Efficiency and the Disposition Effect in NFL Prediction Markets

By Samuel Hartzmark and David Solomon*

This Version: May 2008

First Version: May 2006

* Both authors are from the University of Chicago Graduate School of Business, 5807 S Woodlawn Ave, Chicago, IL, 60637, USA. (773-576–4672). Contact via email at shartzma@chicagogsb.edu and dsolomon@chicagogsb.edu respectively. We would like to thank Arjun Chakravarti, Ricardo Cossa, Andrea Frazzini, Michael Hartzmark, Steven Levitt, Richard Thaler, Justin Wolfers, George Wu and Erez Yoeli for their helpful comments and suggestions. All remaining errors are our own

Efficiency and the Disposition Effect in NFL Prediction Markets

Abstract:

Examining betting contracts for NFL football games at Tradesports.com, we find evidence of mispricing consistent with the disposition effect, where investors prefer closing out positions at a profit than a loss. Prices for a given team to win are too low when the team gets ahead and too high when they get behind. Returns following news events exhibit short-term reversals and longer-term momentum. Higher liquidity games, with larger incentives for arbitrage, exhibit somewhat greater mispricing, without obvious limits to arbitrage. These results indicate circumstances where behavioral biases mean that prediction market prices are less likely to provide accurate forecasts.

1. Introduction

One of the important unresolved questions of economics is the extent to which laboratory evidence on individual-level decision making can be generalized to market behavior. While the psychological literature has documented many heuristics and biases that individuals use when making decisions in laboratory settings, economists have been rightly skeptical in questioning whether in real-world markets such mistakes (should they occur) will necessarily result in mispricing.[1] First, the larger stakes and greater liquidity of real-world financial markets may create a sufficient monetary incentive for arbitrageurs to correct mispricing resulting from the activities of irrational or naïve traders, whereas an experimental market may not. Second, the aggregation of many different opinions in financial markets may cause the individual mistakes of naïve or irrational traders to cancel each other out, or give greater weight to the views of rational investors, thus not affecting overall pricing efficiency.[2] Both of these questions can only be answered with empirical evidence. However the valuable tradeoff in studying experimental markets is that mispricing can often be established much more cleanly than in real-world markets.

In this paper we show that neither aggregation nor greater liquidity can necessarily be relied upon to correct mispricing that arises due to behavioral biases. Moreover, mispricing can persist even in a market that lacks any of the classic conditions limiting the ability of arbitrageurs to operate. We examine the market for betting contracts traded on NFL football games from Tradesports.com, an online limit-order driven exchange. The contracts are a classical contingent claim, paying $10 if the team in question wins the game and $0 otherwise.[3] We find significant evidence of mispricing in this market, and conclude that a behavioral bias known as the disposition effect, the tendency of investors to sell stocks that have risen in value and hold on to stocks that have fallen in value, is the reason for such mispricing. Tradesports exhibits few of the classic limits to arbitrage, as it allows short sales and has low fees (the maximum round trip cost is 0.8%). In other words, unlike many other markets, arbitrageurs in this market easily *could* correct any mispricing, but for whatever reason, they don't.

Tradesports is an important example of how in certain circumstances aggregation of many investors can actually amplify, rather than dampen, the price impact of behavioral biases. In particular, a large proportion of trades (more than 40% of total volume) occur before the game begins, at a price that remains roughly constant. As a result, during the game once the price has moved, a large proportion of the market participants all face the same (paper) gains or losses. Thus uninformed decisions to buy or sell based on the disposition effect have a large price impact, because many investors engage in uninformed buying or selling at the same time. Not only do we find evidence for the disposition effect in price reactions to new information, but also the magnitude of the effect is sufficiently large that prices are significantly inefficient even when simply comparing price with the ex post probability of the team winning. Such inefficiency also allows price-based trading strategies to generate positive expected profits net of transaction costs. More liquid games, where arbitrageurs should have greater financial incentives to correct mispricing, do not exhibit any less mispricing, and by some measures actually have greater mispricing. At least over the range of liquidity considered, the evidence is inconsistent with a hypothesis that greater liquidity and incentives for arbitrage are reflected in more accurate prices.

The results in this paper also give some guidance as to when prices from prediction markets are likely to be more accurate or less accurate. This is important, because an increasing academic and popular literature (such as Andrews, Wolfers and Zitzewitz (2003)) seeks to use prices in prediction markets as unbiased estimates of true probabilities for the purpose of formulating policy. First, prices after periods of news flows are less efficient than those during stable periods. For stable prices before kickoff, we cannot reject the null hypothesis of market efficiency, while prices during games are systematically inefficient. Second, where there are many traders facing the same gains or losses at the same time, price responses to news are likely to be particularly inefficient. This suggests that long periods of stable prices, particularly when accompanied by high volume, might predict inefficient responses to future news since they create a common reference point for those who bought or sold at that price.

We study 525 similarly structured contracts with a reasonable level of participation and liquidity (in aggregate more than $20 million worth of contracts traded). Under minimal assumptions about investor preferences, for the market to be efficient the price of a contract (which pays $10 if the team specified wins and $0 otherwise) should equal 10 times the probability of the event occurring, with anything else resulting in an arbitrage opportunity. By contrast, when we compare observed prices with the ex-post probability of the team winning, we see significant deviations in the pattern of an 'S-shape'. We find that at prices greater than approximately $6 the contract is underpriced relative to the ex-post probability of winning; i.e., for trades during game time, trades between $6.50 and $7.50 won roughly 80-90% of the time. Meanwhile at prices less than approximately $4 the contract is overpriced; i.e., for trades during game time, trades between $2.50 and $3.50 won roughly 10-20% of the time (Figure 3).

This pattern is consistent with disposition-related behavior. Because of the large volume before kickoff, many investors either bought or sold short the security at the same price, usually around $5-$7.[4] During the game, when there is a price increase pushing prices above their pre-game level, disposition-related behavior predicts that those who initially bought the contract will be eager to sell (to realize their gain), while those who initially sold short will be reluctant to buy (to realize their loss). This net uninformed selling pressure pushes prices below their correct level, creating negative short-term returns and positive long-term returns. On the other hand, during the game when there is a price increase and the price is below the pre-game level we make no prediction about what will happen, because those who initially bought the contract still face a loss, so it is unclear whether this changes their desire to sell. Similar predictions are made for negative price changes.

To test this explanation specifically, we examine reaction to news events during the game, similar to Frazzini (2006). Consistent with the disposition effect, we find that if the price is above the pre-game price, when there is a positive shock to price, then prices exhibit significantly negative price changes between -$0.09 and -$0.28 from 4 to 16 minutes after the shock, and significantly positive price momentum between $0.09 and $0.33 from these times and the end of the game. Conversely, if the price is below the pre-game price, the same positive shock causes no statistically significant short-term reversal, and rather than long term momentum we observe long-term reversal. Corresponding patterns also exist for negative shocks to price. These trends are strong evidence for the disposition effect, and are difficult to reconcile with other explanations, such as under- or over-reaction.

To examine other possible causes of the mispricing, we use a panel linear probability model regression. We find that variables associated with higher levels of liquidity on average increase the mispricing. Monday night games (with greater liquidity) have consistently larger pricing errors than non-Monday-Night games. Increases in the total population of the two teams' cities causes increase mispricing for a majority of (but not all) price levels. Moreover, prices are still significantly different from efficient levels after controlling for such factors, suggesting that greater liquidity is not eliminating the behavioral biases. We also examine whether people bet for non-financial reasons, such as supporting a particular team. Variables proxying for an imbalance of support for one team – whether teams were in the top 10 of NFL merchandise sales, and the ratio of the two city populations - do not explain mispricing at all. Moreover, the market is fairly good at incorporating specific fundamental information into price, such as the point difference and whether the home team is the underdog (Levitt (2004)), as none of these variables explain mispricing.

Finally, we document the economic significance of the phenomenon. Simple price-based trading strategies requiring trading for only 10 minutes per game would have generated Sharpe Ratios (net of maximum possible expenses) of 1.064 and profits of $16,400 from roughly 20% of the sample games, or for broader strategies Sharpe Ratios of 0.284 and profits of $65,300 (based on a maximum total investment of $533,500), from roughly 60% of games. Thus the disposition effect is economically significant in this market, and large enough to cause mispricing at the aggregate level.

The results in this paper extend the literature on prediction markets[5]. Wolfers and Zitzewitz (2004) summarizes the main papers in this field. Leigh, Wolfers and Zitzewitz (2003) and Berg and Reitz (2003) examine pricing and reaction to news in a small sample of contracts, while Oliven and Reitz (2004) looks at arbitrage opportunities in a single contract. Gil and Levitt (2007) examine the price response to news events in Tradesports World Cup Football games, and find evidence of a drift in prices after a goal is scored. Most similar to the current paper is Tetlock (2004), who examines sporting vs. financial contracts on Tradesports, and the different levels of mispricing they exhibit. Consistent with Tetlock (2004), we observe S-Shaped patterns in mispricing for sporting contracts. However, our results contrast with Tetlock (2004)'s conclusion that greater mispricing in sporting contracts is due to the non-financial motives of traders in sporting games. We are able to test directly whether non-financial motives for trade have a significant effect on price, and find that they don't. Further, we identify a potentially important alternative explanations for his results - the fact that financial contracts lack any equivalent to 'pre-game' trading, causing the disposition effect to have a reduced price impact. Moreover, the disposition effect serves as a parsimonious explanation of the S-shape in mispricing.

