Applied Financial Economics, 2013
Vol. 23, No. 17, 1407–1418, http://dx.doi.org/10.1080/09603107.2013.829201

Inefficient pricing from holdover bias in NFL point spread markets

Andy Fodor[a], Michael DiFilippo[b], Kevin Krieger[c,*] and Justin Davis[d]

[a]*Finance, Ohio University, Athens, OH, USA*
[b]*PNC Financial Services Group, Cleveland, OH, USA*
[c]*Accounting and Finance, University of West Florida, Pensacola, FL, USA*
[d]*Management and MIS, University of West Florida, Pensacola, FL, USA*

We identify inefficiency in the National Football League (NFL) gambling market indicative of sticky preferences by bettors. NFL teams that qualified for the playoffs in the prior season are favoured by too large a margin in the opening week of the following season. Bettors view these teams as superior though they win only 51.7% of opening week games against teams that failed to make the playoffs in the prior year. Against the point spread, teams that made the playoffs in the prior year win only 35.6% of opening week games played against teams that failed to make the playoffs in the prior year. Systematic betting based on this trend results in significant profitability over the 2004–2012 seasons with an average return over 22% per game. We posit this can be explained by gamblers' tendencies to cling to perceptions of teams formed from observation in the prior season. This confirms research in more traditional markets, suggesting investors can be slow to update asset valuations.

Keywords: sports wagering; efficient markets; holdover bias; irrational investment

JEL Classification: L83; G14; G19

I. Introduction

Basic efficient market theory suggests that gambling spreads serve as unbiased prices for wagers on athletic contests. The best available information regarding a game should be reflected in a spread so that each side of a handicapped contest is equally likely to prevail.[1] Bookmakers are thus able to approximately equalize the funds wagered on either side of a contest and thereby guarantee a riskless profit via commissions. This is the traditional model of bookmaking advanced in the literature (see, e.g. Zuber *et al*., 1985; Sauer *et al*., 1988).

The efficiency of the National Football League (NFL) betting market has been a topic of investigation in the finance literature for over 40 years. Given that this market has grown into a multibillion dollar arena for gambling, the primary stream of research has been aimed at identifying inefficiencies that are exploitable for financial gain. To

*Corresponding author. E-mail: kevinkrieger@uwf.edu

[1] Bookmakers set point spreads, or 'lines', in the most common form of handicapping games. The point spread issued by the bookmaker (often a casino or Internet company) establishes the 'favourite' and the 'underdog' of a game. This point spread serves as a correction based on the perceived likelihood of each team winning a game. The favourite is considered more likely to win a game, and thus the spread is instituted in order to place the two sides of a wager on more equivalent footing. A wager is graded based on subtracting the spread from the favourite's final score and comparing this adjusted figure to the score of the underdog. Whichever side then has the higher score is the winning team of the 'against the spread' wager. The team that wins an against-the-spread wager is said to have 'covered' the game or the spread.

this point, findings from these efforts have been mixed, at best. Numerous authors have identified trends in the betting markets that, if exploited over certain periods, would have led to significant profits. However, other authors have often found that such results fail to persist out of sample. Some of the findings of systematic profitability generate from strategies lacking a strong underlying theory, thus exacerbating data-mining concerns.

Pankoff (1968) provided the initial study of efficiency in the NFL betting market, finding that the market was, overall, efficient. However, the first study considering abnormal profitability based on specific strategic wagering in the NFL was undertaken by Sturgeon (1974), who found that betting against the previous week's biggest winner was profitable. In subsequent years, numerous studies have been conducted which consider various betting strategies in search of similar profitability.

Many of these studies report success in this endeavour. Vergin and Scriabin (1978) consider a number of betting rules of thumb and report that betting on large underdogs (those of five or more points) generated approximately a 5% profit over a 5-year period. Zuber *et al*. (1985) find NFL games (from the 1983 season) to be profitably predictable based on a number of team measures such as rushing yards, passing yards, fumbles and interceptions, though they do not find statistical significance, in part due to the small sample size available. Gandar *et al*. (1988) reconsider the betting rules of Vergin and Scriabin (1978) and introduce new rules of thumb to study. Similar to Zuber *et al*. (1985), they document profitability based on trading rules, but no statistical significance. Golec and Tamarkin (1991) describe biases in NFL lines against underdogs potentially leading to a profitable betting opportunity, though the economic impact is not great. More recently, Paul and Weinbach (2012) find betting against big favourites to be profitable and statistically significant. This is confirmed independently by Wever and Aadland (2012), who note that wagering on large underdogs from the 1985 through 2010 NFL seasons would have proven to be a profitable approach.[2]

On the other hand, the documentation of profitability from such 'inefficiencies' has been called into serious question. Sauer *et al*. (1988) describe how wagering based on the findings of Zuber *et al*. (1985) leads to substantial losses in out-of-sample testing. Dare and Holland (2004) find that previously documented inefficiencies of profitable betting on home underdogs are not consistent from season to season. Gray and Gray (1997) find in-sample profitability based on trading rules but acknowledge, in line with the findings of Sauer *et al*. (1988), that results are considerably mixed out of sample. Vergin (1998) finds that profitable trading rules developed by Lacey (1990), based on the 1984–1986 NFL seasons, did not hold for the subsequent 1987–1995 period.

The difficulty in confirming the persistence, out of sample, of the profitability of betting approaches should not be unexpected. As Burkey (2005) notes, authors who search for profitable trading rules which, ex post, prove to have been successful in earlier periods, will surely succeed if they investigate enough strategies. In extreme cases, dozens of betting rules may be investigated in one study. For example, Woodland and Woodland (2000) reject 7 of 48 null hypotheses based on a 10% significance level. Badarinathi and Kochman (1996) consider 116 null hypotheses and reject 7 based on a 5% significance level. An inference of profitability of a specific strategy, based on such a widespread approach, would very likely be committing a type I error. Findings of profitability are particularly tenuous when little theoretical development is undertaken in developing a hypothesized strategy.

We consider the question of whether bettors in gambling markets display irrationally sticky preferences for wagers in a manner similar to investors who are unwilling to update their asset preferences to reflect new information. Brown and Cliff (2005) document that investors are willing to pay unusually high premia for assets they hold in high sentiment. These high prices result in subsequent abnormally low returns, even in the presence of controls. Haruvy *et al*. (2007) document that individuals' beliefs about prices adapt over time, but they are based in part on past experiences. As a more specific example, numerous authors document the willingness of investors to pay high prices for Internet stocks in the midst of the tech bubble due to sentiment (see, e.g. Cooper *et al*., 2002).

