of future stock prices? The analogy to sports betting markets is quite strong—can we condition on recent winning or losing “streaks” to say something about the (conditional) distribution of future outcomes? Moreover, there is a growing behavioral literature on the perceived

existence of runs in random series. Camerer (1989) examines this issue in the context of sports betting markets. Even if there are no runs in a series of game outcomes, if the market perceives that there are runs, spreads could be biased, in which case profitable betting strategies may exist.

The balance of the article is organized as follows. Section I discusses the structure of the NFL betting market and the conditions necessary for an efficient market. In Section II we outline our econometric methods and report the results. Section III concludes.

I. Efficiency in Point Spread Betting Markets

The structure of the Las Vegas football betting market is relatively simple. The betting establishment posts a spread at which it is willing to transact. The bettor chooses which team to favor and receives a payoff of \$21 for every \$11 bet in the event that his team beats the spread. This is known as the “11 for 10 rule” which requires that bettors must pick winners in 52.4 percent of bets to break even.¹ The marginal dollar is commission or “vigorish,” enabling the betting house to earn a profit. Assuming that the betting house doesn’t want to take a position on the outcome of the game, their incentive is to set the spread so that equal numbers of bettors favor each team. If this is done, the betting house is guaranteed a profit of \$1 for every \$22 bet. If not, the trading house bears risk. Therefore, the game-time spread should reflect the aggregate information in the market—it is a market clearing price. To examine whether the market efficiently incorporates all publicly available information into the price, we can simply test whether any publicly available information is useful in predicting game outcomes, beyond the information contained in the spread.

If the market is efficient, the spread is an unbiased predictor of the result. Accordingly, previous statistical tests of efficiency have centered around whether the spread captures all publicly available information.² Typically, this is done in an OLS framework:

$$R_i = \beta_0 + \beta_1 S_i + \beta_2' X_i \quad (1)$$

where R_i is the actual spread (the outcome of the game), S_i is the game time spread, and X_i represents all other conditioning information. If the spread captures all relevant information, $\beta_1 = 1$, $\beta_0 = 0$, and all of the elements of β_2 are zero. Market efficiency is typically examined by testing this joint hypothesis. In particular, if β_2 is significantly different from zero, the information in X_i can be used to predict whether or not the spread will be beaten.

Even if statistical tests of efficiency indicate that the spread is consistently biased, the decisive test of efficiency is whether such a bias yields positive profits after commissions. In the market efficiency literature, especially as it relates to sports betting, two definitions of efficiency have been used. The

¹ See, for example, Vergin and Scriabin (1978) and Gandar, Zuber, O'Brien, and Russo (1988).

² See, for example, Golec and Tamarkin (1991).

narrow view, which is primarily of academic interest, posits that the expected loss from any betting strategy should approximate 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 commission) on average. In this article, as in the majority of preceding work, we focus on the latter definition. We test for efficiency by searching for betting strategies that yield significantly positive returns, on average. If the market is efficient, there exists no such strategy, as the spread captures all relevant information.

II. Econometric Methods and Results

A. Data

The NFL data was obtained from the *Handicapper's Pointspread Notebook* and the *Goldsheet*, published by Nation-Wide Sports Publications, Inc. The data consist of Las Vegas game time spreads, outcomes, game dates, and game locations for the 28 NFL teams from 1976 to 1994, a total of 4219 games. In this sample, there are 96 "no-bets," approximately five per season. These are games for which the result exactly matches the spread so that neither team beats the spread, in which case all bets are returned. The outcome is as if no bets were placed, so all of these games are removed from the sample, leaving 4123 usable observations. For the majority of games, the spread is not a whole number (e.g., 4.5 points) so that a no-bet is impossible.

For each game, the point spread and the result can be defined relative to either team. Suppose team A is given 10.5 points spread against team B. Then the spread is +10.5 from the perspective of team A and -10.5 from the perspective of team B. We denote the team from whose perspective the spread and result is defined as the *team of record*. There is no single correct way of choosing the team of record, and three different methods have been used.³ First, the team of record can be defined to be the favorite. Second, the team of record can be defined to be the home team. Third, the team of record can be chosen randomly, avoiding any systematic effects. Assigning the team of record on the basis of a specific attribute may uncover a conditional bias, such as favorites and away teams receiving too high a spread, for example. Randomly assigning the team of record is unlikely to uncover any bias, since there is nothing systematic about the teams of record.

