

PRICE BIASES IN A PREDICTION MARKET: NFL CONTRACTS ON TRADESPORTS

*Richard Borghesi**

Texas State University, McCoy College of Business Administration, Department of Finance and Economics, 601 University Drive, San Marcos, TX 78666

We examine deviations between the prices and values of binary options listed on Tradesports.com, an online prediction market. Our analysis shows that NFL sides contracts are overpriced on average, indicating that this market may be characterized by a shortage of sellers. We also find that overpricing is more pronounced immediately after information shocks occur, especially when the news is negative. Additionally, while prior research suggests that differences between asset prices and values should be symmetric around the market-price-midpoint of \$50, we find that this divergence is instead larger for low-priced contracts. Finally, we demonstrate that a simple rule designed to exploit the identified biases enables a highly profitable trading strategy.

JEL Codes: G12, G14

I. INTRODUCTION

The size and importance of prediction markets have grown rapidly throughout the past decade. In particular, public exchanges such as the Iowa Electronic Markets (IEM) and Tradesports have received much attention from researchers and politicians alike. In general, studies show these markets to be effective information aggregators and event predictors. For instance, Berg, Nelson, and Rietz (2003) show that the IEM project, a real-money online exchange in which participants wager on the outcomes of political events, is an extremely precise forecaster of election results. Not only has the IEM consistently outperformed opinion polls in accuracy (Berg, Nelson, and Rietz, 2000), but similar exchanges perform well even when play money is used (Servan-Schreiber, Wolfers, Pennock, and Galebach, 2004).

Another example of the reliability of public prediction markets is the accuracy of wagers on box office revenues and Hollywood award winners (Pennock, Lawrence, Giles, and Nielsen, 2001). Furthermore, private betting exchanges can be valuable tools in corporate settings. For example, Chen and Plott (2002) show that an intra-firm exchange can accurately forecast future equipment sales, while Ortner (1998) provides evidence that a within-firm prediction market is useful at estimating the likelihood of meeting project deadlines.

Public exchanges such as Tradesports are valuable settings in which to test trader behavior, as one can compare clearly observable outcomes to market expectations for a wide variety of events. Because the true value of each prediction market asset is revealed with certainty upon contract expiry, the joint hypothesis problem is mitigated. In addition, expectations can be readily quantified because, as Wolfers and Zitzewitz (2005b) demonstrate, prediction

*Tel.: (512) 245-1733. Email: rickborghesi@txstate.edu

market prices approximate the central tendency of the distribution of trader beliefs. Therefore, in markets where contracts expire at either \$0 or \$100, the price (in dollars) that a trader is willing to pay for a contract is expected to equal her subjective estimate of the probability (in percent) that contract will expire at \$100.

So, if prediction markets are efficient, a plot of contract prices vs. realized win rates should approximate a 45° line (a 'win' is to be defined as expiry at \$100). In studying a variety of exchanges, Wolfers and Zitzewitz (2005c) find that market prices rapidly respond to new information and that few arbitrage opportunities exist. However, Wolfers and Zitzewitz (2005a) also show that observed win rates of contracts traded on the IEM differ from expectations within certain price bands. Thus, it remains an open question whether prediction markets are efficient.

This study is one of the first to examine the within-game financial-market-style of betting that occurs on Tradesports, a setting that may be particularly useful at revealing trader preferences and illustrating a potential framing problem associated with short sales. The contracts considered represent bets on outcomes of National Football League (NFL) games. But, unlike traditional gambling markets which prohibit betting after an underlying event begins, Tradesports allows participants to place bets at any point before game outcomes are decided.

Furthermore, because each NFL game is broadcast live via a variety of media, all traders gain simultaneous access to an identical set of news innovations (within-game scores) as they transpire. Thus, the Tradesports exchange provides a unique opportunity to quantify how quickly and accurately markets respond to information shocks. For instance, we can examine whether traders react efficiently to the news that a team has scored a touchdown by comparing contract prices to underlying asset values following such occurrences.

We demonstrate that asset values (measured by mean expiry payouts) differ significantly from prices only between kickoff and contract expiry, a period during which significant news arrives relatively rapidly. Additionally, while the divergence between values and prices is predicted to be symmetric around \$50, we find evidence suggesting that they deviate furthest from each other when prices are between \$20 and \$60. We then show that assets are more likely to be overpriced immediately after information shocks arrive, and propose a set of trading rules to exploit these price biases.

Our findings are potentially interesting because widespread participation in prediction markets is a relatively new phenomenon, and many fundamental questions remain unanswered. Section II describes the data used in this study. Section III provides an analysis of the data and results, and concluding remarks are presented in Section IV.

II. DATA

The online exchange operated by Tradesports.com is a continuous double auction in which traders can place market and limit orders. All assets are

binary options with prices ranging from \$0 to \$100 for a 10-contract lot (henceforth a ‘contract’), and the expiry value of each contract is either \$0 (henceforth a ‘loss’) or \$100 (henceforth a ‘win’). Contract expiry values depend upon the outcomes of various underlying events such as equity price movements, election results, or sporting event scores. For instance, Tradesports recently listed a contract titled ‘DOW.22JUNE.-100.’ The contract expires at \$100 if the Dow Jones Industrial Average decreases by more than 100 points on June 22, and at \$0 otherwise.

Our dataset consists of execution times, prices, and quantities detailing 336,299 trades of sides contracts associated with 999 preseason, regular season, and postseason football games from the 2002, 2003, and 2004 NFL seasons.² During this period, total wealth wagered across all Tradesports NFL sides contracts was \$31,400,000, and the exchange charged a commission of \$0.40 per contract upon trade execution and expiry.

Sportsnetwork, a provider of streaming sports information, employs observers to monitor live game feeds of all NFL games and to log play-by-play information such as down, distance, yards gained, penalties, etc. On all scoring plays, an observer immediately enters an initial log that briefly states which team has scored, and this data is automatically stamped with a date and time. For instance, the text keyed into the log might be ‘NY Giants TD,’ and the system records the event as ‘09/05 22:17:23 ET NY Giants Touchdown – Details to follow.’ A subsequent entry would contain additional descriptive information such as the number of yards gained on the scoring play and the name(s) of the player(s) involved.