This framework motivates our study of the holdover bias of NFL bettors from one season to the next. Specifically, we consider the case of the early weeks of American professional football seasons. Teams that qualified for the NFL playoffs in the prior season are systematically favoured by too many points in Week 1 of the next season when they play opponents who did not qualify for last season's playoffs (or, in the rare cases where they are underdogs, they are underdogs by too few points). Simply wagering, against the spread, on every Week 1 opponent of last season's playoff teams would result in an average return of 22.6% per game over the 2004–2012 NFL seasons. Unlike many other studies of profitability from

[2] The dynamics of European betting markets can sometimes vary from American betting markets as legal sports betting is considerably more prevalent in Europe. 'Hometown bias', for example, where bettors in an area are offered relatively unfavourable prices for wagers on local teams seems to persist in some European cases, as with local, illegal bookmakers in America (Braun and Kvasnicka, 2013). However, in other cases, as many legal European bookmakers exist, such books sometimes compete for business of wagers by offering more attractive odds to more popular European soccer teams (Forrest and Simmons, 2008).

betting strategies, our results are statistically significant as well.

We posit that the more detailed explanation of these findings lies in the inability of the participants of this market, the bettors, to adjust their preferences. This theory is in line with the findings of Brown and Cliff (2005) and Haruvy *et al*. (2007) for more traditional markets. Many gamblers have a perception of a club based on the previous season. When a playoff participant faces a Week 1 opponent that did not qualify for last season's playoffs, a substantial number of bettors are apt to frame the contest as a match-up between a 'successful' and an 'unsuccessful' team and wish to wager accordingly. Bookmakers, cognizant of this irrational preference by gamblers, may adjust the lines they set on a game, forgoing the traditional riskless profit model in search of an expected higher profit.[3] If spreads of such games are, indeed, set in order to entice bettors to cater to their preferences at an effectively higher price, wagering *against* last season's playoff participants in early weeks of NFL seasons should prove profitable.

In related work, Sapra (2008) documents that teams that perform well (poorly) in one season against the spread tend to revert in performance against the spread the following year. This is described as 'overreaction' by bettors to previous impressions. Our findings are perhaps most similar in vein to those of Vergin (2001), who documents that NFL gamblers develop preferences for teams that perform well in the previous game, 2–5 games, or season. He also characterizes this effect as 'overreaction' by bettors from the 1981–1995 seasons. Our results provide out-of-sample verification of such bettor behaviour for the 2004–2012 period. However, our results are also different in some important ways. Unlike Sapra (2008) or Vergin (2001), we focus on the specific case of holdover preferences from one season to the next and study how these preferences are specifically manifested in a season's opening week.[4] In doing so, we document strong profitability levels and, in addition, our results are statistically significant, a threshold few other studies of gambling profitability meet. We focus on only one hypothesis in our formulation of this article, unlike studies that consider multiple betting rules of thumb, and we test our results out of sample for robustness.

Our Week 1 results may emerge because longer periods of time work to more forcefully instil investor (or bettor) preferences. NFL gamblers typically adjust their impressions of teams on a weekly basis throughout a season. The final impression left at the end of one regular season, however, has over 8 months to linger in the minds of bettors before gambling on Week 1 of the following season commences. The psychology literature notes that 'attitude strength', analogous here to an investor preference, is greater when an opinion has been held for a longer time (see, e.g. Holland *et al*., 2003; Glasman and Albarracin, 2006). Savvy bookmakers may exploit these preferences in search of higher expected profits, and thus bettors willing to take 'against the herd' positions may profit. Bettors update their evaluations of teams in the opening weeks of a season, and thus the opportunity to profit based on sticky beliefs regarding last season dissipates after Week 1.

The 22.6% return demonstrated by our study is unprecedented in the NFL wagering literature, even compared to strategies developed in papers which consider dozens of potential betting rules. Vergin and Scriabin (1978) report that betting on large underdogs (those of five or more points) generated approximately a 5% profit over their 5-year study period. Zuber *et al*. (1985) win 59% of wagers in the latter half of NFL seasons based on models built from team performance and a number of variables measured in the first half of seasons. This translates to a 12.3% return. Gandar *et al*. (1988) analyse four mechanical-based and three behaviour-based rules in an effort to profit in the NFL betting market. The most profitable of their strategies involves wagering on underdogs playing a favourite who easily covered the spread in the previous week and translates to an 11.1% return. Golec and Tamarkin (1991) find that betting home underdogs over the 1973–1987 period returned 6.1%. Gray and Gray (1997) utilize a probit model to choose which teams to wager on in NFL games. Using the highest cutoff probability of winning as a threshold for wagering on a team, they demonstrate 7.7% returns in sample and (surprisingly) 16.1% returns out of sample. Vergin (2001) analyses 70 potential NFL betting strategies, and of these 70 strategies the most profitable returns 18.7%. Using data from 1985 to 2008, Wever and Aadland (2012) find that betting on large home underdogs and even larger road underdogs yields an 11.2% return.

While the magnitude of returns we report for the Week 1, prior playoff approach is newfound, the strategy's effectiveness is actually more impressive when considered in full context. We consider only one hypothesis in our study, and this hypothesis is soundly developed theoretically, in contrast to efforts which briefly introduce numerous potential betting rules. The returns of the approach are statistically significant, a threshold that almost no other

[3] Theoretical development and empirical validation of this modified, profit maximization framework was first undertaken by Levitt (2004) and confirmed by Paul and Weinbach (2007) in out-of-sample findings. We elaborate further on this framework in the discussion section.

[4] Vergin (2001) analyses a trading rule (1 of 11 he considers) which bets against all playoff teams in the following season but does not consider the holdover bias question specifically for the early portion of next season. His result regarding last year's playoff teams is that they are no more or less likely to cover spreads in the following season.

paper considering the NFL gambling market meets. The returns of the approach are not statistically significant when investigated, out of sample, via seasons in the distant past; however, the returns of the strategy are still positive in these earlier seasons, even after factoring in bookmaker commission, a standard few other studies reach out of sample. Furthermore, we detail explanations for why the approach, though recently successful, may not have been as relevant in the long-ago seasons used for out-of-sample analysis. Overall, the effectiveness, development and robustness of the approach are a significant contribution to the literature.

The rest of this study is presented as follows. First, we provide theoretical support for our suggested hypothesis. Next, we describe our data collection methods and the statistical methodology used. We then present the results of the study. For robustness, we describe the performance of our strategy out of sample. In doing so we also discuss the reasons why profitability of the strategy may continue going forward. We conclude in the final section.