Table I contains summary statistics relating to game scores and winning margins. Summary statistics of point spreads and game outcomes are reported in Table II. The team of record is defined in the three different ways described above. There appears to be some overconfidence in the favorite, as the favorite wins by less than expected (according to the spread). On average, the favorite gives up 5.62 points spread, but wins by only 5.20 points. Also, the value of the home field advantage appears to be discounted. On average, the home team

³ See Golec and Tamarkin (1991).

Table I
NFL Match Statistics

This table is based on data from 4219 NFL games from 1976 to 1994. Winning Score refers to the number of points scored by the winning team. Losing Score refers to the number of points scored by the losing team, and Winning Margin refers to the difference in scores.

Statistic	Mean	Std. Dev.	Range
Winning score	25.94	9.11	3-62
Losing score	14.21	8.00	0-48
Winning margin	11.73	9.31	0-59

Table II
Summary Statistics for NFL Games

This table contains summary statistics for 4219 National Football League (NFL) games from 1976 to 1994. The Point Spread columns represent expected outcomes, based on the spread. The Outcome columns refer to actual results. The Favorite column refers to the average points scored by the favorite in excess of the underdog. The Home column refers to the average points scored by the home team in excess of the away team. The Random column refers to the average points scored by a randomly selected team in excess of the opponent.

Parameter	Favorite	Home	Random
Point Spread			
Mean	5.62	2.56	0.07
Std. Dev.	3.68	6.21	6.72
Outcome			
Mean	5.20	2.99	0.20
Std. Dev.	13.92	14.68	14.98
Difference			
Mean	0.42	-0.43	-0.13
Std. Dev.	13.36	13.49	13.50

gives up 2.56 points, but wins by 2.99 points. Obviously, there are no systematic biases when the team of record is defined randomly.

B. Statistical Tests of Market Efficiency

These preliminary results are suggestive of potential systematic biases, or inefficiencies. To formally examine whether the market efficiently incorporates all publicly available information into the spread, we can simply test whether any publicly available information is useful in predicting game outcomes, beyond the information contained in the spread. In this article, we use a probit model rather than the OLS methodology that has been used previously. 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 was large, therefore, have a greater impact on the slope coefficient. In the NFL betting market, the only relevant question is *whether* a certain team beat the spread.

Table III

Parameter Estimates—Simple Probit Model

The simple probit model is $Y_i^* = b_0 + b_1\text{HOME}_i + b_2\text{FAV}_i + \varepsilon_i$, where $Y_i = 1$, indicating that the team of record beat the spread, if $Y_i^* > 0$ and $Y_i = 0$ otherwise. The data consist of 4123 National Football League (NFL) games from 1976 to 1994. Standard errors are reported in parentheses.

Parameter	Estimate
b_0	0.0615 (0.0445)
b_1	0.0750 (0.0574)
b_2	-0.1532 (0.0263)
Correct Predictions	52.51%

The *amount* by which the team beat the spread is irrelevant. Therefore, rather than using the winning or losing margin itself as our dependent variable, we use a discrete choice variable, which is defined as follows:

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

Statistical efficiency requires that no observable information can help to predict Y_i in a statistically significant sense.

We begin by examining the home team and underdog biases documented for the NFL by Golec and Tamarkin (1991) and suggested by Table II. In particular, we examine the probit model:

$$Y_i^* = b_0 + b_1\text{HOME}_i + b_2\text{FAV}_i + \varepsilon_i \quad (2)$$

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 is defined above and

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

$$\text{FAV}_i = \begin{cases} 1 & \text{if the team of record is the favorite} \\ 0 & \text{otherwise.} \end{cases}$$

The results of estimation of this probit model appear in Table III.