The paper is organized as follows: Section I describes the various theoretical explanations of the disposition effect. Section II briefly describes the issues relating to the tradesports.com market structure. Section III describes the transactions data. Section IV presents the main results, and Section V provides our conclusions.

2. The Disposition Effect and Market Efficiency

2.1 The Disposition Effect

The disposition effect, coined by Shefrin and Statman (1985), refers to the tendency of investors to sell winning stocks and hold losing stocks. At an individual level, this behavior poses some challenges to standard assumptions of rationality. In a stock market setting, first it is at odds with optimal tax loss selling and second there is evidence that stock returns are persistent via the momentum effect (Jegadeesh and Titman (1993)), which would suggest the opposite behavior ought to yield greater returns. More generally however, the disposition effect represents uninformed buying and selling, and thus poses a challenge to broader questions of market efficiency to the extent that it impacts market prices.

The existence of the disposition effect has been documented extensively, in individual investors (Odean (1998)), mutual fund managers (Wermers (2003) and Frazzini (2006)) and futures traders (Locke and Mann (2000)). While the evidence for the existence of the disposition effect is strong, the underlying cause for it remains unclear. Initial explanations focused on prospect theory and mental accounting (Kahneman and Tversky (1979) and Thaler (1985)), however more recent papers have questioned whether this can actually cause disposition behavior (see Barberis and Xiong (2006)). Anecdotally, traders in this market actually talk about the disposition effect in the chat rooms visible below the trading screen. Closing out some of your bet once the price has moved in your favor is referred to as 'covering' your position, and viewed as a smart thing to do.

2.2 Prediction Markets and Price Efficiency

In this paper, we rely on the argument that for a binary outcome contract, the price of a $1 contract must equal the probability of the event occurring. Under minimal assumptions about investor preferences and the market in general, the price must equal the true probability of the event for there to be a lack of arbitrage[6]. Stated briefly, the fact that the underlying event is a football game means that it is too brief for intertemporal substitution and discount rates to be significant factors in price determination. Also, for most investors the outcome of the game does not vary in a meaningful way with their consumption or wealth. Hence the price must equal the true probability of the event occurring. Thus if prices in the market are efficient in the sense argued by Fama (1970), we obtain the three central predictions of rational asset pricing in this context:

1. At any point in time, the price of a contract should equal the true probability of the event occurring.

2. Other than the price, no information available at the time should have significant explanatory power over the probability of a team winning.
3. Any trading strategy based on buying and selling contracts should not be able to generate positive expected returns.

One of the common critiques of prediction markets is that we cannot use prediction markets to learn about other financial markets, because even if prices are inefficient the results will not be generalizable to other settings since are investors are betting for reasons other than maximizing profit. While we cannot rule such possibilities out entirely, we can nonetheless distinguish between investor preferences that will distort prices directly and those which will not.

Specifically, to distort prices it must be the case that investors utility from gambling affects *how* they bet, not merely the decision *to* bet in the first place. If investors obtain a pure utility from gambling (such as if they like gambling specifically, if gambling heightens their enjoyment of a football game they were already going to watch etc.), but having decided to gamble still attempt to maximize profits, then this should not affect price efficiency, as having entered the market they will then behave as standard rational agents. It may affect prices indirectly through the number of participants, but this is not a problem specific to gambling settings.

One example of preferences that would affect the interpretation of prices as probabilities would be if investors prefer betting on a particular team, such as the team they support. In this paper we are able to test this directly in section 4.3.3, and find no evidence that they do. Another possibility would be if traders have different risk aversions such that they prefer buying at a particular price (and thus securing a particular distribution of payoffs). While we do not test this directly, it must be the case that the *marginal* investor prefers bets of a particular type – if investors with different risk preferences are able to hedge out each other, then this again will not affect prices. However, to the extent that one believes that investors have specific preferences over types of bets that they are unable to hedge with each other, this may affect the interpretation of the results.

3. Data

The data used in this paper comes from the online exchange Tradesports.com. This is a limit-order driven exchange that trades 23 hours a day, allows short sales, and has no centralized bookmaker or specialist setting prices.[7] The data we have includes each time-stamped transaction price for each individual contract for professional football games during the 2003-04, 2004-05 and 2005-06 seasons.[8] We use 'money-line' contracts where the payoff is based only on a favorite winning or losing (rather than 'spread' contracts where the winner is determined by whether or not the favorite wins by more than a specified number of points').

In a game of Team A versus Team B, in advance of the initiation of contract trading, Tradesports chooses whether the contract is for Team A to win (e.g., Team A is the Tradesports pick) or Team B to win (e.g., Team B is the Tradesports pick). Tradesports picks are based on the Vegas betting line when the contract starts trading, usually a few days before the game. If a buyer and seller agree on a price p and the Tradesports pick wins, then the contract purchaser receives $10 and the seller receives $0, giving a total payoff to the buyer of $(10-p) and a total payoff to the seller of $(p-10). If the Tradesports pick loses then at termination the contract purchaser receives $0 and the seller receives $0, giving a total payoff of -$p to the buyer and $p to the seller. Alternatively, positions may be closed out prior to the end of the game by taking an opposing position.

Once trading commences, these contracts can be bought (a long position created) or sold (a short position created). Therefore, there is no fixed number of contracts and, like futures or options contracts, the number of open contracts fluctuates depending on the total number of buyers and sellers. Transaction costs are low - there are no fees for execution of limit order, $0.04 per $10 contract for executing a market order, and a further fee of $0.04 for whoever is paid out at the termination of the contract.

We choose to study football contracts as they are the most actively traded contracts of a recurring nature on Tradesports. This leaves us with data on 525 games, comprising 151,853 separate trades, 2,068,807 total contracts traded and $20.7 million in dollar volume. Summary statistics are presented in Table I. We obtain city MSA populations from the 2000 census, and data from www.nfl.com about game length, quarter-by-quarter scores[9], and team merchandise sales.

[Insert Table I here]

4 Results

4.1 Tests of Mispricing

First, we test proposition 1 in section 2.2, namely *that the contract price at any time should represent an unbiased estimate of the true probability of the event occurring.* While we cannot examine the true ex ante probability of a particular contract directly, we can measure the ex-post proportion of contracts that win at each contract price. To test the accuracy of these prices, for each trade we generate a dummy variable, Win_i, that equals 1 if the Tradesports pick wins the game and 0 if the Tradesports pick loses. We then take the volume-weighted average value of *Win* over all trades at a particular price level to give the ex-ante proportion of contracts that actually win for a given price. Figure 1 plots this proportion versus all of the contract prices for in-game trading of all of the games in the sample.

[Insert Figures 1 and 2]

Figure 1 shows one of the consistent trends of contract prices in this market – an 'S-Shape', where the probability of winning is higher than expected for prices above about $6 (and contracts are undervalued), and the probability of winning is less than expected for prices (and contracts are overvalued) below about $4. This effect does not appear to be driven by thin trading. Figure 2 shows the relationship between the volume of contracts and transactions prices, indicating that the areas of large mispricing do not have unusually low volumes.[10] Prices cluster significantly towards 10c increments rather than 5c (the minimum tick size), and to multiples of 50c more than other numbers. This is consistent with Aitken et al. (1996) who argue that clustering is greater in markets where price discovery is less efficient.

Figure 1 also shows the result of fitting a curved functional form to the data (with Price/10). Specifically, we use non-linear least squares to fit a model of the form:

$$P(win) = \frac{\delta Price^{\gamma}}{\delta Price^{\gamma} + (1 - Price)^{\gamma}}$$

When δ=1 and γ=1 the function is linear. γ controls the curvature, and δ controls the elevation (roughly speaking, where the curve intersects the 45° line). We choose this particular functional form (used in Tversky and Fox (1995)) because it has the advantage of fitting a curved function that allows for different levels of curvature and elevation, while also giving a specific interpretation of whether there are signification deviations of the parameter estimates from their values under linearity. When we fit this function to the data, we find that the estimated γ is 1.37, representing the S-shape in the fitted values, and the estimated δ is 0.849, indicating that the S-shape intercepts the 45° line to the right of $5, at a value of about $6.

While both of these values are highly significantly different from 1, meaningful inferences are not possible from this because of the large significant positive serial correlation between prices of consecutive trades within the same game. To test the significance of the S-shape, we need to obtain a sample of prices that are independent of each other. While we could take only one price per game, this would result in a very small sample. To circumvent this, we take multiple observations of prices within a game that are sufficiently far apart that they do not demonstrate significant autocorrelation. While multiple prices in a game will not be truly independent, multiple prices at sufficiently long intervals apart may appear statistically independent. We take four prices per game: the 10-minute volume-weighted average prices at kickoff, quarter-time, half-time, and three-quarter-time. Using a bootstrap methodology described in the Appendix, we fail to reject the null hypothesis that prices across four quarters are independent.