To match the arrival of news innovations with corresponding fluctuations of market price, we supplement our contract dataset with Sportsnetwork’s logs describing the 4,599 touchdowns scored in the NFL from the 2002 through 2004 seasons. While NFL restrictions prevent immediate dissemination of game information, all events are logged into Sportsnetwork’s database and time stamped within seconds of occurring. Sportsnetwork’s Director of Operations estimates that, on average, the amount of time that passes between the time that an NFL score occurs and the time that the initial log is recorded is approximately five seconds, while the maximum delay is 15 seconds. For all analyses in this study, we use the initial log time stamps to match game events with corresponding trades.

III. ANALYSIS AND RESULTS

In exploring the relationship between prices and values, we begin by plotting execution prices (grouped into bands \$2.50 wide) against observed win rates for contracts bought or sold on Tradesports *before* kickoff, the point at which underlying events become active. If this market were perfectly efficient, then we should observe that contracts purchased at a price of $\$P$ win at a rate not different from $P\%$ on average.




FIGURE 1. Pre-kickoff contract prices and win rates. This plot is generated from 119,735 NFL contract trades occurring in the five-day span before kickoff. Each data point represents the observed mean win rate of bets placed at each price level (grouped into \$2.50 bands). A win is defined as expiry at \$100. Error bars indicate 95% confidence limits after adjusting standard errors for the effects of clustering.

The relationship emerging from our data is contrasted against the 45° efficiency line in Figure 1, and error bars indicate 95% confidence limits after adjusting standard errors for the effects of clustering (Woodruff, 1971).³ While test power is limited due to relatively thin trading, in the five-day period prior to kickoff we find little evidence that contract win rates systematically deviate from prices. This result differs from that documented for the IEM political exchange, where Wolfers and Zitzewitz (2005a) show that contract win rates deviate from expectations in a quasi-sinusoidal fashion, and are significantly lower (higher) than expected around \$25 (\$75).

Figure 2 illustrates that once underlying events become active (upon kickoff), prices deviate from values in a manner similar to that demonstrated on the IEM. But, unlike IEM contracts, NFL contracts expire at \$100 less frequently than expected on average. The mean observed price movement from purchase to expiry for contracts traded while associated games are underway is -4.54 ticks with a standard error of 1.76 ticks, where a tick is equivalent to \$1. This result does not obtain when we examine only trades executed before kickoff (mean = -1.39 ticks; standard error = 1.87 ticks). The difference between these two means is significant (p -Value = 0.0182), suggesting that there are important dissimilarities between pre-event and within-event market dynamics.

The primary differences between these two distinct trading periods are the rate of information flow and the significance of arriving news packets. The flow rate is potentially important because information (and thus price) moves slowly before underlying events begin, so traders may be able to engage in cross-exchange arbitrage operations (via other Internet sites or




FIGURE 2. Post-kickoff contract prices and win rates. This plot is generated from 216,564 NFL contract trades occurring after kickoff. Each data point represents the observed mean win rate of contracts purchased at each price level (grouped into \$2.50 bands). A win is defined as expiry at \$100. Error bars indicate 95% confidence limits after adjusting standard errors for the effects of clustering.

traditional sports books). Alternatively, the extra time elapsing between consecutive pre-kickoff news events may provide participants with a superior opportunity to reevaluate prior beliefs, and thus more accurately estimate asset value. News importance also varies because, before kickoff, any arriving information is likely to be a less meaningful signal concerning game outcome than would be a within-game touchdown.⁴ We next further explore the systematic price biases that emerge after kickoff.

A. *Utility maximization*

As discussed in Wolfers and Zitzewitz (2005b), the quasi-sinusoidal pattern appearing in Figure 2 may be a product of exchange structure. For prediction markets having a \$0 to \$100 asset pricing format, and for constant relative risk aversion (CRRA) between 0 and 1, an equilibrium occurs in which price is between the market’s mean subjective valuation and \$50. The result is that contracts costing less (more) than \$50 may be overpriced (underpriced). To illustrate the general case, suppose that agents have CRRA utility function $u(x) = x^\alpha$ and that the price of asset A is p . If agent i with wealth w_i has belief b_i that event A will occur, then her expected utility is:

$$(1) \quad E[u(q_i)] = b_i(w_i + (1 - p)q_i)^\alpha + (1 - b_i)(w_i - pq_i)^\alpha,$$

where q_i is the quantity of asset A purchased by agent i . The optimal bet size q_i is the solution to:

$$(2) \quad \frac{\partial E}{\partial q} = b_i[\alpha(w_i + (1 - p)q_i)^{\alpha-1}(1 - p)] + (1 - b_i)[\alpha(w_i - pq_i)^{\alpha-1}(-p)] = 0.$$

Solving, we obtain:

$$(3) \quad q_i = \frac{b_i^{(1-\alpha)^{-1}} (1-p)^{(1-\alpha)^{-1}} - (1-b_i)^{(1-\alpha)^{-1}} p^{(1-\alpha)^{-1}}}{b_i^{(1-\alpha)^{-1}} (1-p)^{(1-\alpha)^{-1}} p + (1-b_i)^{(1-\alpha)^{-1}} p^{(1-\alpha)^{-1}} (1-p)} \cdot w_i.$$

This expression reveals that when agents' beliefs diverge, supply may not equal demand. For instance, if agent 1 has belief $b_1 = 0.3$ that event A will occur, then her asset demand at $p = 0.40$ is $q_1(0.3) = -0.767w_1$. If agent 2 has belief $b_2 = 0.5$ that event A will occur, then her asset demand at $p = 0.40$ is $q_2(0.5) = 0.833w_2$. Thus, the price of asset A needs to rise if supply is to equal demand, so the equilibrium price in this case will be above the mean of the agents' beliefs. Figure 3 depicts the resulting supply and demand at all prices when the mean of the agents' beliefs is 0.40. We next offer a generalized example to show why prices may differ from expected values within particular price bands on the Tradesports exchange.