The results for this probit model indicate that home teams and underdogs are more likely to beat the spread than are away teams and favorites. The coefficient on our home team dummy variable (b_1) is approaching significance, and the coefficient on our favorite dummy variable (b_2) is statistically significant. A likelihood ratio test (LRT) of the joint significance of b_1 and b_2 yields a statistic of 15.41 (χ_2^2) which is significant at any usual level. The probit model

correctly predicts 52.51 percent of NFL game outcomes, which is barely enough to warrant pursuing this strategy given the transaction cost.⁴

Next, we expand the set of conditioning information to capture the effect of perceived runs in a series of game outcomes. We examine whether the recent past performance of the teams is efficiently captured by the spread. The motivation for examining this effect is twofold. First, the momentum/contrarian investment literature posits that the past performance of a stock provides a signal of demand for that stock and hence future performance. Second, there is a strand of the applied psychology and economics literature that demonstrates that perceived “streaks” by sports teams really don’t exist. For example, Camerer (1989) finds some evidence of a belief in the “hot hand” in spreads for professional basketball games. Teams with long winning streaks tend to perform worse than expected, reverting to the mean rather than continuing their streak.

We examine this issue by estimating the following probit model:

$$Y_i^* = b_0 + b_1HWP_i + b_2AWP_i + b_3HL4_i + b_4AL4_i + b_5FAV_i + \varepsilon_i \quad (3)$$

where HWP_i is the overall winning percentage (relative to the spread) of the home team in the current season, AWP_i is the overall winning percentage of the away team (relative to the spread) in the current season, $HL4_i$ is the number of times the home team has beaten the spread in the last four games played (a number between 0 and 4), $AL4_i$ is the number of times the away team has beaten the spread in the last four games played (a number between 0 and 4), and FAV_i takes a value of 1 when the home team is the favorite and 0 otherwise. The construction of these “recent form” or “streak” variables results in a loss of the first four observations for each team in each season. The resulting sample size is 3068 games. In this case, we define the team of record to be the home team to aid in the interpretation of results and to economize on parameters. Therefore Y_i takes a value of 1 when the home team beats the spread. This model can, therefore, be thought of as a tool to predict the probability that the home team will beat the spread. Table IV contains results of this augmented probit model.

The coefficients on HWP and AWP are positive and negative respectively. Although b_1 does not reach statistical significance the LRT statistic for the joint significance of these coefficients is 6.14 (χ_2^2 under the null) which is significant at the 95 percent level. Moreover, since these variables are naturally related, we offer a joint interpretation. In particular, a team that has been performing well (relative to expectations), over the course of the season to date is more likely to beat the spread than a team that has been performing poorly. This is consistent with the idea that the market is slow to realize that a particular team is having a particularly good (or bad) year.

Moreover, the coefficients on HL4 and AL4 are negative and positive respectively—the reverse of the coefficients discussed above. Again, although only b_4

⁴ Recall that under the 11 for 10 rule, a winning percentage of 52.4 is required to break even.

Table IV
Parameter Estimates—Augmented Probit Model

The augmented probit model is $Y_i^* = b_0 + b_1\text{HWP}_i + b_2\text{AWP}_i + b_3\text{HL4}_i + b_4\text{AL4}_i + b_5\text{FAV}_i + \varepsilon_i$, where $Y_i = 1$, indicating the home team beat the spread, if $Y_i^* > 0$ and $Y_i = 0$ otherwise. The data consist of 4123 National Football League (NFL) games from 1976 to 1994. In fitting the augmented probit model, games from early in the season are lost due to the construction of our “recent form” variables so that the available sample consists of 3068 games. Standard errors are reported in parentheses.

Parameter	Estimate
b_0	0.2446 (0.0623)
b_1	0.0008 (0.0019)
b_2	-0.0036 (0.0016)
b_3	-0.0575 (0.0450)
b_4	0.0753 (0.0369)
b_5	-0.1659 (0.0478)
Correct Predictions	54.56%

is significant individually, the LRT statistic for their joint significance is 12.27 (χ_2^2 under the null) which is significant at the 95 percent level. Since these variables are also naturally related, we again offer a joint interpretation. In particular, a team that has been performing well (relative to expectations), over the last four games is less likely to beat the spread than a team that has been performing poorly over the last four games. This is consistent with the idea that the market overreacts to recent form, discounting the performance of the team over the season as a whole.

The coefficient on the variable FAV is negative and statistically significant. This is consistent with our earlier finding that home-team underdogs are relatively likely to beat the spread. The negative coefficient here indicates that the less favored the home team is, the more likely it is to beat the spread. Finally, note that the constant term, b_0 , is positive, suggesting that home team underdogs are more likely to beat the spread.