Using this smaller independent sample of the quarter-time, half-time and three quarter-time prices per game (giving a total of 1023 observations), we can thus test whether the non-linearity is statistically significant. We find that the estimated values are γ=1.32 (with a t-statistic of 2.95 for the null that γ=1) and δ=0.96 (with a t-statistic of –0.56 for the null that δ=1)[11]. By contrast, kickoff prices, which are a function of the Vegas line, have a different pattern where we cannot reject the assumption of linearity: γ =0.95 (t=0.2) and δ=1.248 (t=1.1). This is consistent with the disposition effect explanation, because the relatively stable prices before the game mean that few investors face meaningful gains or losses, upon which the disposition effect relies. For in-game trading, the relationship between the price and the probability of winning is significantly non-linear, violating market efficiency. This curve is plotted in Figure 3

[Insert Figure 3 here]

The simplest interpretation of this S-Shape is that the contracts are overvalued for prices below $4, and undervalued for prices above $6 (though not for prices above $9.50[12]). Equivalently, the market views games as more uncertain than they actually are. When the Tradesports pick gets ahead, the market views the non-pick team as having a greater likelihood of coming back to win, and thus underprices the Tradesports pick, and vice versa.

Another feature of Figures1-3 is that when adjustment is made for the different axes, the graphs are similar in shape to the probability weighting function in Tversky and Kahneman (1992). The two-sided nature of the contract makes the interpretation inexact, because the same contract price is a long-shot for one party and a favorite for the other party, and hence Figures 1 and 3 might be expected to be more symmetric than Tversky and Kahneman (1992)'s Figure 1 weighting function (p310 in original). However, the large differences in fitted values between kickoff prices (for which linearity cannot be rejected) and in-game prices (where it can) are difficult to reconcile with an explanation that investors are simply applying a prospect theory weighting function. There is no particular reason to assume that investors' weighting functions ought to change over the course of the game. On the other hand, the difference in coefficients between in-game and kickoff prices is consistent with information flows affecting mispricing, such as through the disposition effect. While prospect theory may have a role in explaining the disposition effect, a prospect theory weighting function alone does not seem to explain the pattern in prices. A simple prospect theory weighting function would not predict the observed reactions to news events, discussed below in section 4.2.

4.2 Reactions to News Events

To further examine whether the disposition effect is driving the S-shape, we examine the returns available after event-driven price changes (i.e., price changes which may occur due to events like touchdowns, intercepts, etc[13]). Frazzini (2006) looks at reactions to earnings announcements and finds that the disposition effect causes investors to 'under react' to news in terms of subsequent returns. Because we have trade-by-trade prices after news events, we can take this test one step further and examine the mechanism for this apparent under reaction. 41.3% of total contracts are traded before kickoff, and prices are very stable – the standard deviation of pre-game prices[14] is 12.1¢, whereas the standard deviation of in-game prices is $1.524. Because the security is in zero net supply, we know that at the start of the game, an equal number of contracts were bought and sold short at the pre-game price. Specifically, if the disposition effect after positive news events causes an excess supply of contracts – eager sellers and reluctant buyers – then the price must fall and we should observe that after positive news short-term returns are negative, while longer term returns are positive as the price is eventually corrected as the game approaches its conclusion. This creates a pattern of initial overreaction and subsequent underreaction that is difficult to reconcile with other stories.

In the context of the Tradesports market we predict that when the price is above the pre-game price and there is favorable news pushing up the price, there will be an excess disposition-related supply at the current price, causing short-term negative returns due to selling pressure and prices being pushed below their fundamental value. This will be followed by positive long-term returns as prices return to their true equilibrium probability. On the other hand, if the same positive shock occurs when the price is below the pre-game price, then this reduces the size of gains (to short traders) and losses (to long traders), and there is no reason to predict returns of a particular sign. Similarly, a negative shock pushing down the price when the price is already below the pre-game price will produce an excess disposition-related demand at the current price, resulting in short term positive returns due to buying pressure and prices being pushed above their fundamental value. This will be followed by long-term negative returns as the price returns to its true equilibrium probability. Finally, a negative shock above the pre-game price makes no prediction about subsequent returns.

[Insert Table II here]

In Table II, we test these predictions and find evidence supporting all of them. Table II examines price changes following shocks or large price changes. We calculate volume-weighted prices for each 3-minute period of the game. We define a price shock as occurring when the price change from the subsequent period was greater than 50¢, excluding certain defined cases.[15] As an intuitive check on whether these are actual shocks, we find that after the exclusions, 50¢ changes of either sign occur roughly four times per game. We split shocks into positive and negative price changes, and then further segment the shocks into cases where the price in the period prior to the shock was above the pre-game price, and cases where it was below the pre-game price.

Consider t to index the start of a three-minute period in which the price shock occurs. If the average price during period (t-3, t) is at least 50c different from the average price during period (t, t+3) then the event is presumed to occur some time before t+3 (most likely between t and t+3, although it is possible it occurred slightly before t). The exact timing is not important, provided we are capturing returns after the event and not those due to the news itself. We then skip a minute (to eliminate the effect of the bid-ask bounce) and examine the change in price between the average price in the period subsequent to the shock (t+4, t+7) and the average price two, three and four periods subsequent to the shock (that is, for periods (t+7, t+10), (t+10, t+13) and (t+13, t+16)). We also examine price changes between those latter three periods and the end of game price (which must be either $0 or $10 depending on which team won). For the purposes of inference, we examine whether price changes, rather than returns, are significantly different from zero. This is because returns are significantly non-normal due to the price being bounded by $0 and $10[16].

Table II presents the results of these tests. Consistent with our predictions, positive shocks when the price is above the pre-game price and negative shocks when the price is below the pre-game price both demonstrate short term price reversals, and long-term price momentum. After a positive shock of 50¢ when the price is above the pre-game price, prices decrease by -$0.09, -$0.15 and -$0.28 in two, three and four periods after the shock, all statistically significant at a 1% level. Positive shocks when the price is below the pre-game price exhibit no significant reversal. Similarly, negative shocks when the price is below the pre-game price exhibit significant subsequent reversal in the following three periods, of $0.08, $0.12 and $0.12 in two, three and four periods after the shock. Negative shocks when the price is above the pre-game price exhibit no significant reversal.

Corresponding patterns of subsequent momentum are observed when comparing the price changes to the end of the game from 2, 3 and 4 periods after the shock. Positive shocks above the pre-game price show weakly significant positive price changes to the end of the game for prices in the (t+13, t+16) period (i.e., four periods after the shock), of $0.33. For negative shocks below the pre-game price, price changes to from periods 2, 3 and 4 to the end of the game are $-0.30, $-0.35 and -$0.32 respectively, all statistically significant. Positive shocks below the pre-game price actually show long-term reversals, rather than positive momentum.

These results strongly support the predictions that prices will exhibit initial reversals and subsequent momentum, depending on the price relative to the pre-game price. For negative shocks, the disposition-related buying appears to be largely confined to the second 3-minute period, as the total price change post-shock does not vary much after that. For positive shocks above the pre-game price, the subsequent selling seems to occur for a longer period[17]. Thus the evidence is mixed as to exactly how long the disposition trading occurs for. On the other hand, periods with larger reversals also exhibit larger subsequent momentum, suggesting that the reversal is indeed pushing prices beyond efficient levels. The results in Table II and IV indicate that investors reverse some of their initial response to news, but reverse it by too much, leading to an ultimate continuation in returns. An alternative story would need to explain not only this dual pattern of reversal and continuation, but also why this pattern should depend on both price and the pre-game price. It is not obvious what such an alternative story would be, and thus the responses to news indicate strong support for the disposition effect.

4.3 Regression Tests of Efficiency

In this section we test the second proposition implied by pricing efficiency, namely that *no information other than the price should have significant explanatory power over the probability of a team winning*. In particular, we seek to examine other explanations that may account for the mispricing observed earlier, such as investors trading for non-financial reasons, or mispricing due to insufficient liquidity, or investors failing to incorporate specific fundamental information about a game. Whereas the earlier tests of efficiency were designed to test whether probabilities were correct ex post at a given price level, here we take the cross section of prices at the end of each quarter and use linear probability regressions to determine if the price is able to completely explain variation in $Win_{i,}$ a dummy variable that equals 1 if the Tradesports pick wins game i, and 0 otherwise.

Thus for each game at time t, (where t is either quarter time, half time, or three quarter time) we estimate the panel linear probability regression:

$$Win_i = c + \beta_1 Price_{i,t} + \beta_2 (Controls)_i + \beta_3 (TimeFE) + \varepsilon_i$$

where the main explanatory variable $Price_{i,t}$ is the average price divided by 10 (so as to lie on an interval between 0 and 1) for game i for the 10 minutes around time t. We include a number of exogenous control variables, designed to test for variation in liquidity, non-financial reasons for trade, and fundamental information. Under the null hypothesis of market efficiency, we should observe that c=0, β_1= 1 β_2= 0, and β_3= 0.[18]

Table III presents the results of series of regressions for variables related to Liquidity (Panel A), Non-Financial Reasons for Trade (Panel B) and Fundamental Information (Panel C). Table IV shows the marginal impacts of the statistically significant non-price variables from Table V.

[Insert Table III and Table IV here]

4.3.1 Significance of Mispricing

We first observe that in nearly every specification the market is significantly inefficient. When comparing just price and *Win* in the first column of panel A, β_{Price}=1.109 and the F-test of H_0: $\beta_{\text{Price}}=1$ gives a p-value of 0.028, while the constant is –0.059 with a t-statistic of –1.68. Similar results are obtained across the different specifications, with the constant significantly negative at or near a 10% level in most specifications and β_{Price} significantly greater than one at or near a 5% level in most cases. While the addition of other variables (such as the Price Above Pre-game Price dummy in Panel B) may remove the residual mispricing, this does not resuscitate market efficiency, since in an efficient market price should be the only variable that explains the chance of a team winning. These results also support the conclusions reached earlier that the mispricing in this market is significant.