To illustrate the expected excess demand at all price levels when agents' beliefs differ, suppose that Tradesports agent 1 is pessimistic ($b_1 = p - 0.10$) and Tradesports agent 2 is optimistic ($b_2 = p + 0.10$). Here, p is the true probability of an event occurring and is equal to $Price/\$100$. Figure 4 shows the expected supply, demand, and excess demand on the Tradesports exchange at each price level in the presence of pessimistic agent 1 and optimistic agent 2. For example, if each agent's wealth is \$100 and the market price is \$40, then agent 1 is willing to supply 1.92 contracts while agent 2 demands 2.08 contracts.




FIGURE 3. Expected demand for prediction market contracts. This plot illustrates the differences between supply and demand under a CRRA utility model with $u(x) = x^{0.5}$ when price = 0.40, the beliefs of two agents are $b_1 = 0.30$ and $b_2 = 0.50$, and each agent i has wealth $w_i = 1$.

PRICE BIASES IN A PREDICTION MARKET

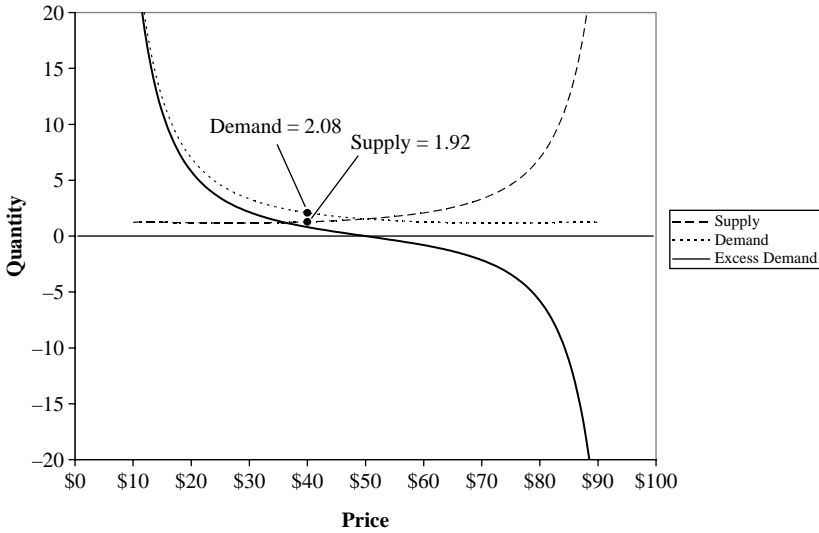


FIGURE 4. Expected excess demand for contracts on tradesports. This plot illustrates the predicted differences between supply and demand under a CRRA utility model with $u(x) = x^{0.5}$. The probability of an event occurring is $p = \text{Price}/\$100$, pessimistic agent 1 has belief $b_1 = p - 0.10$, optimistic agent 2 has belief $b_2 = p + 0.10$, and each agent i has wealth $w_i = \$100$. Quantity is expressed as the number of 10-contract lots supplied and demanded at each price level.

Thus, the market is characterized by a supply shortage and contracts purchased at \$40 are overpriced (the mean win rate is below p). Importantly, the resulting plot illustrates that the imbalance between supply and demand (the solid curve) under the CRRA model should be symmetric around \$50, and is expected to be greatest when prices diverge furthest from \$50. We next contrast the predictions of this model with observed price behavior on the Tradesports NFL exchange when underlying events are active, a period in which market participants face the rapid arrival of significant news innovations.

B. Asset mispricing

Selecting the appropriate metric to quantify post-score price movements is imperative. Rather than attempt to establish a baseline asset price at $t = 0$ (the moment that the score is realized), we choose measure the amount of price movement from trade execution to contract expiry. There are several reasons for this decision. The prices of transactions occurring just prior to $t = 0$ would not provide a meaningful baseline because important information has yet to be revealed. Utilizing the price of the first post-score trade as the baseline would also be problematic because several minutes sometimes pass before this first trade occurs, thus stale prices become an obstacle.

Furthermore, as we later demonstrate, price tend to drift downwards after $t = 0$ so using an average of post-score prices over the t minutes following each score would be problematic. The metric *Returns to Expiry* would again be suboptimal because all contracts that expire at \$0 would have return rates of -100% , and return rates of contracts that expire at \$100 would be positively

skewed. To circumvent these difficulties, we choose to use the metric *Ticks to Expiry* which is the difference between the expiry and execution prices of each contract.

As shown in Figure 2, assets that are priced far below (above) \$50 win less (more) often than expected. This is consistent with the utility maximization model presented above. However, the plot in Figure 2 is not symmetric around \$50, and assets within the band \$50 to \$60 appear to be overpriced. Furthermore, the win rates for all NFL contracts are too low on average, as mean *Ticks to Expiry* is -4.54 , which is significantly different from zero (p -Value = 0.0099). None of the latter three characteristics is true of assets traded on the IEM exchange, nor predicted by the utility maximization model. These unexpected findings can potentially be explained by the framing of Tradesports contracts, gambler preferences, and structural differences between the assets in prediction markets and those in traditional casino-type betting markets.

Each NFL contract listed on Tradesports is framed with respect to the favorite, and prior research shows that unsophisticated football bettors prefer to take the stronger team when betting against the spread (Levitt, 2004). Additionally, as sports prediction markets are relatively new, it is reasonable to believe that most traders using the Tradesports NFL exchange are accustomed to traditional betting markets (*e.g.*, casino sports books and illegal bookmakers). As such, the vast majority of traders are experienced (inexperienced) in buying (selling) bets. Furthermore, selling a bet is less straightforward because one must perform a more complex calculation to determine the possible returns.

In traditional markets, gamblers who expect Team A (B) to cover will buy bets on Team A (B). On the Tradesports exchange, those who expect Team A to cover the spread will buy bets on Team A, but those who expect Team B to cover must *sell* bets on Team A. It is plausible that reframing bets in this manner causes confusion among unsophisticated traders and thus creates a reluctance to sell.⁵ If this is the case, then the resulting supply suppression would force contract win rates to fall relative to expectations. In addition, price inefficiency could sustain while underlying events are active because the high rate of information flow during this time would make across-online-market arbitrage problematic. The supply suppression hypothesis would be supported if we find that prices are most inflated when unsophisticated traders dominate the market.