The full probit model correctly predicts 54.56 percent of NFL game outcomes, suggesting that it could be used profitably in the NFL betting market. The economic efficiency of this market is examined below.

C. Specification Tests

To further document the performance of our probit model, we compare the model probability of success with the actual success rate. For example, we take all games for which the probit model reports that the probability that the team

Table V
Probit Probabilities and Actual Success Rates

Probit Probability refers to the probability that the team of record will beat the spread, generated from the augmented probit model. Proportion refers to the proportion of games in the sample, and Success Rate refers to actual success rates in the sample. The data consist of 3068 National Football League (NFL) games from 1976 to 1994.

Probit Probability	Proportion (%)	Success Rate (%)
$p < 42.5$	0.68	42.86
$42.5 \leq p < 45.0$	5.25	38.51
$45.0 \leq p < 47.5$	14.90	46.39
$47.5 \leq p < 50.0$	19.85	47.29
$50.0 \leq p < 52.5$	22.29	54.24
$52.5 \leq p < 55.0$	19.56	54.17
$55.0 \leq p < 57.5$	11.05	56.05
$57.5 \leq p$	6.42	56.35

of record will beat the spread is between 50 percent and 52.5 percent. We then report the proportion of successes for these games in Table V.

From Table V it is clear that the probabilities generated by the model correspond quite closely to the actual frequency of success. It is only for the case of extreme probabilities that the model and actual probabilities are not closely aligned. In these cases, however, the sample sizes are relatively small. In the 45 percent to 55 percent range, which covers the majority of the sample, the probabilities from the probit model are remarkably accurate.

D. Economic Tests of Market Efficiency

This section considers the economic efficiency of the NFL betting market by examining a number of trading strategies directly. The results of Table V suggest a novel betting strategy: place bets only on those games where the probability of success is greater than X percent. Under this rule, X is used as a filter, so that bets are placed only when there is a better than even chance of success. Table VI reports the success rate and profitability of this strategy, and some other, simple betting rules. The average return is computed using the 11-for-10 rule, where a successful \$11 bet collects a gross payout of \$21. Table VI also reports two *p*-values, which examine the predictability of game outcomes and the profitability of betting strategies, respectively. Predictability is examined by assessing the significance in relation to a binomial distribution with $p = 0.5$. The resulting *p*-value is the probability that a strategy of betting on a randomly selected team would yield a success rate higher than the observed success rate of the particular strategy. Profitability is examined in relation to a binomial distribution with $p = 0.524$ incorporating the 11-for-10 rule.

In Table VI, the simple betting strategies (i.e., betting on home teams, underdogs, or home-underdogs) are examined over the entire sample, and over the reduced sample that is used for constructing probit estimates. This allows

Table VI
Tests of the Profitability of Betting Strategies

The Total Sample consists of 4123 National Football League (NFL) games from 1976 to 1994. 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 3068 games. This is referred to as the Probit Sample. The Success Rate is the proportion of times that the particular strategy correctly predicts the outcome. The Average Return is the net percentage return per dollar bet, incorporating the 11-for-10 rule. In-Sample refers to cases where the model is both estimated and tested over the probit sample. Out-of-Sample refers to cases where the model is estimated using data from 1976 to 1992 (2727 games) and tested using data from 1993 and 1994 (341 games). Predictability is examined by assessing the significance in relation to a binomial distribution with $p = 0.5$. The resulting p -value is the probability that a strategy of betting on a randomly selected team would yield a success rate higher than the observed success rate of the particular strategy. Profitability is examined in relation to a binomial distribution with $p = 0.524$ incorporating the 11-for-10 rule.