4.3.2 Analysis of Changes in Liquidity

The regression results in Panel A of Table III provide evidence that pricing accuracy is affected by variables related to market depth and liquidity. We are interested in variables that potentially cause an increase in market participation for reasons unrelated to game events. Variables like volume are not clean measures of liquidity (which we might expect to correct mispricing), because they might also increase with greater news (which under the Disposition Effect explanation can make mispricing worse). The variables instead are dummy variables for Monday Night Games (which are screened on TV nationally and are the focal NFL game that week), dummy variables for which season the game was in (as Tradesports as a whole became more liquid over the three seasons) and the log of the total population of the cities playing, as bettors may watch games involving teams that they personally follow.[19] Pre-Game Volume is also a measure of liquidity largely unrelated to news.[20] Other than Pre-Game Volume, these variables show statistically significant effects on prices, although the impact of these variables does not eliminate the mispricing entirely.

To interpret whether greater liquidity is bringing prices closer to or further from efficient levels, Table IV shows marginal effects and associated pricing errors of the significant variables in Table III. The nature of the regression means that an increase in an independent variable may improve price efficiency for some price levels but worsen it for others. The 25th, 50th and 75th percentile of prices are \$5.00, \$7.00 and \$8.59. For Monday Night Games, Table IV shows that the estimated probabilities of winning at such prices are 0.38, 0.604 and 0.783, respectively, versus 0.513, 0.738 and 0.917 for other games. Such probabilities imply pricing errors at these prices of \$1.204, \$0.956 and \$0.758 for Monday Night Games, versus \$0.134, \$0.382 and \$0.580 for all other games. In other words, the more liquid Monday Night Games exhibit greater mispricing across the interquartile range of prices than other games.

For the Season 2003 Dummy, a similar analysis shows that the mispricing is roughly symmetric – the low-volume 2003 season, when compared with the base high-volume 2005 season, is more accurate at the 25th percentile of price, roughly equally accurate at the 50th percentile, and less accurate at the 75th percentile. For Total Population, at the 25th, 50th and 75th percentile of price, the 25th Percentile of Total Population exhibits pricing errors of \$0.314, \$0.110 and \$0.052 respectively, while the 75th percentile of Total Population exhibits pricing errors of \$0.215, \$0.419 and \$0.581 respectively. Thus an interquartile increase in Total Population increases mispricing at the 50th and 75th percentile of prices, but somewhat decreases mispricing at the 25th percentile of prices.

The interpretation of these results is that the there is evidence that games with exogenously higher levels of bettor interest (particularly as measured by Monday Night games) exhibit somewhat greater mispricing. At a minimum, there is no evidence that higher liquidity is correcting any mispricing in the market.

4.3.3 Analysis of Non-Financial Motives For Trade

The regression results in Panel B of Table III provide very little evidence that bettors are motivated in their trading decisions by non-financial motives. The coefficients that measure imbalance in the size of city populations, and Dummy Variables for which teams are popular in terms of merchandise sales, show no significant effect on the team's chances of winning (after price is considered). This evidence is not consistent with Tetlock's (2004) explanation of sports prices being driven by non-financial motives for trade. Variables proxying for the most obvious non-financial motives relating to team loyalty do not show any effect. This is important, because the evidence appears inconsistent with a hypothesis that that people are betting in this market merely because they obtain non-financial enjoyment out of supporting a certain team.

4.4.4 Analysis of Fundamental Information

Panel C shows the effect of fundamental information. The point differential between the two teams and whether the home team is an underdog (see Levitt (2004)) both show no significant effect. The Period Quantity/Pre-game Quantity (proxying for the amount of news in that quarter) is insignificant. The Above Pre-Game Price Dummy (which equals 1 if the price is greater than the pre-game price, and 0 otherwise) is significant, which is not an indication of fundamental information, but is suggestive of a role for the pre-game price in terms of the disposition effect.

The results in Table III show that the mispricing in the market is not driven by low levels of liquidity, traders betting on the teams they support, or on markets not incorporating some of the more obvious sources of fundamental information about the game.

4.5 Tests of Profitability of Trading Strategies

Another more direct test of the economic impact of the disposition effect (and also the economic significance of the Tradesports mispricing generally) comes from testing proposition 3, namely that *trading strategies should not be able to generate positive expected returns.* In Table V, we evaluate strategies (based on the results in section 4.1) that involve buying contracts for some range of prices above \$5, and/or selling short contracts for some range of prices below \$5. In Panel A we measure profitability using the Sharpe Ratio of the strategy, namely:

$$SR = \frac{E(R)}{\sigma(R)}$$

where R is the return on the strategy until the end of the game.

The range of contract prices in question is given in the left two columns. Since these price ranges may not be available in all games, we report under each Sharpe Ratio the number of games in the sample (out of 525) that the strategy would have been available to be implemented. For the Sharpe Ratios, expected returns are given net of the maximum possible expenses incurred. In Panel B, we present the dollar profit that could have been obtained by taking every transaction observed in the given price range, and underneath give the total amount of money required to establish the positions.

There are a number of caveats necessary in interpreting these numbers. Firstly, there is the issue of the cost of the bid/ask spread. Unfortunately we do not have data on spreads, and placing market orders will incur a spread cost. However, the prices given are all actual transaction prices, not bid/ask midpoints, and given the limit-order nature of the market, the profits described were necessarily realized by some investors in the market. The second question is that of depth and market impact. The profits described are the actual profits that existed to investors who took up positions at the prices given. On the other hand, if additional demand or supply had been introduced by arbitrageurs following this strategy, this will eventually impact prices. The volume traded does give some indication of what kind of depth is available at these prices, however. In interpreting the available volume, the dollar profits are an overstatement in one sense, because they require the investor to take up every possible trade at a given price. On the other hand, they understate the profits available in other dimensions, as they relate to trading only for a 10-minute period surrounding each point of the game, rather than continuously trading throughout the game. Table V also greatly overstates the amount of money needed to establish the position, because in reality an investor could roll over profits from one game to establish the position in the subsequent game, rather than taking it to be an additional investment.

Given the above caveats, the results in Table V indicate that over the period in question there were significant profit opportunities from the mispricing. From a Sharpe Ratio standpoint, the most lucrative strategies involve shorting the contract at prices between \$1 and \$2 or between \$4 and \$5, and buying the contract between \$8 and \$9. The \$1-2 short and \$8-9 long strategy generates a Sharpe ratio of 1.064 at Three-Quarter time (in 107 games) and 0.447 at half time (in 98 games). The simplest strategy, of buying above \$5 and selling below \$5, generates a Sharpe Ratio of 0.284 at Three-Quarter time. To put these numbers into perspective, the stock market has a Sharpe Ratio of about 0.5, and the average bet on the NFL in Vegas has a Sharpe Ratio of between–0.03 and –0.1, depending on the size of the vig. Panel B shows the dollar profits available from such strategies. A strategy based on buying for all prices greater than $5 and selling for all prices less than $5 would have generated profits of $65,300 on a maximum required investment of $533,500, based on trading for only 10 minutes per game. Moreover, this is not simply the result of a small number of games – this strategy can be implemented in 326 out of 525 games (65% of the sample).

5 Conclusion

In this paper, we examine the efficiency of NFL betting markets at Tradesports.com and find significant evidence of the disposition effect. Prices do not reflect the true probability of a team winning, but instead follow an S-Shape, with contracts being overpriced for prices between $0 and $4, and underpriced for prices between $6 and $10. We examine the reaction to news events, and find patterns of apparent overreaction followed by underreaction that depend on whether the price is above or below the pre-game price. This evidence is very difficult to reconcile with explanations other than the disposition effect.

By analyzing the cross section of prices at different points in the game using linear probability regressions, we can reject the alternative possibilities that the mispricing is driven by structural factors such as fundamental information or non-financial reasons for trade, such as participants betting to support a particular team. Moreover, variables associated with greater liquidity actually make mispricing somewhat worse, suggesting that greater monetary incentives for arbitrage are not acting to correct the mispricing.

These results provide a comparatively clean test of arguments that prices in financial markets should be efficient based on aggregation and large monetary stakes. Aggregation of many different investor opinions need not cancel out individual mistakes when the market structure causes certain events and prices to be salient for all investors. Greater stakes due to more liquidity will not correct mispricing if the additional investors who enter the market suffer from the same behavioral biases as existing investors. The Tradesports NFL market exhibits an array of evidence in favor of the disposition effect, and difficult to reconcile with alternative explanations. Inasmuch as we observe the effects of different liquidity levels and the aggregating effect of markets, neither of them acts to correct the behavioral biases that drive the disposition effect.

Notes

[1] See for instance Kahneman (2003) for a summary of some of the research on biases and heuristics in decision-making

[2] These two criticisms are both examples of the Levitt and List (2006) critique of the generalizability of laboratory experiments. The first point reflects the fact that increasing the liquidity and stakes in laboratories will affect the self-selection of individuals into the market (Levitt and List (2006)'s point 3), while the second point indicates that the structure of the marketplace may affect process and context in which the decision is embedded (Levitt and List (2006)'s point 2).

[3] Prior to the conclusion of the NFL game an investor can pay a certain amount between \$0 and \$10 for the uncertain \$10 payout if his/her choice wins. Therefore, the value of a contract at any time reflects the probability of the favorite winning. For example, if the investor pays \$6.15 for Team A, implicitly the investor (and the market) is revealing the expectation that Team A has a 61.5 percent chance of winning.