To better understand the relationship between information flow rates and price biases, we next briefly examine the Tradesports National Basketball Association (NBA) exchange. Scores occur more frequently in NBA basketball games than they do in NFL football games, but each contributes comparatively little information about contract expiry value. By contrasting the price/value deviations in these two markets, we hope to infer whether the magnitude of information shocks is an important determinant of mispricing.

Figure 5 demonstrates that observed NBA contract win rates (the dashed curve) more closely approximate IEM political contract win rates. That is, compared to the NFL exchange, the NBA exchange is characterized by smaller deviations between prices and values, and these deviations are also

PRICE BIASES IN A PREDICTION MARKET




FIGURE 5. Win rates of NFL and NBA contracts. This plot is generated from 216,564 (229,627) NFL (NBA) contract trades occurring after kickoff (tipoff). Each data point represents the observed mean win rate of contracts purchased at each price level (grouped into \$2.50 bands). A win is defined as expiry at \$100.

more symmetric around the market-price-midpoint of \$50. This result suggests that the relatively large magnitude of arriving information shocks may be one source of systematic NFL contract price biases. We next explore this idea further by focusing our attention on price fluctuations after touchdowns, events which represent significant news innovations.

In the NFL, the majority of scoring drives span several minutes from beginning to end. A team typically takes possession of the football relatively far away from, and attempts to move into, scoring position. As the offense advances, the likelihood that the drive will produce a touchdown increases on average. A plot of price movement in the neighborhood of a score should therefore exhibit considerable movement in the minutes just prior to $t = 0$ as the offense drives downfield followed by a flat line thereafter once the score occurs. Even when the opposing team takes possession on the ensuing kickoff, the price should not move on average since bettors know that this change of possession will occur. Thus, a price spike would appear at $t = 0$ only if the event is unexpected.

Figure 6 illustrates the mean win rates for contracts purchased at any point after a touchdown occurs.⁶ If the most recent news event was negative (positive), then it is used to construct plot 6A (6B). We define a negative (positive) touchdown as one that makes an associated contract less (more) likely to expire at \$100. Contrasting the two plots, post-score mean contract win rates appear depend partly upon the direction (negative vs. positive) of the prior news event. The mean *Ticks to Expiry* for all contracts purchased following negative (positive) touchdowns is -6.63 (-3.63), each of which is statistically

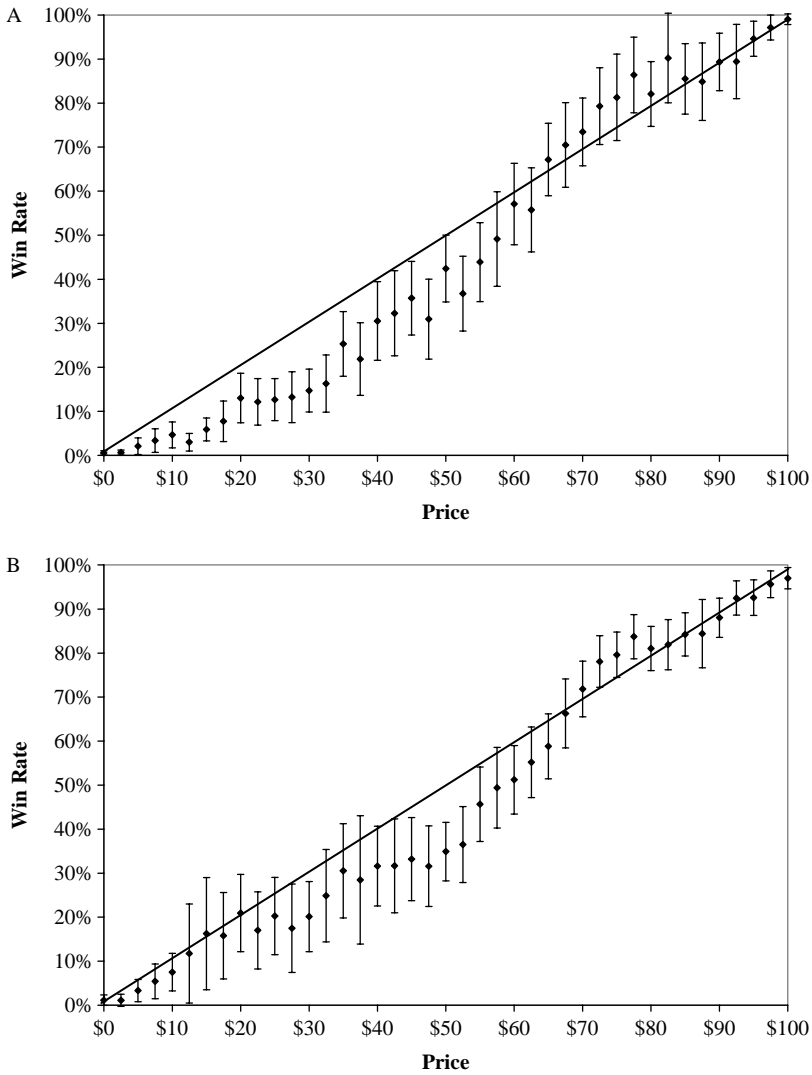


FIGURE 6. (A) Contract prices and win rates after negative news events. This plot is generated from 85,810 NFL contract trades occurring after negative news events. A news event is defined to be negative if it makes the contract less likely to expire at \$100. Each data point represents the observed mean win rate of contracts purchased at each price level (grouped into \$2.50 bands). A win is defined as expiry at \$100. Error bars indicate 95% confidence limits after adjusting standard errors for the effects of clustering. (B) Contract prices and win rates after positive news events. This plot is generated from 100,381 NFL contract trades occurring after positive news events. A news event is defined to be positive if it makes the contract more likely to expire at \$100. Each data point represents the observed mean win rate of contracts purchased at each price level (grouped into \$2.50 bands). A win is defined as expiry at \$100. Error bars indicate 95% confidence limits after adjusting standard errors for the effects of clustering.

different from zero (p -Value = 0.0003 and p -Value = 0.0757, respectively), and the difference in means is also significant (p -Value = 0.0761).