II. Hypothesis Development

In the traditional model of bookmaking, bookmakers attempt to set a line for a game so that equal amounts of money will be wagered on each side, and the bookmaker may then claim a riskless profit due to the 10% commission that is charged on winning bets (e.g. Zuber *et al*., 1985; Sauer *et al*., 1988).[5] A refined model has been developed in recent years, claiming that bookmakers are willing to instead take some level of risk in order to increase their expected profitability. Bookmakers may do so by setting lines so that naïve bettors will flock to the side of a contest which is less than 50% likely to prevail in a wager.

Avery and Chevalier (1999) note that a great deal of ‘dumb money’ exists in the betting market and that bettors make errors in their wagering based on preferences for teams, visibility of teams and momentum beliefs. A bookmaker might ignore this fact and set a line that achieves equal betting on each side of a game, or the bookmaker may shift the line in order to intentionally allow more of the betting to occur on the side of a contest less likely to win. Thus, by assuming some risk on individual contests, the bookmaker may increase profits over time. We offer the following example for clarification. A perfectly knowledgeable bookmaker handicaps the actual spread of a game between the New Orleans Saints and Cleveland Browns to be the Saints – 9 points. The perfect bookmaker also knows most bettors would prefer to bet on the Saints (or perhaps, against the Browns), and they are so eager to do so that betting would not be equalized between the two sides unless the line were Saints – 11 points. The bookmaker may set the line at Saints – 10 points and therefore attract more than 50% of the money wagered on the Saints and be more than 50% likely to win these bets. This approach would result in an expected profitability greater than the 5% guaranteed from the riskless, commission-only approach. Given a large enough bankroll, over time, a large number of games handicapped thusly will minimize any risk to the bookmaker.[6]

The profit maximization approach is discussed by Strumpf (2003), who notes that hometown bookies may set spreads which local teams are notably less than 50% likely to cover. Levitt (2004) conducted a seminal study confirming this approach by bookmakers in the NFL, based on special data from an NFL handicapping contest at the Las Vegas Hilton, and this finding was later supported by Paul and Weinbach (2007).

One might ask if informed bettors (often professional gamblers known as ‘sharps’ or ‘wiseguys’) might not take advantage of such lines and nullify the bookmaker’s profitability by wagering on the ‘correct’ side of a contest. While sharps undoubtedly seek out such opportunities, their impact is limited by three factors. First, a sharp cannot be certain that his handicapping of any one game is superior to the bookmaker, and thus risk aversion dictates prudence. Second, the bankroll of a sharp is typically considerably less than that of the bookmaker, and here again risk aversion prevents sharps from too heavily exploiting even a spread they confidently identify as biased. Third, the impact a sharp can make in a betting market like the NFL is saturated by ordinary, naïve bettors. Betting limits placed on gamblers by Internet companies or casinos further limit this impact.[7]

Given the documentation of this framework, we posit that bookmakers intentionally set spreads in such a manner as to attract naïve bettors to wager on teams for which they are biased. We further describe next why this bias is likely to be based on sentiment for teams perceived as ‘successful’ as demonstrated by their post-season presence in the previous NFL year. Bettors willing to take the contrarian approach may thus be able (like the sharps) to place wagers with positive expected value, even factoring in bookmaker commissions.

[5] As another example, Lee and Smith (2002) note, ‘Bookies do not want their profits to depend on the outcome of the game. Their objective is to set the point spread to equalize the number of dollars wagered on each team..’.

[6] While this structure supposes a ‘perfect’ bookmaker, anecdotal documentation exists that real bookmakers are, indeed, willing to take such an approach to improve profits (see Millman, 2001).

[7] Typical maximum bets range between $10,000 and $25,000 for one side of an NFL contest. Any wagers larger than the maximum amount must be broken into separate wagers, allowing the bookmaker to change the spread in the interim.

Barberis *et al*. (1998) build a model describing how investor sentiment may impact asset returns. They cite the conservatism bias noted by Edwards (1968) in the psychology literature that holds that people are slow to update beliefs based on information. When beliefs *are* altered, the updates typically undershoot the true value that should be reflected. In our study, we hypothesize a direct parallel for participants in the betting market for NFL games. In particular, we believe that bettors are cognizant that the end of one NFL season marks an important shift in the quality of a team. Substantial personnel turnover of both players and coaches occurs in between seasons, and many teams, particularly unsuccessful ones from the previous year, are likely to make dramatic strategic shifts in the 8-month offseason. Nevertheless, many investors may hold too tightly to their perceptions of success from the previous season and be willing to wager accordingly. Thus, the prices offered by bookmakers, in the form of lines, may be set at specific levels in order to take advantage of these biases.

Broad empirical evidence supports Barberis *et al*.'s (1998) theory that investor sentiment inflates asset prices and therefore yields lower future returns. Neal and Wheatley (1998) consider three different measures of investor sentiment: odd-lot sales frequencies, closed-end fund discounts and net mutual fund redemptions. They note that the first two proxies help predict stock returns. Fisher and Statman (2000), rather than studying different proxies of sentiment, consider three different groups of stock market participants and conclude that the preferences of all three are negatively linked to future returns. Wall Street strategists and individual investor sentiment are each significantly negatively linked to the performance of stocks in the future while the sentiment of newsletter writers is insignificant but still negative in direction.

Brown and Cliff (2005) note that it is difficult to demonstrate inefficiency due to sentiment in traditional markets. They conduct a survey of investors and note new evidence of inefficiency due to sentiment. Excessive optimism of investors is linked to future periods of market-wide overvaluation. The result is economically significant as well. In a laboratory setting, Haruvy *et al*. (2007) show that investors update prices over time, but prices are based on past trends that traders have experienced in markets. This is analogous to our theory that bettors hold on to perceptions from the prior NFL season. While prices converge to fundamentals in Haruvy *et al*.'s experiment, bubbles do persist for a time.

The evidence supporting the negative link between investor sentiment and asset returns is not restricted to domestic stocks. Nor is the impact of sentiment demonstrated only by identifying which stocks are held in highest esteem and separating those stocks for comparison. Schmeling (2009) empirically confirms the sentiment theory in the international marketplace by studying 18 industrialized countries and noting that, across nations, investor sentiment is negatively linked to future stock returns. The findings are particularly strong in those countries more prone to overreaction. The empirical findings of the impact of sentiment are not restricted to stock markets either. Han (2008) demonstrates that beyond factors typically modelled, index option prices reflect the sentiment surrounding the equity market. We suggest that this investor sentiment argument is also applicable in the sports betting framework.