[4] Because Tradesports chooses the team predicted to win as the Vegas favorite at the time the contract opens, the price is usually above \$5 when the game commences.

[5]

[6] Provided that the Law of One Price and Absence of Arbitrage hold, the price will only differ from the probability of the event occurring by the amount Cov(m,X), where X is the security payoff and m is the stochastic discount factor. In many rational asset pricing models, the stochastic discount factor is based around individual consumption risk. Apart from the owners of the respective teams, even amongst hard-core fans the consumption risk of a football game is limited to at most a few post-game drinks. Thus it is reasonable to assume that under the Law of One Price and Absence of Arbitrage the price must be equal to the true probability of the event occurring.

[7] Gil and Levitt (2007) argue that although there are de facto market makers, they do not affect prices.

[8] Ideally we would like to obtain trader level data recording each individual's buying and selling, but Tradesports.com's policy is to not disclose trader id numbers.

[9] In our data, we are able to obtain from http://www.nfl.com/ actual calendar times for kickoff and the end of the game, but point scores are only given in terms of game time, not calendar time. To estimate quarter time, half time and three quarter time, we take the total playing time, subtract 15 minutes in the middle for the half time break, and

divide the remaining period into quarters. Because this is necessarily inexact, prices listed as 'quarter time price' represent the volume-weighted average price during the 10-minute period surrounding the designated quarter time etc.

[10] The reason that trade volumes are higher for higher prices is the average starting price is $6.63 (shown in the figure as the vertical line) and hence the median (not mean) price movement will be positive even under an efficient market, making it likely that there will be more trades occurring at these prices.

[11] The inferences are similar and slightly stronger when using non-linear least squares with weights equal to the number of observations of each price (γ=1.31 (t=3.2) and δ=0.91 (t=-1.1)). For smaller, truly independent samples taking only one quarter's prices per game, the coefficients are similar but the significance is weaker due to the reduced sample size (γ=1.45 (t=1.97) and δ=0.93 (t=-0.55) for quarter-time, γ=1.27 (t=1.52) and δ=0.90 (t=-0.84) for half-time, and γ=1.30 (t=1.81) and δ=0.94 (t=-0.47) for three-quarter-time).

[12] The mispricing around $9.50 is possibly due to the unusual incentives created by the Tradesports fee structure at these prices - traders have an incentive to close out positions in order to avoid the 4c fee for holding the contract at the end of the game.

[13] We examine price changes, rather than game-specific events such as touchdowns, because it is not obvious that all touchdowns have significant news content, or that other events (such as interceptions) are less newsworthy. By simply examining large price changes, we can be agnostic about which events are important to a team's chances of winning, and instead rely on the market to tell us what is important.

[14] From here on, 'pre-game price' refers to the volume weighted average price between the start of trading and one hour before the kickoff.

[15] In order to exclude price changes right at the end of games, and trades based on the fee-induced effects at very high and very low prices, we eliminate shocks that occur at prices less than $0.50 or more than $9.50, and those that occur less than 15 minutes before the end of the game. The three minute average reduces the effect of spurious shocks resulting from a small volume.

[16] Returns between (t+7, t+10) and the end of the game are mechanically either -100% or some non-negative number, but cannot lie between 0% and –100%. This presents a serious obstacle to the assumption of normality. Price changes, on the other hand, can take on a range of positive and negative values.

[17] There are problems of inference in determining exactly when the reversal begins to occur, because taking longer and longer time horizons after the shock decreases the number of shocks for which there are actually still trades, which may create selection biases in terms of which types of shocks we are selecting

[18] The linear probability model allows for a very convenient interpretation of both the intercept and the slope in terms of market efficiency, whereas other models such as a Logit do not.

[19] The total population of the two cities indicates how many people may be watching the game for exogenous interest. If people actually bet disproportionately on their own team, then this will affect prices through the imbalance of the two city populations (measured later as a ratio) rather than the total number, as greater populations that are evenly split between two cities can simply hedge each other's risk. Including this as an additional control in the Total Population regressions has no impact on the coefficient or significance for Total Population.

[20] Regressing Total Volume on these variables shows that for the most part they do increase volume. The Monday Night Dummy has a coefficient of 4640 (t-stat of 6.96), the Season 2003 Dummy has a coefficient of –2430 (t-stat of –3.59), although Total Population has an insignificant effect, with a coefficient of –50.5 (t-stat of –0.14), indicating that it is a noisy measure of a team's popularity. Regressing In-Game Volume on Pre-Game Volume gives a coefficient of 0.2318 (t-stat of 18.99).

References

Aitken, Michael, Philip Brown, Christine Buckland, H. Y. Izan, and Terry Walter, 1996, “Price clustering on the Australian Stock Exchange”, *Pacific-Basin Finance Journal*, 4(2-3), 297-314

Barberis, Nicholas and Wei Xiong, 2006, “What drives the disposition effect? An analysis of a long-standing preference-based explanation”, Working Paper

Berg, Joyce and Thomas Rietz, 2003, “Prediction Markets as Decision Support Systems,” *Information System Frontiers*, 5(1), 79-93.

Fama, Eugene F., 1970. “Efficient Capital Markets: A Review of Theory and Empirical Work,” *Journal of Finance*, 25(2), 383-417

Frazzini, Andrea, 2006, “The Disposition Effect and Under-Reaction to News”, *Journal of Finance*, 61, 2017-2046

Gil, Ricard and Steven D Levitt, 2007, “Testing the Efficiency of Markets in the 2002 World Cup”, *The Journal of Prediction Markets,* 1(3), pp. 255-270

Jegadeesh, Narasimhan and Sheridan Titman, 1993, “Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency,” *Journal of Finance*, 48(1), 65-91.

Kahneman, Daniel, 2003, "Maps of Bounded Rationality: Psychology for Behavioral Economics", *American Economic Review*, 93, 1449-75.

Kahneman, Daniel, and Amos Tversky, 1979, "Prospect Theory: An Analysis of Decision under Risk", *Econometrica*, 47 , 263-291

Leigh, Andrew, Justin Wolfers, and Eric Zitzewitz, 2003, "What Do Financial Markets Think of War in Iraq?" *National Bureau of Economic Research Working Paper 9587*.

Levitt, Steven D., 2004, "How do markets function? An empirical analysis of gambling on the National Football League," *Economic Journal* 114, 495: 2043-2066.

Levitt, Steven D., and John A. List, 2006, "What do Laboratory Experiments Tell Us About the Real World?", *Working Paper,* University of Chicago

Locke, Peter R., and Steven C. Mann, 2000, "Do Professional Traders Exhibit Loss Realization Aversion?" Working paper, Texas Christian University.

Maddala, G.S. and Shaowen Wu, 1999, "A comparative study of unit root tests with panel data and a new simple test," *Oxford Bulletin of Economics and Statistics*, Special Issue, 631-641.

Odean, Terrance, 1998, "Are Investors Reluctant to Realize Their Losses?," *Journal of Finance*, 53(5), 1775-1798.

Oliven, Kenneth and Thomas Reitz, 2004, "Suckers are Born, but markets are made: Individual Rationality, Arbitrage and Market Efficiency on an Electronic Futures Market," *Management Science,* 50:3, 336-351.

Shefrin, Hersh and Meir Statman, 1985, "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence," *Journal of Finance*, 41(4), 951-974.

Shleifer, Andrei and Robert W Vishny, 1997. " The Limits of Arbitrage", *Journal of Finance*, 52(1),

35-55

Tetlock, Paul, 2004 ,"How Efficient Are Prediction Markets? Evidence from and Online Exchange" Working Paper.

Thaler, Richard H., "Mental accounting and consumer choice" (1985) *Marketing Science*, 4 , 199-214

Tversky, Amos, and Craig R. Fox, 1995, "Weighing Risk and Uncertainty", *Psychology Review*, 102, 269-283

Tversky, Amos, and Daniel Kahneman, 1992, “Advances in Prospect Theory: Cumulative Representation of Uncertainty”, *Journal of Risk and Uncertainty*, 5, 297-323

Wermers, Russ R. 2003, “Is Money Really 'Smart'? New Evidence on the Relation Between Mutual Fund Flows, Manager Behavior, and Performance Persistence,” *SSRN Working Paper #414420*.

Wolfers, Justin and Eric Zitzewitz, 2004, “Prediction Markets,” *Journal of Economic Perspectives*, 18(2).

Appendix - Testing for Independence Among 4 Quarters of Prices Per Game

It is necessary to determine whether prices at the four quarters of the game (using the 10 minute volume weighted average price at each point) are sufficiently far apart that they do not demonstrate significant autocorrelation. In other words, they will not be truly independent, but they will be sufficiently far apart that they appear statistically independent. Standard tests of autocorrelation such as the Dickey-Fuller test are inappropriate here, because an unknown autocorrelation structure is induced by the fact that prices must mechanically finish at either zero or one at the end of the game and is bounded by zero and one at all times, and thus the distribution of prices is not constant over the course of the game.

To test whether prices exhibit significant autocorrelation across these four periods, we take the sample of the 94 games that have non-missing prices at all four points in the game, and use as a test-statistics the autocorrelation function with a lag of 1: $ACF(1) = \frac{E[(\mathrm{P}_i - \mu)(\mathrm{P}_{i-1} - \mu)]}{\sigma^2}$.