On the other hand, we find that the mean *Ticks to Expiry* for contracts purchased after kickoff but before a touchdown is scored (all transactions that

are not included in 6A or 6B) is -1.67 , which is not statistically different from zero. The mean elapsed time over this period is 28.25 minutes. Findings thus far indicate that this market is inefficient at incorporating information into prices but, as we later demonstrate, there are several important determinants for which we must adjust before drawing conclusions.

To better understand the patterns appearing in Figure 6, we isolate transactions occurring in the time window around scores. Under the null hypothesis of market efficiency, we expect to find that the average post-score price movement of any asset from purchase to expiry is zero. Examining all transactions occurring less than five minutes after touchdowns, we instead observe that the mean *Ticks to Expiry* is -7.10 , which is statistically different from zero (p -Value = 0.0003).

In Figure 7A (7B), we plot the mean observed *Ticks to Expiry* and volume for all trades in each one-minute period within a 30-minute neighborhood of negative (positive) touchdowns.⁷ Contrasting 7A and 7B, evidence shows that the magnitude of the price/value deviation also depends upon the direction of the prior news event. The mean *Ticks to Expiry* for contracts purchased fewer than five minutes after negative (positive) scores is -9.47 (-4.97), and the difference between these two means is statistically significant (p -Value = 0.0040). Each mean is also reliably different from zero (p -Value < 0.0001 and p -Value = 0.0234, respectively). Because transaction timing appears to be an important determinant of subsequent price movements, we next divide the data into within-game subgroups.

To establish whether post-score returns vary from the start to the end of each underlying event, we calculate mean post-score *Ticks to Expiry* by game quarter. Again, only those assets that are traded fewer than five minutes after touchdowns are included. As shown in Table 1, contracts bought immediately after negative first-quarter scores are significantly overpriced, but those bought following positive first-quarter scores are not. The difference in means is highly significant (mean = -15.28 ; p -Value = 0.0003). The difference in means is also statistically significant in the second quarter (mean = -6.53 ; p -Value = 0.0271). In the final two quarters, assets purchased immediately after scores exhibit negative returns to expiry in 16 of 20 of the 24 price/quarter subgroups, but prior score direction is not a significant determinant of subsequent returns. We also find that whenever the execution price is between \$20 and \$60, post-score returns are negative regardless of quarter.

To formally test for these price biases, we next specify several OLS models in which the dependent variable is *Ticks to Expiry*. Initially, we again include only those trades occurring less than five minutes after scores, but later extend the cutoff to incorporate all trades executed less than 20 minutes after scores. Table 2 describes the post-score price movements of the group of contracts used to produce Table 1 after adjusting for both the elapsed time since the previous score occurred and total contract volume.