III. Data and Methodology

Our data cover 2304 games played in the 2004–2012 NFL regular seasons. All betting lines and game results were collected from Sunshine Forecast[8] which provides free downloadable statistics and spread information regarding the NFL.

We hypothesize that inefficiencies may exist in gambling markets in the early weeks of an NFL season due to bettors overemphasizing team performance in the prior year. We choose to explore the implications of prior season playoff appearances (and thus utilize the terminology 'prior playoff teams') because bettors might classify a team as having a successful year when it qualifies for the postseason. After one NFL season ends, this perception of success (or failure, in the case of nonplayoff teams) could remain the impression of a club, for bettors, until a new NFL season begins. Our study considers whether such impressions may push a market to inefficiency. To search for such a 'holdover bias', we examine whether lines are systematically errant when prior playoff teams face prior nonplayoff teams in the following season.

For each week, we calculate the percentage of prior playoff teams which are favoured when playing prior nonplayoff teams.[9] We compare this to the weekly percentage of prior playoff teams that actually won these games. We test the significance of these differences with a difference of proportions *z*-test.

We also calculate the percentage of prior nonplayoff teams which cover spreads, by week, when playing prior playoff teams. We then test if this percentage is significantly different than 52.38% via a two-sided test. This is accomplished with a one-sample proportions *z*-test. It is common convention in sports betting literature to compare win percentages to 52.38% rather than 50.0% because, due to bookmaker commission, for a bettor to break even, based on a number of equally sized bets, he or she must win more than 52.38% of these wagers.

[8] www.repole.com/sun4cast
[9] 974 of these games match prior playoff teams with prior nonplayoff teams.

We also calculate line errors for all games where prior playoff teams play prior nonplayoff teams. Line errors are calculated as

$$\begin{aligned}\text{Line Error} = {} & \text{Prior Playoff Team Score} \\ & + \text{Prior Playoff Team Point Spread} \\ & - \text{Prior Nonplayoff Team Score} \end{aligned} \tag{1}$$

Prior Playoff Team Score and Prior Nonplayoff Team Score are the realized scores from games pitting prior playoff teams against prior nonplayoff teams. Prior Playoff Team Point Spread is the point spread relative to the team that made the playoffs in the prior year. The calculated Line Error reflects the number of points by which the spread of the game is incorrect, relative to the prior playoff team. A negative (positive) Line Error means the playoff team was favoured by too many (few) points and hence failed to cover the spread (covered the spread). For example, if previous playoff qualifier New England was favoured by 10 points over nonqualifier Detroit, and the final outcome of a game between the teams was New England 24, and Detroit 17, then the Line Error would be –3, reflecting that New England, the prior playoff team, was favoured by three points too many, relative to perfect foresight.

We next perform a regression analysis designed to test the efficiency of NFL lines. We estimate the following specification, as presented in Zuber *et al*. (1985):

$$\text{Score Difference}_i = \alpha + \beta(\text{Point Spread}_i) + \varepsilon_i \tag{2}$$

where Score Difference_i is the score of the prior playoff team less the score of the prior nonplayoff team, Point Spread_i is the closing line for the game relative to the prior playoff team, and ε_i is the error term. One test of betting line efficiency is a test of the joint hypothesis that $\alpha = 0$ and $\beta = 1$. We thus estimate the coefficients from regressions for each NFL week (aggregated over all years in the sample period). We then conduct an F-test, based on restricting α and β to be 0 and 1, respectively. As detailed by Zuber *et al*. (1985), statistical rejection of line efficiency, based on this specification, is difficult due to the relative lack of power of such a test.

Finally, we present specific, season-by-season performance of a strategy which wagers against prior playoff teams in the early weeks of a new NFL season, in order to benefit from the holdover bias of ordinary bettors. This is the most important metric in determining the benefit of a potential wagering strategy. Summary statistics of line errors and regression approaches provide insight into the bookmaking dynamic, but such results can be influenced greatly by large discrepancies between lines and final game scores in a few, select games, and the most important question is how often any strategy succeeds versus how often it fails, not how much it succeeds or fails by, on average.

We calculate the overall, in-sample performance as well and compare the winning frequency of these wagers to the breakeven level of 52.38%. While a record better than 52.38% indicates economic profitability of the strategy, a mark heavily emphasized in the literature, we further test whether the strategy is statistically significantly better than 52.38% via a one-sample proportions z-test. We utilize a one-sided significance threshold as we hypothesize that our strategy will exceed the profitability breakeven mark. Statistical significance of economically profitable returns would be an almost unprecedented finding for a betting market strategy, given the rather small samples that permeate the literature. Occasional studies report strategies which are successful more than 52.38% of the time, but very rarely are these strategies successful at a rate that is statistically significantly greater than 52.38%.

IV. Results

Table 1 displays the performance of prior playoff teams, week-by-week, in the subsequent season. We consider how frequently prior playoff teams are favoured, how often they win, and how often they cover the spread when playing prior nonplayoff teams.

If the betting lines create an efficient gambling market, we would expect to see no difference between the percentage of teams that are favoured and the percentage of teams that win. From 2004 to 2012, prior playoff teams, in Week 1, are favoured in 85.0% of games. However, they only win 51.7% of the time. This striking disparity means that, even with a relatively small sample size, the Week 1 difference of proportions test between percentage of prior playoff teams favoured and percentage of playoff teams who win is significant at the 1% level. The analogous difference of proportions test in Week 2 is significant at the 5% level.

Prior playoff teams remain favoured in more games than they actually win throughout almost all the remaining weeks of an NFL season. Table 1 notes that only in Week 16 do prior playoff teams actually win as many games as they were predicted to via the gambling line, though even this result is exact (35 prior playoff teams win games while 35 were expected to, in Week 16, over the 2004–2012 seasons). The full sample of 1084 games with one prior playoff team shows 821 teams that were favoured, but only 677 of these teams actually won.[10] The overall

[10] Exactly 50% of prior playoff teams advance to the playoffs in the following year over our sample period. 37.5% of teams advance to the playoffs in each season, so holdover biases, as compared to naïve assumptions that evaluate all teams equally, appear to have a basis in reality; however, the game-by-game projections of prior playoff team performance are too high, particularly in Week 1.