We calculate the distribution of the autocorrelation function under the null of independence between prices but allowing for differing distributions of prices at each point in time. We use the 94 games to generate distributions of prices at each of the time periods. The conclusions are the same if the distribution of prices at each point in the game is taken from all games, not just those with non-missing prices for all four quarters. For each game we take the starting price as given (since this is primarily a function of the Vegas odds, but will mechanically limit how far prices can move in either direction), and randomly draw three prices from the quarter-time, half-time and three-quarter-time distributions respectively. We then calculate the sample autocorrelation function for those four prices, and repeat this 30,000 times to obtain a distribution of autocorrelation functions for independent subsequent prices conditional on the starting price. Each game is then assigned a p-value based on where the observed game autocorrelation function fits in the bootstrapped distribution. Finally, the 94 p-values are combined into a single test statistic using the Fisher test as described in Maddalla and Wu (1999), whereby $K = \sum_{i=1}^{n} -2\ln p_i$ is distributed according to a Chi-Squared distribution with $2n$ degrees of freedom.

Performing such a test, we find that the value of K from the 94 game-level p-values of the autocorrelation function is 193.47, which corresponds to a final p-value of 0.3806. In other words, when taking prices at the four quarters, we fail to reject the null hypothesis that the prices are independent based on the autocorrelation function. This motivates our use of four prices per game.

Table I: Sumary Statistics of Tradesports.com Contracts

	n	Mean	Std Dev	Max	Min
Price (in $)	151853	6.1	2.3	9.95	0.05
Pre-Game Price	61653	6.6	1.1	9.95	4.00
In-Game Price	90200	5.7	2.7	9.95	0.05
Contracts traded per Game	525	3941	6604	53678	34
Contracts Traded In-Game per Game	525	2314	5228	42592	1
Contracts per trade	151853	13.6	47.4	3900	1
Game Length (hours)	525	3.22	0.26	4.10	2.58
Pre-Game Volatility (std dev., in $)	525	0.121	0.06219	0.542	0.0319
In-Game Volatility (std dev., in $)	525	1.524	0.66764	3.7462	0
Total Games	525				
Total Wins	361				
Games 2003-2004 Season	132				
Games 2004-2005 Season	133				
Games 2005-2006 Season	260				
Total Contracts	2068807				
In-Game Contracts	1214606				
Total Trades	151853				
Earliest Game Date	5/09/2003				
Latest Game Date	15/01/2006				

This table contains summary statistics for trades of money-line contracts (ie "Team A to Win") on tradesports.com during NFL games from the 2003, 2004 and 2005 seasons. Each contract is for a payoff of $10 if the team in question wins, and $0 if the team loses.

Table II - Price Changes after Shocks

Top is Price Change in $ between the two time periods, Middle is t-stat, Bottom is number of shocks

Shock Direction	Price Criteria	(t+4,t+7) to (t+7,t+10)	(t+4,t+7) to (t+10,t+13)	(t+4,t+7) to (t+13,t+16)	(t+7,t+10) to End of Game	(t+10,t+13) to End of Game	(t+13,t+16) to End of Game
Positive	**All**	**-0.09** ***	**-0.09** *	**-0.13** **	**-0.09**	**-0.04**	**-0.02**
		(-2.65)	(-1.82)	(-2.19)	(-0.61)	(-0.32)	(-0.16)
		707	709	706	845	898	915
Positive	**Pre-shock Price > Pre-Game Price**	**-0.09** **	**-0.15** **	**-0.28** ***	**0.09**	**0.27**	**0.33** **
		(-2.56)	(-2.54)	(-3.59)	(0.47)	(1.56)	(2.00)
		359	364	357	448	486	495
Positive	**Pre-shock Price < Pre-Game Price**	**-0.08**	**-0.02**	**0.02**	**-0.29**	**-0.41** *	**-0.44** **
		(-1.46)	(-0.31)	(0.18)	(-1.62)	(-1.87)	(-2.04)
		348	345	349	397	412	420
Negative	**All**	**0.06** **	**0.09** *	**0.06**	**-0.23**	**-0.26** **	**-0.24** *
		(2.06)	(1.91)	(0.96)	(0.28)	(-1.97)	(-1.86)
		758	756	750	913	956	986
Negative	**Pre-shock Price > Pre-Game Price**	**-0.04**	**-0.05**	**-0.25**	**0.09**	**0.10**	**0.11**
		(-0.59)	(-0.40)	(-1.59)	(-1.26)	(0.35)	(0.38)
		130	130	124	173	180	184
Negative	**Pre-shock Price < Pre-Game Price**	**0.08** **	**0.12** ***	**0.12** *	**-0.30** **	**-0.35** **	**-0.32** **
		(2.42)	(2.38)	(1.84)	(-1.93)	(-2.32)	(-2.22)
		628	626	626	740	776	802

This Table presents the price changes after a shock to the price for contracts on NFL games at Tradesports.com for the 2003, 2004 and 2005 seasons. A shock is defined as when a) the volume-weighted average price (VWAP) in one 3-minute period, designated the shock period, is at least 50c different from the VWAP in the previous 3-minute period, and b) at least one minute out of the three in the second period has a volume that is in the top half of volume per minute for that game. The first three columns show the price change (in $) between the VWAP in the 3-min period starting one minute after the end of the shock period, and the relevant subsequent VWAP listed in the heading. Columns 3 to 6 show the price change (in $) between each VWAP listed (with (t+4, t+7) being the period starting one minute after the end of the shock period) and the contract price at the end of the game. The top entry in bold is the price change in dollars, the middle entry is the t-statistic, and the bottom entry is the number of shocks identified that met the criteria given (out of 525 games). *, ** and *** indicate significance at a 10%, 5% and 1% level.

Table III - Panel Linear Probability Regressions of Win/Loss Dummy on Price

Dependent Variable is a Dummy Variable that equals 1 if the team won the game and 0 otherwise

Panel A - Variables associated with Changes in Liquidity

	(1)	(2)	(3)	(4)	(5)	(6)
Price	**1.109** ***	**1.097** ***	**1.104** ***	**1.124** ***	**1.110** ***	**1.108** ***
	(22.30)	(22.05)	(22.24)	(22.69)	(22.29)	(22.33)
Constant	**-0.059** *	**-0.021**	**-0.100** ***	**-0.049**	**-0.056**	**-0.064**
	(-1.68)	(-0.55)	(-2.64)	(-1.40)	(-1.53)	(-1.50)
Season 2003		**-0.097** ***				**-0.073** **
		(-3.06)				(-2.19)
Season 2004		**-0.035**				**-0.031**
		(-1.26)				(-1.12)
Log Total Population			**0.004** ***			**0.004** **
			(2.69)			(2.46)
Monday Night Game				**-0.134** ***		**-0.121** ***
				(-3.94)		(-3.31)
Pre-Game Volume					**-0.001**	**0.003**
(in 1000s)					(-0.25)	(0.49)
R2	**0.328**	**0.334**	**0.333**	**0.338**	**0.328**	**0.346**
Obs	**1023**	**1023**	**1023**	**1023**	**1023**	**1023**
F Stat for H0: Beta(Price)=1	**4.83**	**3.81**	**4.37**	**6.27**	**4.85**	**4.73**
P-Value	**0.028**	**0.051**	**0.037**	**0.012**	**0.028**	**0.030**
Quarter FE	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**

This table presents panel linear probability regressions for tradesports.com contracts on NFL games from the 2003, 2004 and 2005 seasons. The Dependent variable in each regression is a Dummy variable that equals 1 if the Tradesports pick team wins and 0 if the team loses. 'Price' is the price for a $10 contract (divided by 10) for three observations per game, taken as the average for the 10 minutes surrounding the quarter-time, half-time, and three-quarter-time in each game. Panel A presents variables associated with different levels of liquidity: Dummy variables for which season the game is in (Season [ ]), variables for the log of the combined metropolitan population as reported in the 2000 census of the two cities that the teams are from (Total Population), a dummy variable for whether the game was played on Monday night (Monday Night Game) and the number of contracts traded between the opening of trade and one hour before kickoff (Pre-Game Volume). Panel B presents variables associated with non-financial reasons for trade: the ratio of the populations of the two cities (Favorite Population/ Non-Favorite Population), and dummy variables for whether the favorite, underdog, or both were in the top 10 merchandise sales of NFL teams in the previous season (Favorite in Merchandise Top 10, Non-Favorite in Merchandise Top 10, Both in Merchandise Top 10). Panel C presents variables associated with fundamental information: the difference between the team scores at each time (Point Differential), a dummy variable for whether the home team is favored to win before the game (Home Underdog), a dummy variable for whether the period price is greater than the average price before kickoff (Price Above Pre-Game Price) and the volume in the previous quarter divided by the volume before kickoff (Volume/Pre-game Volume). In each row, the top entry in bold is the coefficient, the bottom entry in parentheses is t-statistic for that estimate. *, ** and *** denote significance at a 10%, 5% and 1% level respectively.