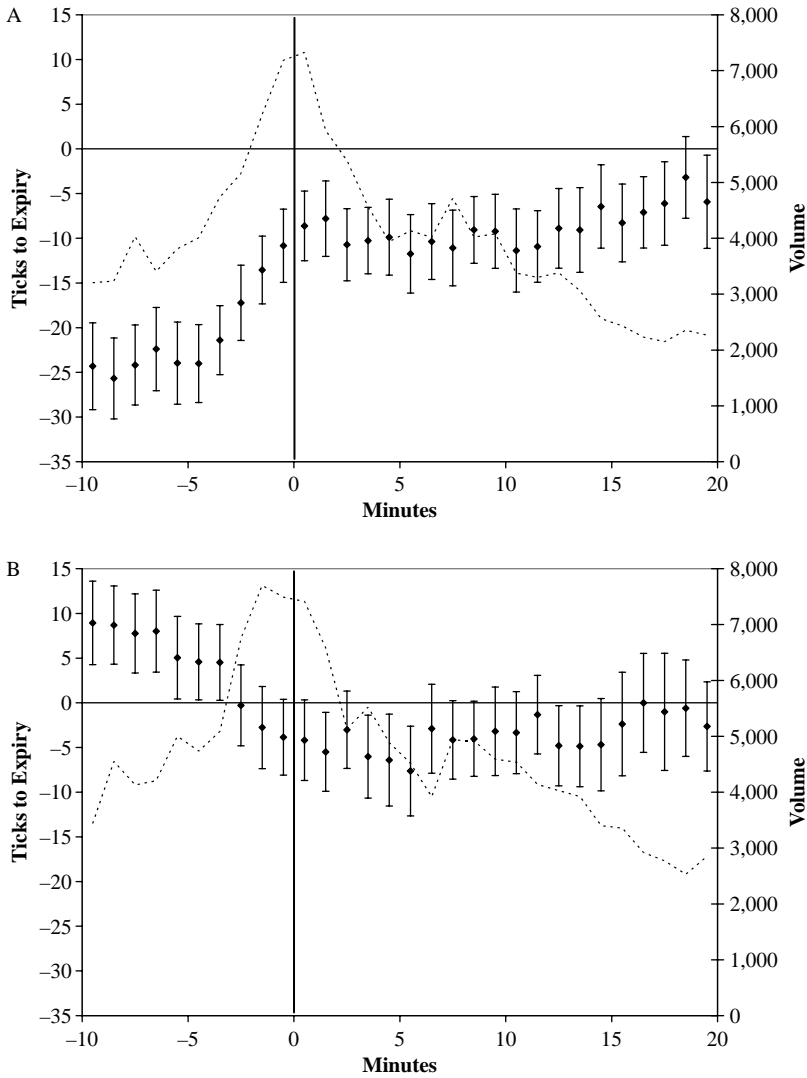


FIGURE 7. (A) NFL contract price movement after negative scores. This plot is generated from 75,398 NFL contract trades occurring within the time neighborhood of negative scores. Negative scores are defined to be those that make a contract less likely to expire at \$100. Each data point represents the observed mean *Ticks to Expiry* of contracts purchased fewer than 10 minutes before and 20 minutes after negative touchdowns (grouped into 60-second bands). *Ticks to Expiry* is calculated by differencing expiry and execution prices. One tick is equal to \$1. Error bars indicate 95% confidence limits after adjusting standard errors for the effects of clustering. The dashed curve indicates volume expressed as the number of 10-contract lots traded within each time band. (B) NFL contract price movement after positive scores. This plot is generated from 89,946 NFL contract trades occurring within the time neighborhood of positive scores. Positive scores are defined to be those that make a contract more likely to expire at \$100. Each data point represents the observed mean *Ticks to Expiry* of contracts purchased fewer than 10 minutes before and 20 minutes after negative touchdowns (grouped into 60-second bands). *Ticks to Expiry* is calculated by differencing expiry and execution prices. One tick is equal to \$1. Error bars indicate 95% confidence limits after adjusting standard errors for the effects of clustering. The dashed curve indicates volume expressed as the number of 10-contract lots traded within each time band.

PRICE BIASES IN A PREDICTION MARKET

TABLE 1
POST-SCORE PRICE MOVEMENTS OF NFL CONTRACTS.

Event Type	Price Range	First Quarter			Second Quarter		
		Volume	Ticks to Expiry	p-Value	Volume	Ticks to Expiry	p-Value
Negative	$\$0 < P \leq \20	239	-3.96	0.6367	1,135	-9.26	0.0000
	$\$20 < P \leq \40	1,832	-19.79	0.0000	2,281	-14.40	0.0011
	$\$40 < P \leq \60	1,926	-12.59	0.0155	2,341	-13.74	0.0143
	$\$60 < P \leq \80	558	-2.80	0.7160	736	-5.95	0.3645
	$\$80 < P < \100	15	-4.55	0.8067	240	6.37	0.1308
	All	4,570	-13.47	0.0002	6,733	-11.16	0.0001
Positive	$\$0 < P \leq \20	0	-	-	61	-14.61	0.0000
	$\$20 < P \leq \40	49	-6.96	0.5995	584	-16.86	0.0106
	$\$40 < P \leq \60	1,219	-6.00	0.3102	1,425	-22.92	0.0000
	$\$60 < P \leq \80	3,344	5.80	0.1473	3,477	-0.24	0.9572
	$\$80 < P < \100	987	-1.96	0.7503	3,126	4.68	0.0462
	All	5,599	1.81	0.5781	8,673	-4.63	0.0722
Event Type	Price Range	Third Quarter			Fourth Quarter		
		Volume	Ticks to Expiry	p-Value	Volume	Ticks to Expiry	p-Value
Negative	$\$0 < P \leq \20	1,418	-4.33	0.0402	5,261	-5.19	0.0000
	$\$20 < P \leq \40	1,760	-9.09	0.1674	2,048	-14.22	0.0024
	$\$40 < P \leq \60	964	-10.21	0.2943	625	-3.31	0.7079
	$\$60 < P \leq \80	319	2.69	0.7499	1,137	-2.45	0.7906
	$\$80 < P < \100	448	2.00	0.6492	1,924	-14.57	0.0336
	All	4,909	-5.74	0.0886	10,995	-7.85	0.0024
Positive	$\$0 < P \leq \20	189	2.97	0.6880	1,649	-2.80	0.2696
	$\$20 < P \leq \40	436	-16.25	0.0128	1,008	-12.53	0.0095
	$\$40 < P \leq \60	1,136	-10.60	0.1601	752	-5.32	0.5819
	$\$60 < P \leq \80	1,603	-6.58	0.2687	2,887	-15.05	0.0587
	$\$80 < P < \100	1,975	2.28	0.3580	3,656	-9.76	0.0141
	All	5,339	-4.75	0.1214	9,952	-9.66	0.0040

This table shows summary statistics describing 36,317 contract trades executed fewer than five minutes after a touchdown occurs. Event type is defined to be negative (positive) if the touchdown makes the contract less (more) likely to expire at \$100. Events are classified into time periods based on when the score occurred. *Volume* is the number of 10-contract lots traded in the five minute post-score period and *Ticks to Expiry* is calculated by differencing expiry and execution prices. One tick is equal to \$1. *p-Value* shows the confidence level for rejecting the null hypothesis that *Ticks to Expiry* is not different from zero.

In this model, *MSLS* is the number of minutes since the last score occurred. *Q1* (*Q2*, *Q3*, *Q4*) is a dummy variable set to 1 when the score occurs in quarter 1 (2, 3, 4), and to 0 otherwise. *Score* is the point value of the observed event [set to -6 (6) if the touchdown makes the contract less (more) likely to expire at \$100]. *Volume* is the demeaned natural log of the number of 10-contract lots traded.⁸ Each *PRICE_X_Y* variable is set to 1 whenever the execution price is greater than \$X and less than or equal to \$Y, and to 0

TABLE 2
SHORT-TERM POST-SCORE MISPRICING OF NFL CONTRACTS.

	Estimate	p-Value	Estimate	p-Value	Estimate	p-Value
Intercept	-3.51	0.2120				
<i>MSLS</i>	-0.65	0.0004	-0.63	0.0009	-0.66	0.0006
<i>Q1*Score</i>	1.17	0.0002	0.87	0.0079	0.85	0.0093
<i>Q2*Score</i>	0.47	0.0706	0.33	0.2000	0.30	0.2552
<i>Q3*Score</i>	-0.08	0.8138	-0.13	0.7003	-0.14	0.6869
<i>Q4*Score</i>	-0.25	0.4350	-0.20	0.5125	-0.17	0.5700
<i>Volume</i>	-0.70	0.0009	-0.68	0.0009	-0.68	0.0010
<i>Price</i>	0.04	0.4297				
<i>Price_0_20</i>			1.04	0.6229		
<i>Price_20_40</i>			-7.04	0.0071		
<i>Price_40_60</i>			-5.81	0.0390		
<i>Price_60_80</i>			1.83	0.4576		
<i>Price_80_100</i>			1.03	0.6915		
<i>Price_0_10</i>					3.07	0.1284
<i>Price_10_20</i>					-1.35	0.5998
<i>Price_20_30</i>					-5.61	0.0564
<i>Price_30_40</i>					-8.14	0.0071
<i>Price_40_50</i>					-5.73	0.0959
<i>Price_50_60</i>					-5.83	0.0769
<i>Price_60_70</i>					-2.32	0.4394
<i>Price_70_80</i>					7.70	0.0067
<i>Price_80_90</i>					3.34	0.2072
<i>Price_90_100</i>					-0.93	0.7624