Table 1. Playoff team favourite and win percentages

Week	*N*	Favoured	Percentage favoured	Win	Percentage win	Favoured percentage – Win percentage	Win against spread	Lose against spread	Percentage win against spread
1	60	51	85.00	31	51.67	33.33***	21	38	35.59*
2	72	57	79.17	45	62.50	16.67**	35	35	50.00
3	71	57	80.28	48	67.61	12.68	34	35	49.28
4	59	45	76.27	39	66.10	10.17	31	28	52.54
5	66	52	78.79	46	69.70	9.09	38	28	57.58
6	59	45	76.27	36	61.02	15.25*	26	31	45.61
7	62	48	77.42	36	58.06	19.35**	30	32	48.39
8	58	44	75.86	39	67.24	8.62	34	23	59.65
9	56	42	75.00	33	58.93	16.07*	29	26	52.73
10	62	49	79.03	42	67.74	11.29	32	29	52.46
11	65	52	80.00	47	72.31	7.69	34	30	53.13
12	63	53	84.13	40	63.49%	20.63**	32	31	50.79
13	59	43	72.88	34	57.63	15.25*	24	31	43.64
14	75	52	69.33	44	58.67	10.67	35	36	49.30
15	62	45	72.58	38	61.29	11.29	30	31	49.18
16	64	39	60.94	39	60.94	0.00	35	28	55.56
17	71	47	66.20	40	56.34	9.86	27	44	38.03
2–17	1024	770	75.20	646	63.09	12.11***	506	498	50.40
All	1084	821	75.74	677	62.45	13.28***	527	536	49.58

Notes: The statistics for each NFL week over the 2004–2012 seasons is presented. Only games where a prior playoff team played a prior nonplayoff team are included. For each week, the number and percentage of previous playoff teams favoured to win the game and the number and percentage of prior playoffs teams who won games are presented. Differences between these proportions are shown with significance levels from a difference of proportions *z*-tests indicated. The results of games against the spread are also presented for each week. We exclude games where the outcome against the spread was a tie. The percentage of games where the prior playoff team won against the spread is presented with significance levels from one-sample *z*-tests for a winning percentage of 52.38% indicated. For conservatism, all tests are two-sided.
***, ** and * indicate significance at the 1%, 5% and 10% levels, respectively.

difference of proportions test is significant at the 1% level. The result is most powerful, however, in Week 1.

Given the organization of gambling markets, the more relevant question is the performance of prior playoff teams after the spread is incorporated in games the subsequent season. In Table 1, we note that in only 21 of 60 (35.6%) Week 1 games does a prior playoff team win against the spread. In 38 of 60 games, the prior playoff team loses against the spread (there is also one tie or 'push' in the sample). This proportion of 64.4% of prior playoff teams failing to cover the spread (excluding the one tied game) is significantly different than 52.38% at the 10% level. This extreme, almost unprecedented result means that systematically wagering against prior playoff teams in Week 1 of an NFL season proves to be very profitable in our sample period (further details are provided in Table 4). While we earlier observed that prior playoff teams underperform throughout the season in terms of wins versus bookmaker projections, the results against the spread are considerably different. In no week, other than Week 1, is the proportion of prior playoff teams covering significantly different than 52.38%. For the combined sample of Weeks 2–17, 50.4% of prior playoff teams cover their games against prior nonplayoff teams.

We next calculate statistics regarding line errors. Table 2 presents the mean and median line errors for games in which exactly one of the participants is a prior playoff team. Line errors are presented which compare the actual difference in the scores of the teams to the bookmaker spreads of the games. The errors in the first column are in reference to the prior playoff team.[11]

Table 2 shows the average line error for Week 1 games in which exactly one participant is a prior playoff team is –3.74 points. This means that, on average, the prior playoff team was 3.74 points short of covering the spread (or in other words, that the spread is 3.74 points higher, on average, than it should be, biased towards prior playoff teams, in Week 1). This result is significantly different than 0 at the 5% level. The Week 1 line error is significantly different than the average line error (0.88 points) of prior playoff teams in Weeks 2–17 at the 1% level. The results are similar when we consider the median errors. The median error is 5.5 points biased towards the prior playoff team, and this result is statistically different than 0

[11] For example, if team A is not a prior playoff team and team B is a prior playoff team, and team B is favoured by three points, and team A wins the game by two points, then the line error is –2–3 = –5 points.

Table 2. Line errors

Panel A: Mean line errors

	One playoff team		No playoff teams		Two playoff teams	
	N	Mean	*N*	Mean	*N*	Mean
$\text{Error}_{\text{Week 1}}$	60	−3.74**	60	0.62	24	3.43
$\text{Error}_{\text{Weeks 2–17}}$	1024	0.88*	837	0.21	299	−0.92
$\text{Error}_{\text{Week 1}} - \text{Error}_{\text{Weeks 2–17}}$		−4.62***		0.42		4.35

Panel B: Median line errors

	One playoff team		No playoff teams		Two playoff teams	
	N	Median	*N*	Median	*N*	Median
$\text{Error}_{\text{Week 1}}$	60	−5.50*	60	−0.50	24	4.00
$\text{Error}_{\text{Weeks 2–17}}$	1024	0.00	837	0.00	299	−0.50
$\text{Error}_{\text{Week 1}} - \text{Error}_{\text{Weeks 2–17}}$		−5.50***		−0.50		4.50

Notes: The mean and median line errors after separating games according to the number of teams in the game that made the playoffs in the prior year are presented. For the one-playoff team sample, line errors are calculated relative to the prior playoff team as prior playoff team score, less prior nonplayoff team score, adjusted for the spread of the game relative to the prior playoff team. A negative line error therefore indicates that prior playoff teams were favoured by too many points (or were underdogs by too few points) and fell short of covering the spread by this amount. A positive line error indicates that prior playoff teams were favoured by too few points (or were underdogs by too many points) and covered the spread by this amount. Line errors for the no-playoff and two-playoff team samples are calculated relative to the favourite team as the favourite team score, less the underdog team score, less the number of points the favourite team is favoured by. A negative line error therefore indicates the favourite team was favoured by too many points and fell short of covering the spread by this amount. A positive line error indicated that favourite team was favoured by too few points and covered the spread by this amount. Significance levels from *t*-tests for means and signed-rank tests for medians are indicated. Results are presented separately for Week 1 and Weeks 2–17. Differences between Week 1 and Weeks 2–17 line errors are presented with significance levels from two sample *t*-tests for mean differences and Wilcoxon rank-sum tests for median differences. The sample period is the 2004–2012 NFL seasons.
***, ** and * indicate significance at the 1%, 5%, and 10% levels, respectively.

at the 1% significance level (which is the median error for games involving one prior playoff team in Weeks 2–17).