Betting Strategy	Proportion of Games	Success Rate	Average Return	P -Value ($p = 0.5$)	P -Value ($p = 0.524$)
Bet on underdog					
Total sample	98.04	52.60	0.41	0.001	0.401
Probit sample	98.11	52.16	-0.42	0.009	0.604
Bet on home team					
Total sample	100.00	50.42	-3.73	0.293	0.994
Probit sample	100.00	51.11	-2.43	0.110	0.924
Bet on home team underdog					
Total sample	31.70	54.63	4.29	0.000	0.054
Probit sample	31.75	54.83	4.67	0.001	0.065
Bet on predicted team					
In-sample	100.00	54.56	4.17	0.000	0.008
Out-of-sample	100.00	56.01	6.93	0.013	0.091
Bet on teams with probit probability of success exceeding 52.5%					
In-sample	57.86	55.32	5.62	0.000	0.007
Out-of-sample	52.20	47.19	-9.91	0.773	0.918
Bet on teams with probit probability of success exceeding 55%					
In-sample	23.40	57.38	9.55	0.000	0.004
Out-of-sample	20.82	50.70	-3.20	0.500	0.657
Bet on teams with probit probability of success exceeding 57.5%					
In-sample	7.11	56.42	7.71	0.029	0.118
Out-of-sample	5.28	61.11	16.67	0.240	0.308

for comparison with previous work and with the probit results below. Table VI shows that the success rate of betting on home teams and underdogs is marginally above 50 percent. Since neither strategy consistently exceeds the 52.4 percent break-even rate, returns are negative on average. The average return reported in Table VI is the net percentage return per dollar bet, incorporating the 11-for-10 rule. The strategy of betting on home team underdogs yields a success rates above 54 percent, which is significantly greater than the break-even rate of 52.4 percent, so this strategy would have been profitable over our sample period. The net return from this strategy is in excess of 4 percent.

All of our probit betting strategies generate high success rates in-sample. The strategy of betting on the predicted team, for example, yields a success rate of 54.56 percent in-sample that is significantly greater than the 52.4 percent breakeven rate. The strategy of placing bets only when the probability of success exceeds 57.5 percent performs extremely well in-sample yielding a success rate of 56.42 percent and a net return of 7.71 percent. For this strategy, when the probit model suggests that the probability that the home team will win is greater than 57.5 percent we bet on the home team, and when the probability that the home team will win is less than 42.5 percent we bet on the away team. Both of these findings are supported by the out-of-sample tests. Betting on the predicted team yields an out-of-sample success rate of 56.01 percent and a statistically significant net return of 6.93 percent. Betting on teams with a probit probability of success exceeding 57.5 percent yields an out-of-sample success rate of 61.11 percent and a net return of 16.67 percent, although this is not statistically significant. For the out-of-sample tests of all of our model-based strategies, the model is estimated using data from 1976 to 1992 (2727 games) and tested using data from 1993 and 1994 (341 games). Thus, the parameter estimates and probit probabilities would have been available at the time the bets were placed. From Zuber, Gandar, and Bowers (1985) and Sauer, Brajer, Ferris, and Marr (1988) it is apparent that betting strategies that perform well in-sample need not perform well out-of-sample. In contrast to this earlier literature, however, the strategy of betting on the team predicted by the probit model does generate significant profits both in-sample and out-of-sample.

There are, however, some inconsistencies among the other out-of-sample results. In particular, there is no monotonic relationship between the probit filter and the out-of-sample success rates. For example, the success rate from betting on teams with a probit probability of success exceeding 50 percent is higher than the success rate from betting on teams with a probit probability of success exceeding 55 percent. Moreover, the net returns from betting on teams with probit probabilities of success exceeding 55 percent is negative, although not significant. We conclude from the mixed out-of-sample results that there is an indication that the possible inefficiencies documented in the in-sample tests may be dissipating over time. This is consistent with the results reported in Table VII.

Finally, in order to compare the simple home-underdog strategy with the results of Golec and Tamarkin (1991) and others, and to document that the effort of implementing our probit strategy (rather than a simpler strategy) is warranted, we document the success rates of (1) betting on home teams, (2) betting on underdogs, and (3) betting on home-underdogs on a year-by-year and month-by-month basis. These results are reported in Table VII. The striking feature of Table VII is the dissipation of the well-documented home-underdog bias over recent years. The very strong bias over the 1970s and early 1980s has apparently been eliminated. In only two of the last seven, and three of the last eleven seasons has the home-underdog strategy yielded success rates above 52.4 percent. The monthly comparisons reveal an apparently large

Table VII
Home-Underdog Success Rates

The data consist of 4123 National Football League (NFL) games from 1976 to 1994. The Success Rate is the proportion of times that the particular strategy correctly predicts the outcome.