This OLS model examines 36,317 contract trades executed fewer than five minutes after a touchdown occurs. The dependent variable is *Ticks to Expiry*, which is calculated by differencing the expiry price of each contract and the price at which the transaction occurred. One tick is equal to \$1. *MSLS* is the number of minutes since the last score. *Q1 (Q2, Q3, Q4)* is a dummy variable set to 1 when the score occurs in quarter 1 (2, 3, 4), and to 0 otherwise, *Score* is the point value of the observed event [set to -6 (6) if the touchdown makes the contract less (more) likely to expire at \$100], and *Volume* is the natural log of the demeaned number of 10-contract lots traded throughout the life of the contract. Each *PRICE_X_Y* variable is set to 1 whenever the execution price is greater than \$X and less than or equal to \$Y, and to 0 otherwise. We adjust standard errors for the effects of clustering.

otherwise. We adjust standard errors for the effects of clustering (Fuller, 1975).

Results for *MSLS* and *Q1*Score* show that information shocks play an important role in prediction market price dynamics. The negative estimate for *MSLS* means that those contracts purchased later (earlier) within the five-minute post-score span are more (less) overpriced than otherwise. So, while the long-term ($t = 5$ to $t = 20$) trends in Figures 7A and 7B indicate that post-score prices tend to drift downwards (towards underlying values), regression results show that prices drift upwards (further away from underlying values) in the short-term.

The coefficient estimates for the variable *Q1*Score* indicate that first-quarter score direction is an important determinant of subsequent price movement. Contracts purchased after negative scores drop more in value compared to those purchased following positive scores. The models employing price dummies confirm that contracts purchased fewer than five minutes after news innovations

while price is between \$20 and \$60 subsequently experience significant declines in value.

Importantly, while prior research suggests that any differences between prices and values should be symmetric around the market-price-midpoint of \$50, and that assets priced above \$50 should be underpriced, the significant estimates for *PRICE_40_60* and *PRICE_50_60* suggest otherwise in the short-term. Results in Table 2 also consistently show that when volume is highest, prices are most above underlying values. Our interpretation is that the most public events attract the least sophisticated traders (or more unsophisticated traders), thus supply suppression is particularly pronounced when popular events underlie Tradesports contracts. This finding is consistent with the supply suppression hypothesis offered earlier.⁹ The final portion of our analysis focuses on the relatively long-term inefficiencies persisting after news innovations.

While asset prices deviate furthest from underlying values in the few minutes immediately following information shocks, they also appear to remain significantly biased for a substantial amount of time thereafter. Specifically, the plots in Figure 7 suggest that in the 20-minute span after scores, prices tend to drift downwards (mean *Ticks to Expiry* becomes less negative), but fail to reach underlying values. Contracts bought during this period drop in price on average by \$6.38 prior to expiry, and this decline is reliably different from zero (p -Value = 0.0007). Those that are purchased after negative (positive) scores drop by \$9.22 (\$3.94) on average, each of which is also statistically different from zero (p -Value < 0.0001 and p -Value = 0.0596, respectively), as is the difference in means (p -Value = 0.0002).

Earlier, when examining short-term price movements, it was reasonable to adjust for time by dividing games into quarters. However, in studying longer time frames, it may be more appropriate to account for time linearly. Figure 8, which shows mean *Ticks to Expiry* within each 10-minute band prior to contract expiration, suggests that the amount of time remaining until contract expiry may be an important factor in explaining long-term variations in post-score price movements.¹⁰

OLS regression results presented in Table 3 provide confirmation that information shocks influence the magnitude of price biases in this market. Even after accounting for time to expiry, volume, and price, contracts bought less than 20 minutes after scores lose \$4.58 on average from execution to expiry. One possibility is that during this time, while many market participants are eager to trade (as suggested by the relatively high transaction volume around $t = 0$ in Figure 7), ask prices generally exceed underlying values. We propose that unsophisticated traders in this market are apprehensive about selling and, when underlying events are active, the resulting supply shortage causes asset prices to become persistently inflated. Before underlying events begin, however, sophisticated traders have a sufficient opportunity to exploit any such biases, so the inefficient behavior of novices does not force a sustained divergence between prices and values. The *Volume* coefficient, which demonstrates that popular assets are more overpriced, is consistent with

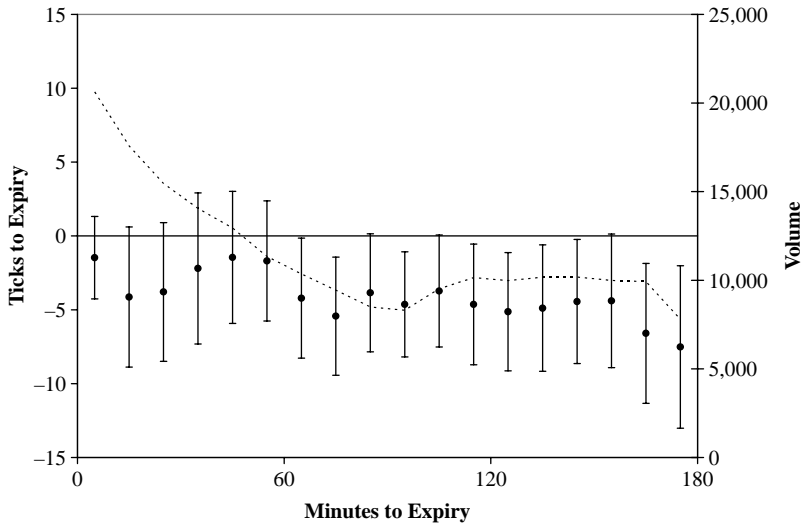


FIGURE 8. NFL contract price movements by minutes to expiry. This plot is generated from 209,052 NFL contract trades occurring after kickoff. Each data point represents the observed mean *Ticks to Expiry* of contracts purchased while fewer than 180 minutes remain until expiry (grouped into 10-minute bands). One tick is equal to \$1. *Ticks to Expiry* is calculated by differencing expiry and execution prices. Error bars indicate 95% confidence limits after adjusting standard errors for the effects of clustering. The dashed curve indicates volume expressed as the number of 10-contract lots traded within each time band.

results from prior specifications and further supports the supply suppression hypothesis.