The abnormally high errors in Week 1 may be attributed to our theory that the betting market places too great an emphasis on prior season performance, which leads to inaccuracies in the spreads. Furthermore, the sharp reduction in line errors between Weeks 1 and Weeks 2–17 matches our expectation that this holdover bias dissipates once new information is available with which to update bettor preferences. As a brief check that the aforementioned bias in Week 1 for prior playoff teams is based in part on the Week 1 status of the game, and not entirely on the prior playoff status, we compare the line errors across other types of games in our data set. Table 2 shows that the average line error is higher in Week 1, as compared to Weeks 2–17 for all games, and when splitting these games into three categories: (1) games with no prior playoff teams, (2) games with one prior playoff team and (3) games with two prior playoff teams. Higher errors for these types of games (all relative to the favourite of the game) confirm that the betting market has particular difficulty in handicapping Week 1 games.

Interestingly, the small sample of 24 observations suggests that in Week 1 games with both participants as prior playoff teams, the favourite of the game is favoured by too few points. We briefly posit an explanation which coincides with our other findings. Specifically, it appears that favourites are over-favoured in Week 1 games with two prior playoff teams. It may be that the ordinary bettors of the gambling market lump all prior playoff teams too closely together in their impressions. Prior playoff status may serve as a type of dummy variable for quality, in the minds of naïve bettors, which then receives too much weight. In games with one prior playoff team, this results in overestimation of that team's quality. In games with two prior playoff teams, this results in overestimation of an underdog's ability, simply because it made the postseason in the prior year.[12] Prior playoff status may distort the market's perception of teams, which leads to high variability in the line errors in Week 1 of the next season.

As a more formal test of gambling market statistical efficiency, we employ the regression method, first noted by Zuber *et al.* (1985), which regresses the final point differential of games on the spread. We utilize our sample

[12] In most games matching two prior playoff teams, the team that advanced further in the previous year's playoffs is the favourite. Thus, the fact that a team made the playoffs in the prior season seems to overshadow the actual gap in quality between the two teams.

Table 3. Regression analysis

Week	N	Intercept (α)	Line (β)	F-stat
1	60	−4.35*	1.133	4.82***
2	72	−2.52	1.631	3.53**
3	71	4.43**	0.576	4.44**
4	59	2.71	0.825	1.22
5	66	2.56	0.970	1.82
6	59	1.48	0.704	0.54
7	62	−0.07	0.945	0.04
8	58	3.63	0.688	2.54*
9	56	−1.07	1.055	0.43
10	62	2.13	0.714	2.31
11	65	2.75	1.239	5.70***
12	63	−4.06*	1.813	5.12***
13	59	−0.95	1.148	0.49
14	75	−0.56	1.295	3.04**
15	62	−0.56	1.182	0.56
16	64	2.88	0.844	1.89
17	71	−2.29	1.303	1.85

Notes: The regression results for each NFL week over the 2004–2012 seasons are presented. The estimated model regresses the realized point differential of games involving one prior playoff team, relative to the prior playoff team, on the spread of the game, relative to the prior playoff team. Only games where a prior playoff team played a prior nonplayoff team are included. Intercepts and coefficients of spreads are presented. The F-stats shown test the joint hypothesis that $\alpha = 0$ and $\beta = 1$.
***, ** and * indicate significance at the 1%, 5% and 10% levels, respectively.

of games involving exactly one prior playoff team and run the regression on a week-by-week basis from our sample of the 2004–2012 NFL seasons. The results are presented in Table 3.

The null hypothesis of an efficient gambling market can be rejected when an F-test of the joint hypothesis that $\alpha = 0$ and $\beta = 1$ provides a significant test statistic. We are able to reject the Week 1 lines as inefficient at the 1% significance level while the Week 2 and 3 lines are rejected as inefficient at the 5% level.

The more important metric of success, given the motivation of gamblers, is economic profitability. A betting strategy is profitable, after factoring in bookmaker commission, if it results in winning wagers more than 52.38% of the time. We test the profitability of wagering on the opponents of prior playoff teams in Week 1 and Week 2 of an NFL season and present our results in Table 4.[13]

The results indicate substantial profitability from betting against prior playoff teams in Week 1 of the following season. In all seasons besides 2012, the approach is profitable and yields an average profit of 22.6% per game. Furthermore, the statistical significance of the profitability level is strong. A one-sided test, based on the direction advocated by our strategy, rejects the null hypothesis that the real winning percentage of this strategy is less than or equal to 52.38% at the 5% significance level. Explicitly considering the statistical significance of the economic profitability level is rare, even amongst those papers noting economic significance (e.g. Zuber *et al*., 1985), and this significance is often nonexistent or weak when it is considered (e.g. Gray and Gray, 1997). Finding statistical economic significance for our only hypothesis is a strong indication of an inefficient gambling market, tied to a holdover bias.

We also note that statistical significance regarding economic profitability is not redundant to the statistical significance found from the regression approach in Table 3. Specifically, the regression approach considers whether the margins of actual game scores compare well with the margins predicted via spreads. This is an interesting statistical question, but it is a different matter to consider the question of profitability, in which each wager is an all-or-nothing proposition. In the regression framework, a large error in one game might affect parameter estimation greatly and alter statistical significance. In the economic approach, all games are of equal importance to our conclusions, regardless of the level of error between spread and actual result. Determining statistical significance via both of these approaches adds to our confidence regarding the findings of holdover bias from one NFL season to the next.

The difference between economic significance and statistical significance from the regression framework is highlighted by our results regarding Week 2 in Table 4. While we reject the null hypothesis of market efficiency based on the regression approach, there is no economic profitability, ex post, from betting against prior playoff teams in Week 2 of an NFL season. Our findings of holdover bias appear to be only relevant in Week 1.

V. Out-of-Sample Robustness and Future Prospects

Given the unprecedented profitability over our 2004–2012 study period for the Week 1 strategy, wagering against prior playoff teams who play prior nonplayoff teams, it is natural to question whether the findings are an in-sample phenomenon or, more cynically, the result of data snooping. We utilize historical point spread data from another source, footballlocks.com, and calculate the effectiveness of the Week 1 strategy over the earlier 1994–2003 period.

A new culture of competitive balance was developing throughout the NFL in the early 1990s. With salaries rising rapidly, players began to more frequently leave their old teams for substantial pay raises, even if this meant departing some of the most successful teams and agreeing to play

[13] Results of profitability for other weeks are all insignificant, as one would expect based on the against-the-spread results of Table 1. We reconsider the cases of Weeks 1 and 2 separately in Table 4 to provide further discussion.