Period	<i>N</i>	Underdog	Home Team	Home-Underdog
1976	193	51.05	55.96	56.98
1977	195	53.61	50.26	55.26
1978	222	57.41	55.86	68.83
1979	226	51.57	54.87	59.70
1980	225	57.66	47.11	57.58
1981	228	54.22	53.07	61.54
1982	139	51.52	48.20	48.65
1983	225	54.63	46.67	53.45
1984	225	50.45	46.22	45.95
1985	225	45.62	55.56	51.95
1986	235	53.65	48.09	51.85
1987	208	55.39	47.60	54.55
1988	230	52.23	48.70	51.61
1989	222	53.42	48.20	52.24
1990	229	46.61	52.84	49.28
1991	217	51.85	53.00	57.33
1992	222	50.23	50.90	51.28
1993	230	55.11	47.83	55.07
1994	227	52.68	47.14	50.88
September	944	52.01	48.20	52.53
October	1039	53.61	51.88	59.55
November	1061	50.33	49.95	51.30
December	946	51.27	49.89	53.13
January	133	44.36	62.41	76.19

home-underdog bias in the month of January. However since most January games are playoff games, and because favorites are more likely to play playoff games at home, our sample has only 21 games in this category.

III. Conclusions

This article examines the in-sample and out-of-sample performance of various NFL betting strategies. In-sample, probit-based strategies perform extremely well and a number of relatively simple strategies are able to generate positive returns, in excess of commissions. In particular, the strategy of betting on home-team underdogs averages returns of over 4 percent, in excess of commissions, which is statistically significant, although this anomaly has been substantially attenuated over recent years. In particular, in only three of the last eleven seasons has the home-underdog strategy generated returns in excess of the breakeven success rate of 52.4 percent. Moreover, betting on teams predicted by the probit model and teams with a probit probability of success exceeding 57.5 percent generate out-of-sample returns of 6.93 percent

and 16.67 percent respectively. In contrast, other probit-based strategies fail to replicate significant positive in-sample returns out-of-sample. Also, while the strategy of betting on the predicted team generates significant positive returns out-of-sample, the relationship between the probit filter and actual success rates is not monotonic, and some probit-based strategies generate negative net returns (although not significant). These findings are also consistent with the notion that the apparent inefficiencies documented in-sample may be dissipating over time.

Our probit model indicates that the market overreacts to a team's recent performance and discounts the overall performance of the team over the season to date. This suggests that a profitable strategy involves betting on teams that have performed well over the season as a whole, but which have performed poorly in recent weeks. This result has a close analogy with financial markets, amounting to a contrarian strategy.

One avenue for future research involves expanding the set of conditioning information used by the probit model. Exogenous variables such as weather conditions as well as fundamentals such as rushing/passing yards and goal-kicking success rates could be added to increase the predictive power of the model. We are also currently investigating the way in which markets react to evidence of apparent inefficiencies, attenuating apparent excess returns over time.

REFERENCES

- Brown, William O., and Raymond D. Sauer, 1993, Fundamentals or noise? Evidence from the professional basketball betting market, *Journal of Finance* 48, 1193–1209.
- Camerer, Colin F., 1989, Does the basketball market believe in the "Hot Hand"?, *American Economic Review* 79, 1257–1261.
- Gandar, John, Richard Zuber, Thomas O'Brien, and Ben Russo, 1988, Testing rationality in the point spread betting market, *Journal of Finance* 43, 995–1008.
- Golec, Joseph, and Maury Tamarkin, 1991, The degree of inefficiency in the football betting markets, *Journal of Financial Economics* 30, 311–323.
- Thaler, Richard H., and William T. Ziemba, 1988, Parimutuel betting markets: Racetracks and lotteries, *Journal of Economic Perspectives* 2, 161–174.
- Sauer, Raymond D., Vic Brajer, Stephen P. Ferris, and M. Wayne Marr, 1988, Hold your bets: Another look at the efficiency of the gambling market for NFL games, *Journal of Political Economy* 96, 206–213.
- Vergin, Roger C., and Michael Scriabin, 1978, Winning strategies for wagering on National Football League games, *Management Science* 24, 809–818.
- Woodland, Linda M., and Bill M. Woodland, 1994, Market efficiency and the favorite-longshot bias: The baseball betting market, *Journal of Finance* 49, 269–279.
- Zuber, Richard A., John M. Gandar, and Benny 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.