Panel B - Variables associated with Non-Financial Reasons for Trade					
Price	**1.110** ***	**1.108** ***	**1.110** ***	**1.109** ***	**1.108** ***
	(22.31)	(22.24)	(22.29)	(22.29)	(22.21)
Constant	**-0.065** *	**-0.065** *	**-0.062** *	**-0.060** *	**-0.077** **
	(-1.85)	(-1.81)	(-1.73)	(-1.70)	(-2.02)
(Favorite Population) / (Non-Favorite Population)	**0.003**				**0.003**
	(1.13)				(1.07)
Favorite in Merchandise Top 10		**0.018**			**0.026**
		(0.73)			(0.88)
Non-Favorite in Merchandise Top 10			**0.011**		**0.019**
			(0.41)		(0.56)
Both in Merchandise Top 10				**0.009**	**-0.033**
				(0.23)	(-0.62)
R2	**0.329**	**0.329**	**0.328**	**0.328**	**0.330**
Obs	**1023**	**1023**	**1023**	**1023**	**1023**
F Stat for H0: Beta(Price)=1	**4.87**	**4.68**	**4.86**	**4.84**	**4.66**
P-Value	**0.028**	**0.031**	**0.028**	**0.028**	**0.031**
Quarter FE	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**

Panel C - Variables associated with Fundamental Information and the Disposition Effect					
Price	**1.098** ***	**1.113** ***	**0.987** ***	**1.121** ***	**1.049** ***
	(17.83)	(22.35)	(13.19)	(21.85)	(12.76)
Constant	**-0.052**	**-0.071** **	**-0.021**	**-0.069** *	**-0.056**
	(-1.29)	(-1.98)	(-0.53)	(-1.83)	(-1.20)
Point Differential	**0.001**				**0.000**
	(0.32)				(0.03)
Home Underdog		**0.037**			**0.031**
		(1.39)			(1.11)
Price Above Pre-Game Price			**0.079** **		**0.047**
			(2.18)		(1.22)
Volume/Pre Game Volume				**0.048**	**0.053**
				(0.49)	(0.53)
R2	**0.328**	**0.329**	**0.331**	**0.336**	**0.338**
Obs	**1023**	**1023**	**1023**	**956**	**956**
F Stat for H0: Beta(Price)=1	**2.52**	**5.11**	**0.03**	**5.55**	**0.36**
P-Value	**0.112**	**0.024**	**0.864**	**0.019**	**0.550**
Quarter FE	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**

Table IV - Marginal Effects of Statistically Significant Variables from Linear Probability Regressions

Pctile of Price	25th	50th	75th	25th	50th	75th
Price ($)	**5.00**	**7.00**	**8.59**	**5.00**	**7.00**	**8.59**
	Estimated Prob(win)			**Implied Pricing Error (in $)**		
Total Population (mil.)						
Min = 1.33	0.457	0.678	0.853	0.429	0.221	0.056
25th Pctile = 4.93	0.472	0.693	0.869	0.276	0.069	0.097
Median = 8.43	0.487	0.708	0.883	0.128	0.080	0.245
75th Pctile = 12.10	0.503	0.724	0.899	0.028	0.236	0.401
Max = 42.40	0.631	0.852	1.028	1.313	1.520	1.685
Monday Night Game						
Min = 0	0.513	0.738	0.917	0.134	0.382	0.580
Max = 1	0.380	0.604	0.783	1.204	0.956	0.758
Season 2003						
Min = 0	0.528	0.747	0.922	0.028	0.047	0.063
Max = 1	0.431	0.651	0.825	0.069	0.049	0.034
Price Above Base Price						
Min = 0	0.473	0.670	0.827	0.027	0.030	0.032
Max = 1	0.552	0.750	0.907	0.052	0.050	0.048

This Table presents the fitted values of the linear probability regressions in Table III and the pricing errors implied by such values, for tradesports.com contracts on NFL games from the 2003, 2004 and 2005 seasons. In each case, the coefficients are taken from the regression in Table III of a dummy of whether the team won the game on Price and the Independent Variable listed at the left. Columns titled 'Estimated Prob(win)' are the fitted values of the regression using the price listed at the top of the column at the value of the independent variable at the left of that row. The columns titled 'Implied Pricing Error' equals 10*|Price/10 - Estimated Prob(win)| , equal to the difference between the actual price and the efficient price.

Table IV: Sharpe Ratios and Profits Calculated For Various Trading Strategies

Panel A: Sharpe Ratios

(Sharpe Ratio net of maximum possible expenses above, Number of games for which strategy is possible below)

Short	Long	Kickoff			1st Quarter			Half Time			Third Quarter		
		S	L	S+L	S	L	S+L	S	L	S+L	S	L	S+L
P<5	**P>5**	**0.100**	**0.134**	**0.129**	**0.323**	**0.205**	**0.250**	**0.030**	**0.109**	**0.083**	**0.381**	**0.196**	**0.284**
		31	450	463	73	280	325	93	282	358	120	233	326
4<P<5	**5<P<6**	**0.087**	**0.288**	**0.226**	**0.225**	**0.255**	**0.239**	**0.295**	**-0.047**	**0.073**	**0.950**	**0.413**	**0.516**
		31	168	182	57	93	122	46	68	98	54	50	79
3<P<4	**6<P<7**		**0.047**	**0.055**	**0.272**	**0.282**	**0.278**	**-0.214**	**0.218**	**0.066**	**0.231**	**0.036**	**0.181**
			195	196	27	89	116	37	81	118	62	60	116
2<P<3	**7<P<8**		**0.173**	**0.173**	**0.873**	**0.454**	**0.620**	**-0.221**	**0.192**	**0.016**	**0.445**	**0.243**	**0.382**
			135	135	17	106	123	27	71	98	46	64	110
1<P<2	**8<P<9**		**-0.363**	**-0.363**	**17.644**	**-0.023**	**0.003**	**2.021**	**0.356**	**0.447**	**1.450**	**0.836**	**1.064**
			52	52	2	77	79	18	80	98	35	72	107
0<P<1	**9<P<10**		**0.870**	**0.870**		**0.326**	**0.326**	**-4.658**	**0.543**	**0.191**	**0.206**	**-0.168**	**-0.131**
			9	9		32	32	2	70	72	25	110	135

Panel B: Total Possible Profits

(Profit Net of Maximum Possible Expenses x $1000 above, Maximum amount of capital required to implement strategy below)

Short	Long	Kickoff S	Kickoff L	Kickoff S+L	1st Quarter S	1st Quarter L	1st Quarter S+L	Half Time S	Half Time L	Half Time S+L	Third Quarter S	Third Quarter L	Third Quarter S+L
P<5	**P>5**	**4.4**	**28.5**	**32.9**	**23.1**	**23.1**	**46.2**	**1.2**	**17.0**	**18.2**	**44.4**	**20.9**	**65.3**
		41.3	355.8	397.1	111.9	225.1	337.0	88.9	259.5	348.4	259.6	274.0	533.5
4<P<5	**5<P<6**	**3.9**	**22.1**	**26.0**	**7.7**	**8.8**	**16.5**	**6.8**	**-1.6**	**5.1**	**8.7**	**14.4**	**23.2**
		40.5	92.1	132.6	41.9	40.5	82.4	27.0	41.9	69.0	16.1	45.2	61.3
3<P<4	**6<P<7**	**0.5**	**3.8**	**4.3**	**3.2**	**5.7**	**8.9**	**-3.4**	**6.8**	**3.3**	**11.9**	**0.6**	**12.5**
		0.8	95.6	96.4	18.5	31.3	49.8	22.7	47.8	70.5	81.5	23.1	104.6
2<P<3	**7<P<8**		**10.8**	**10.8**	**11.4**	**7.5**	**19.0**	**-3.7**	**3.8**	**0.2**	**14.4**	**3.4**	**17.8**
			122.5	122.5	47.6	41.3	89.0	23.8	39.7	63.5	80.4	29.8	110.2
1<P<2	**8<P<9**		**-8.5**	**-8.5**	**0.8**	**-0.5**	**0.3**	**2.3**	**5.7**	**8.0**	**8.8**	**7.6**	**16.4**
			42.6	42.6	3.9	79.8	83.6	14.6	67.3	81.8	56.6	63.3	119.8
0<P<1	**9<P<10**		**0.2**	**0.2**		**1.5**	**1.5**	**-0.7**	**2.3**	**1.5**	**0.6**	**-5.1**	**-4.5**
			2.9	2.9		32.2	32.2	0.8	62.8	63.5	25.0	112.7	137.7

This table presents the Sharpe Ratios and Total Possible Profits for various trading strategies on Tradesports.com NFL contracts during the 2003, 2004 and 2005 seasons. The Sharpe Ratio is defined as E(R)/σ(R), where E(R) is the expected return from the strategy and σ(R) is the standard deviation of returns for the strategy. Sub-columns 'S' shorts the contract if the price criteria in the 'Short' column is met, Sub-column 'L' buys the contract if the criteria in 'Long' is met, and sub-column 'S+L' shorts the contract if the 'Short' criteria is met and buys the contract if the 'Long' criteria is met. The strategies are evaluated using all traded contracts during the 10-minute interval surrounding each of the time points given (kickoff, 1st Quarter, Half Time, and 3rd Quarter. In Panel A, in each row the top value in bold is the Sharpe Ratio of the strategy, the lower value is the number of contracts where the strategy was able to be implemented. In Panel B, the top value is the total amount of money possible from the strategy over that 10 minute period, and the bottom value is that maximum amount of money needed to establish the position under Tradesports.com's margin requirements.