Table 4. Profitability of betting against prior playoff teams

Season	N	Lose against spread	Push	Win against spread	Percentage lose against spread	Percentage return
Panel A: Week 1						
2004	6	4	0	2	66.7	27.3
2005	10	7	0	3	70.0	33.6
2006	10	6	0	4	60.0	14.5
2007	4	3	0	1	75.0	43.2
2008	8	5	0	3	62.5	19.3
2009	6	3	1	2	60.0	12.1
2010	4	3	0	1	75.0	43.2
2011	6	4	0	2	66.7	27.3
2012	6	3	0	3	50.0	−4.5
Total	60	38	1	21	64.4**	22.6
Panel B: Week 2						
2004	8	3	0	5	37.5	−28.4
2005	8	4	0	4	50.0	−4.5
2006	10	4	0	6	40.0	−23.6
2007	6	5	0	1	83.3	59.1
2008	12	4	1	7	36.4	−28.0
2009	6	3	0	3	50.0	−4.5
2010	8	5	0	3	62.5	19.3
2011	6	4	0	2	66.7	27.3
2012	8	3	1	4	42.9	−15.9
Total	72	35	2	35	50.0	−4.4

Notes: The results of prior playoff teams when playing against prior nonplayoff teams for our sample period of the 2004–2012 NFL seasons are presented. The sample consists of Week 1 and Week 2 games and is divided according to season. Against-the-spread results are presented for wagering on the prior nonplayoff teams. The percentage of games where the prior playoff teams loses against the spread is calculated excluding pushes. Due to bookmaker commission, on a bet of \$110, a win pays \$210, a push pays \$110 and a loss pays nothing; thus, to be profitable, wagers must win 52.38% of the time. Statistical significance of betting against prior playoff teams in Weeks 1 and 2 is measured by comparing winning percentages to 52.38%. Statistical significance is tested for total Week 1 and Week 2 results only, given the small sample sizes for individual years and the nature of our strategy.
** denotes statistical significance at the 5% level. We utilize one-sided significance levels due to our posited findings. Percentage return is calculated as total return on equal bets on each game over the total invested and incorporates the bookmaker commission.

for some of the least successful.[14] Furthering this shift towards increased league-wide balance, the first salary cap and floor were introduced to the NFL in 1994. These rules placed a cap on the total money used to pay player salaries by each team, as well as a minimum amount that a club could spend on player salaries. Restrictions of total team pay level were designed to place all franchises on a more equivalent footing. The NFL desired this increased parity across the league in order to garner more widespread interest and thereby increase advertising, ticket and merchandising revenues via broad exposure. This intended change manifested shortly thereafter: 12 different franchises have won the NFL championship in the 17 seasons since the salary cap's introduction. In the 17 seasons prior, only eight different franchises won an NFL championship.

We effectively study whether a by-product of this dynamic, in the betting markets, has been the development of overly sticky beliefs regarding team quality, based on playoff participation, from one season to the next. From 1970 to 1993, with no salary cap in place, 61.2% of teams that qualified for the previous year's postseason returned to the playoffs, so such beliefs may have been somewhat warranted. This was the case even though the NFL playoffs consisted of only 8 teams until 1978 and 10 teams from 1979 to 1989.[15] The strength of such beliefs, however, may have been too great once the salary cap era arrived. For our out-of-sample robustness period of 1994–2003, the start of the salary cap era, only 53.3% of teams that qualified for the previous year's playoffs returned the following season (all these in the new,

[14] For a short time, the owners of NFL clubs resisted this development via the implementation of 'Plan B'. In 1989, this rule gave teams the right to 'protect' up to 37 players on their rosters. This protection guaranteed teams the opportunity to match any offer from another team to a current player approaching free agency and thereby retain that player's services. This rule was found to be a violation of antitrust laws in by a US Federal Court in 1992.
[15] League membership was smaller in these seasons, but not by an amount proportional to the smaller playoff fields.

12-team playoff era). For our sample period of 2004–2012, this figure shrinks further to 50%. We posit that it took some time for the salary cap's effects to fully emerge, with teams learning, over time, to more aggressively solicit players from opposing clubs and players more frequently seeking pay maximization.

For our out-of-sample robustness period of 1994–2003, the strategy of betting against prior playoff teams that face nonplayoff teams in Week 1 games results in a winning rate of 54.3%, commensurate with a return of 3.6% after factoring in the typical bookmaker commission. Unlike our in-sample result, this economic finding is not *statistically* significant; however, we note that the result still rises to the threshold of economic profitability out of sample, which is the most important viewpoint to consider.[16] Most all prior studies reporting 'effective' betting strategies (e.g. Zuber *et al*., 1985; Lacey, 1990) do not actually demonstrate statistically significant profitability (win rates statistically better than 52.4%) *even in sample*. It is testament to the strength of our sample result that we find performance which is not just historically profitable but which meets the almost unprecedented standard of statistically significant economic profitability.

Furthermore, we believe that the course of the development of increased parity of competition in the NFL betting market also supports the improvement of the strategy from the robustness seasons to the sample seasons. In more recent seasons, roster turnover has continued to increase, highlighting the fallacy in believing that last year's teams will remain intact and therefore achieve similar results in the following season. There is little reason to believe this trend will reverse. Ordinary bettors have yet to overcome the temptation to overvalue last season's playoff performance even though this trend has been developing for many years. We suspect they may never do so. Betting against prior playoff teams has actually become more profitable in the most recent NFL seasons, our sample period. We anticipate the sticky preferences of naïve bettors will continue to offer profitably opportunities for informed bettors, who are able to take contrarian positions, in the future.[17]

VI. Conclusion

In our study, we consider a unique sentiment question. We hypothesize that bettors in a gambling market are overly influenced by impressions from a prior time, rather than fully adjusting their beliefs and pricing their wagers accordingly. The holdover bias, based on a classification of a team as successful because it qualified for the prior season's playoffs, allows gamblers to pay too high a price in order to cater to their biases. The structure of the betting marketplace, the overabundance of naïve participants and the limits on entry of informed bettors allow for the existence of inefficient prices. It is particularly easy to profit from these 'high price' spreads in a betting market. It is not necessary to arrange for permission to sell short or to open special accounts for margin or derivatives. Contrarian bettors need only wager against prior playoff teams. The statistical evidence confirms that bettors who take the contrarian viewpoint by simply wagering against prior playoff teams in Week 1 of an NFL season may profit from the effectively cheap prices.

Some of the prior literature suggests that the NFL betting market does not contain any inefficiencies that can be profitably exploited by informed bettors. Those studies which have reported such inefficiencies have frequently been criticized as under-motivated, unlikely to persist out of sample, and driven by data snooping. Over the 2004–2012 NFL seasons, we find unprecedented profitability based on taking a contrarian strategy of wagering against prior playoff teams in the following season's opening week. The line errors of these games indicate a statistically inefficient market, utilizing the low-power tests of Zuber *et al.* (1985). Most impressively, the profitability level of the strategy is itself found to be statistically significant, a hurdle cleared by virtually no other study of the sports betting market.

We consider only one hypothesis in our study, we develop its theoretical underpinnings through the financial and psychological literature and we measure its performance out of sample. In out-of-sample testing, we demonstrate that the profitability of our Week 1 strategy holds for other time periods, though at reduced levels relative to the unprecedented 2004–2012 period. Each season, as more information about the current teams becomes available and stale impressions from the previous season can be more thoroughly updated, our analysis shows a learning effect that occurs in the betting markets. As the season progresses, the line errors for games involving exactly one prior playoff team stabilize.

References

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

Badarinathi, R. and Kochman, L. (1996) Football betting and the efficient market hypothesis, *The American Economist*, **40**, 52–5.

[16] Also, like our sample period, our robustness period is quite limited in sample size of Week 1 playoff team versus nonplayoff team games, making statistical significance very difficult to achieve.

[17] We also note that the emergence of the Internet, in recent years, has made sports wagering a more acceptable and accessible activity; thus an increased number of ordinary bettors may now exist, giving more credence to the Week 1 inefficiency theory.

Barberis, N., Shleifer, A. and Vishny, R. (1998) A model of investor sentiment, *Journal of Financial Economics*, **49**, 307–43.
Braun, S. and Kvasnicka, M. (2013) National sentiment and economic behavior: evidence from online betting on European football, *Journal of Sports Economics*, **14**, 45–64.
Brown, G. and Cliff, M. (2005) Investor sentiment and asset valuation, *Journal of Business*, **78**, 405–40.
Burkey, M. (2005) On 'arbitrage' and market efficiency: an examination of NFL wagering, *New York Economic Review*, **36**, 13–28.
Cooper, M., Dirmitrov, O. and Rau, P. R. (2002) A rose.com by any other name, *Journal of Finance*, **56**, 2371–88.
Dare, W. H. and Holland, A. S. (2004) Efficiency in the NFL betting market: modifying and consolidating research methods, *Applied Economics*, **36**, 9–15.
Edwards, W. (1968) Conservatism in human information processing, in Kleinmutz, B. (Ed.) *Formal Representation of Human Judgment*, Wiley, New York.
Fisher, K. and Statman, M. (2000) Investor sentiment and stock returns, *Financial Analysts Journal*, **56**, 16–23.
Forrest, D. and Simmons, R. (2008) Sentiment in the betting market on Spanish football, *Applied Economics*, **40**, 119–26.
Gandar, J., Zuber, R., O'Brien, T. *et al*. (1988) Testing rationality in the point spread betting market, *The Journal of Finance*, **43**, 995–1008.
Glasman, L. and Albarracin, D. (2006) Forming attitudes that predict future behavior: a meta-analysis of the attitude-behavior relation, *Psychological Bulletin*, **132**, 778–822.
Golec, J. and Tamarkin, M. (1991) The degree of inefficiency in the football betting market, *Journal of Financial Economics*, **30**, 311–23.
Gray, P. K. and Gray, S. F. (1997) Testing market efficiency: evidence from NFL sports betting market, *The Journal of Finance*, **52**, 1725–37.
Han, B. (2008) Investor sentiment and option prices, *Review of Financial Studies*, **21**, 387–414.
Haruvy, E., Lahav, Y. and Noussair, C. (2007) Traders' expectations in asset markets: experimental evidence, *The American Economic Review*, **97**, 1901–20.
Holland, R., Verplanken, B. and van Knippenberg, A. (2003) From repetition to conviction: attitude accessibility as a determinant of attitude certainty, *Journal of Experimental Social Psychology*, **39**, 594–601.
Lacey, N. (1990) An estimation of market efficiency in the NFL point spread betting market, *Applied Economics*, **22**, 117–29.
Lee, M. and Smith, G. (2002) Regression to the mean and football wagers, *Journal of Behavioral Decision Making*, **15**, 329–42.
Levitt, S. (2004) Why are gambling markets organized so differently from financial markets?, *Economic Journal*, **114**, 223–46.
Millman, C. (2001) *The Odds: One Season, Three Gamblers, and the Death of Their Las Vegas*, Public Affairs (Perseus Books Group), New York.
Neal, R. and Wheatley, S. (1998) Do measures of investor sentiment predict returns?, *Journal of Financial and Quantitative Analysis*, **33**, 523–47.
Pankoff, L. (1968) Market efficiency and football betting, *Journal of Business*, **41**, 203–14.
Paul, R. and Weinbach, A. (2007) Does sportsbook.com set pointspreads to maximize profits? Tests of the Levitt model of sportsbook behavior, *Journal of Prediction Markets*, **1**, 209–18.
Paul, R. and Weinbach, A. (2012) NFL bettor biases and price setting: further tests of the Levitt hypothesis of sportsbook behaviour, *Applied Economics Letters*, **18**, 193–97.
Sapra, S. (2008) Evidence of betting market intraseason efficiency and interseason overreaction to unexpected NFL team performance 1988–2006, *Journal of Sports Economics*, **9**, 488–503.
Sauer, R. D., Brajer, V., Ferris, S. P. *et al*. (1988) Holding your bets: another look at the efficiency of the gambling market for national football league games, *Journal of Political Economy*, **96**, 206–13.
Schmeling, M. (2009) Investor sentiment and stock returns: some international evidence, *Journal of Empirical Finance*, **16**, 394–408.
Strumpf, K. (2003). *Illegal Sports Bookmakers*, Department of Economics, University of North Carolina, NC.
Sturgeon, K. (1974) *Guide to Sports Betting*, Harper & Row, New York.
Vergin, R. C. (1998) The NFL pointspread market revisited: anomaly or statistical aberration, *Applied Economics Letters*, **5**, 175–79.
Vergin, R. C. (2001) Overreaction in the NFL point spread market, *Applied Financial Economics*, **11**, 497–509.
Vergin, R. C. and Scriabin, M. (1978) Winning strategies for wagering on national football league games, *Management Science*, **24**, 809–18.
Wever, S. and Aadland, D. (2012) Herd behaviour and underdogs in the NFL, *Applied Economics Letters*, **19**, 93–7.
Woodland, B. and Woodland, L. (2000) Testing contrarian strategies in the National Football League, *Journal of Sports Economics*, **1**, 187–93.
Zuber, R. A., Gandar, J. M. and Bowers, B. D. (1985) Beating the spread: testing the efficiency of the gambling market for national football league games, *Journal of Political Economy*, **93**, 800–6.