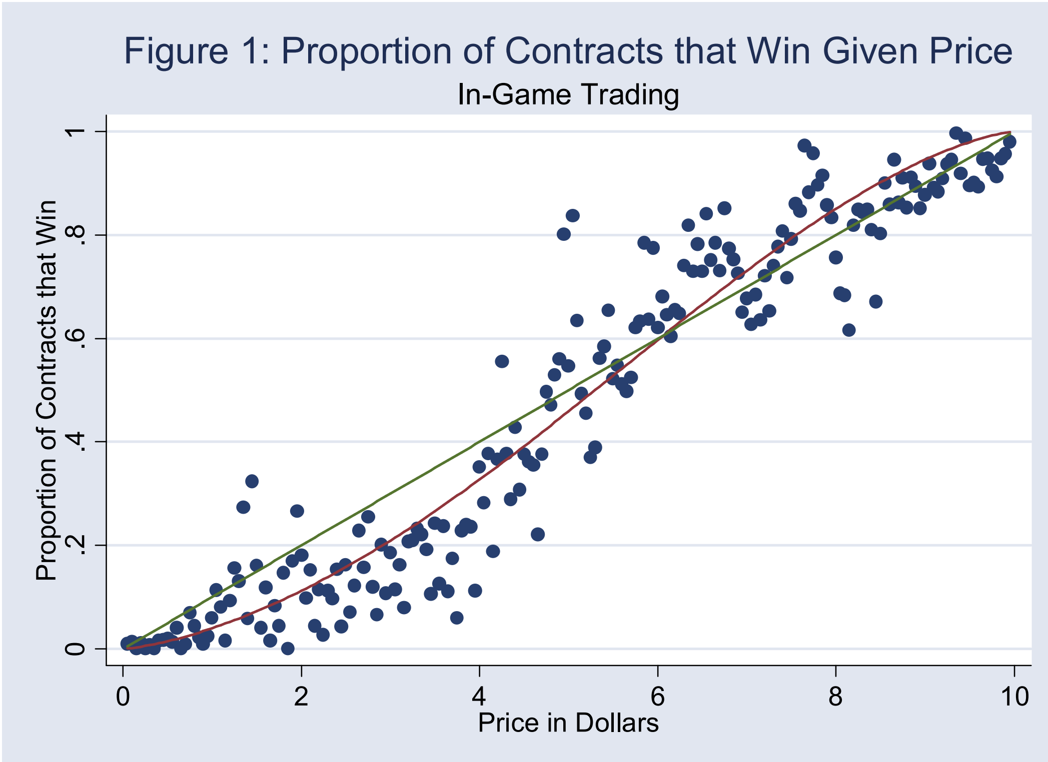

Figure 1 presents data Tradesports.com of contracts from the 2003, 2004 and 2005 seasons of NFL football of prices of all trades during game play. The plotted points are show for the price (in increments of 5c) on the x-axis vs. the fraction of cases at that price where the team actually won the game over all trades at the given price level on the y-axis (that is, the mean value of a dummy variable that equals 1 if the team in question won the game). The straight line is the 45° line, which is what prices should be under the null hypothesis of market efficiency. The curved line is the fitted values of a weighting function of the form $P(win) = \frac{\delta Price^{\gamma}}{\delta Price^{\gamma} + (1 - Price)^{\gamma}}$

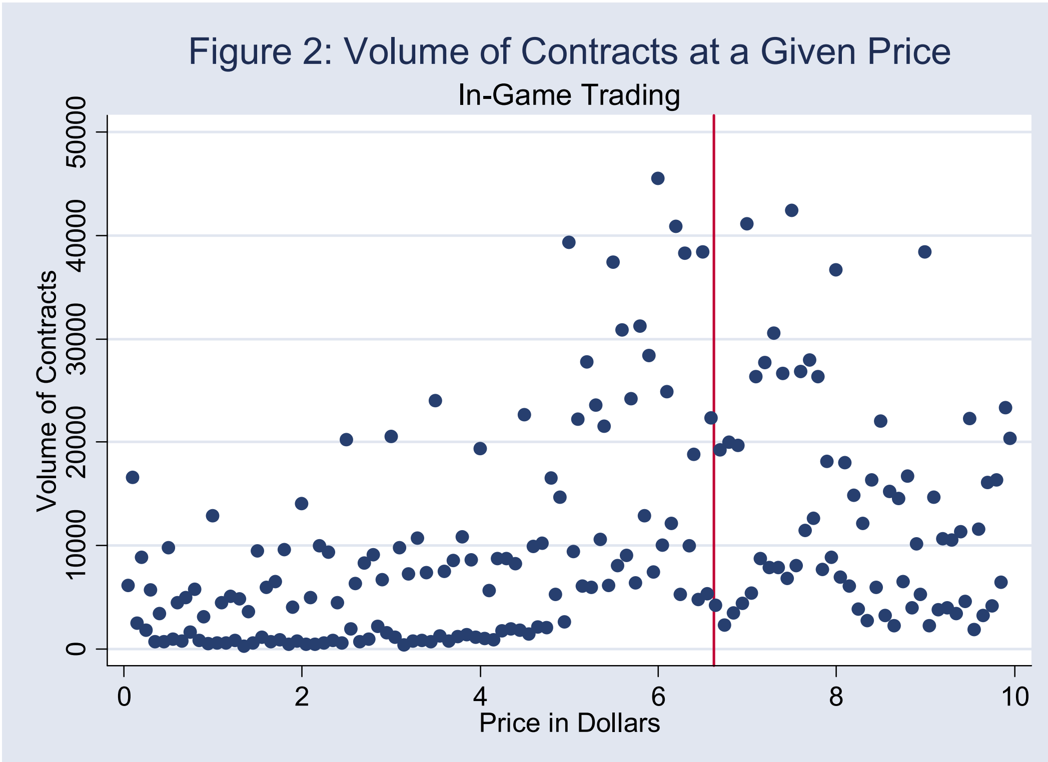

Figure 2 compares price and volume from Tradesports.com contracts from the 2003, 2004 and 2005 seasons of NFL football for trades that happened during game play. The plotted points are shown for the price (in increments of 5c) on the x-axis vs. the total volume traded at that price level on the y-axis. The vertical line is the average Pre-Game price, $6.63.

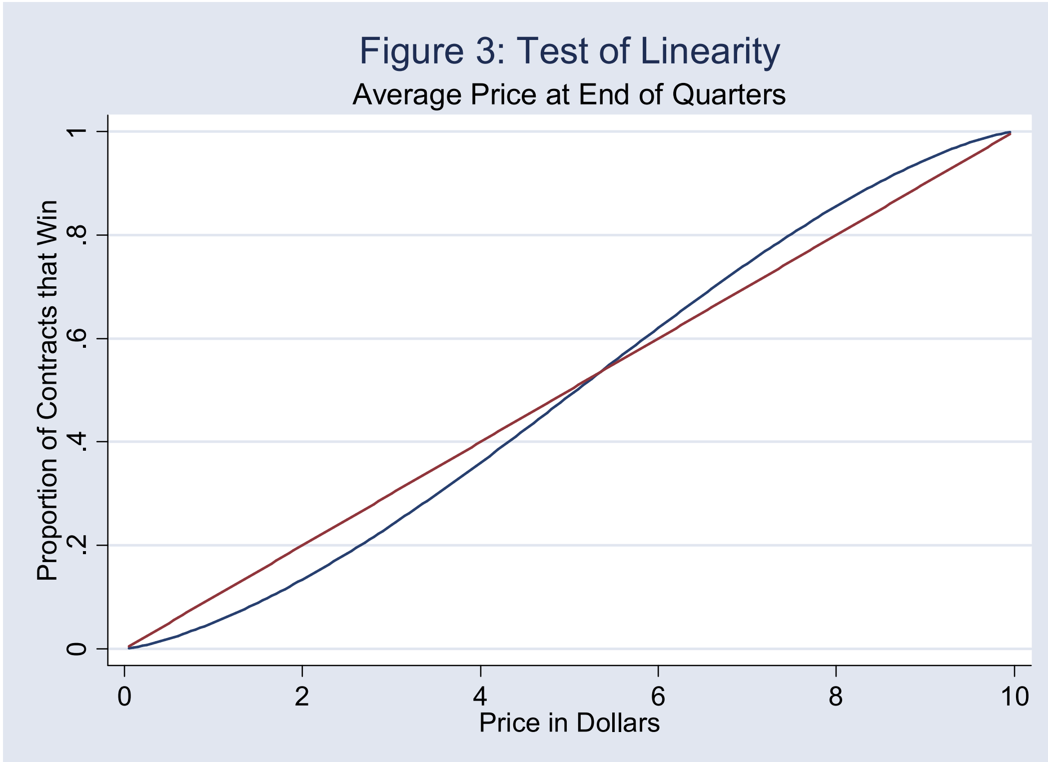

Figure 3: Test of Linearity

Figure 3 presents fitted values from a weighting function fitted to Tradesports.com data of contracts from the 2003, 2004 and 2005 seasons of NFL football. For each game, prices are taken as the 10-minute volume weighted average price at estimated times corresponding to quarter-time, half-time and three-quarter-time. The weighting function is fitted to price (as the x variable) vs. a dummy variable that equals 1 if the team in question won the game (as the y variable). The curved line is the fitted values of a weighting function of the form $P(win) = \frac{\delta Price^{\gamma}}{\delta Price^{\gamma} + (1 - Price)^{\gamma}}$. The straight line is the 45° line, which is what prices should be under the null hypothesis of market efficiency.