The estimates for *MSLS* and *Score* are also significant. The former estimate confirms that contract overpricing becomes less pronounced with time throughout the 20-minute post-score period, while the latter indicates that contracts bought following negative scores subsequently lose more value than those bought after positive scores. The positive sign of the *Score* estimate is consistent with the proposition that Tradesports participants dislike selling in general, and provides evidence that they may prefer to hold (rather than sell) contracts that have recently lost value. However, this effect is diminished once we account for non-linearities in the relationship between price and *Ticks to Expiry*. Using dummy variables, we again find that assets priced between \$20 and \$60 lose the most value after scores. After discretely correcting for price, both *MSLS* and *Volume* remain important determinants of post-score price movement.

To further test the relationship between purchase prices and subsequent losses, we also divide contracts into 10 price bands (column 3). Results are consistent with those from the five-price-band model, and the significant coefficient estimate for *Price_50_60* provides further support for two potentially interesting findings. First, imbalances between supply and demand are not symmetric around the market-price-midpoint of \$50. Second, while the utility maximization model predicts that assets priced above \$50 are undervalued, we observe that the opposite is true within a relatively wide price range. These

PRICE BIASES IN A PREDICTION MARKET

TABLE 3
LONG-TERM POST-SCORE MISPRICING OF NFL CONTRACTS.

	Estimate	p-Value	Estimate	p-Value	Estimate	p-Value
Intercept	-4.58	0.0656				
<i>MSLS</i>	0.15	0.0466	0.16	0.0342	0.16	0.0408
<i>MRUE</i>	-0.04	0.0364	-0.02	0.3374	-0.02	0.4214
<i>Score</i>	0.34	0.0241	0.24	0.1087	0.24	0.1252
<i>Volume</i>	-0.49	0.0161	-0.49	0.0144	-0.49	0.0138
<i>Price</i>	0.05	0.2100				
<i>Price_0_20</i>			-0.83	0.6689		
<i>Price_20_40</i>			-8.36	0.0079		
<i>Price_40_60</i>			-9.24	0.0073		
<i>Price_60_80</i>			1.46	0.6722		
<i>Price_80_100</i>			-0.22	0.9382		
<i>Price_0_10</i>					0.53	0.7575
<i>Price_10_20</i>					-2.60	0.3864
<i>Price_20_30</i>					-7.79	0.0141
<i>Price_30_40</i>					-9.10	0.0149
<i>Price_40_50</i>					-9.17	0.0157
<i>Price_50_60</i>					-9.67	0.0138
<i>Price_60_70</i>					-1.89	0.6234
<i>Price_70_80</i>					5.20	0.1469
<i>Price_80_90</i>					1.26	0.7022
<i>Price_90_100</i>					-1.71	0.5161

This OLS model examines 105,737 contract trades executed fewer than 20 minutes after a touchdown occurs. The dependent variable is *Ticks to Expiry*, which is calculated by differencing the expiry price of each contract and the price at which the transaction occurred. One tick is equal to \$1. *MSLS* is the number of minutes since the last score, *MRUE* is the number of minutes remaining until expiry, *Score* is the point value of the observed event [set to -6 (6) if the touchdown makes the contract less (more) likely to expire at \$100], and *Volume* is the natural log of the demeaned number of 10-contract lots traded throughout the life of the contract. Each *PRICE_X_Y* variable is set to 1 whenever the execution price is greater than \$X and less than or equal to \$Y, and to 0 otherwise. We adjust standard errors for the effects of clustering

results, in conjunction with earlier evidence showing that assets are overpriced in aggregate, suggest that this particular exchange, and possibly others like it, may be dominated by unsophisticated traders who unwittingly suppress asset supply. We next explore the possibility that our findings can be used to implement an abnormally profitable within-game short-selling strategy.

C. Trading strategy

One obvious profit-seeking approach would be to exploit the anticipated value decline of assets within the price range \$20 to \$60. For instance, we find that the mean expiry price for all contracts purchased within the price band $\$27.50 \pm \1.25 is \$16.91 (N = 3,110). The return rate of a strategy that sells contracts in this range and buys them back at expiry is 13.43% after accounting for round-trip trading fees.¹¹ Even if the realized rate at which such contracts expire at \$100 were two standard errors above the mean, a trader would expect to earn a 5.12% profit over the course of just a few hours.

Perhaps more interestingly, it would also be possible to exploit the market without conditioning on price, instead taking advantage of the combined effects of aggregate supply suppression and contract misvaluation after news innovations. Our proposed trading strategy is as follows. Following a negative (positive) score, we sell one contract at a price equal to that of the first transaction occurring at least one (two) minute(s) after the score, as long as that trade did not occur longer than two (three) minutes after the score. We choose these particular one-minute periods to be conservative, as our assumed selling prices are less desirable than those available immediately after scores.¹²

We also measure the effects imposing of one of two liquidity constraints. The more conservative *Restriction 1* requires that we sell a contract only if it has traded in each of the 10 one-minute spans prior to the score. *Restriction 2* is less stringent, and requires that we sell only when a contract has traded in at least five of the 10 one-minute spans before the score occurs. All assets sold using either strategy are held to expiry.

Under *Restriction 1*, we sell 72 (85) contracts approximately one (two) minute(s) after negative (positive) events. Mean *Ticks to Expiry* for this portfolio after round-trip transactions costs is -10.23 , with a standard error of 3.99. The mean selling (expiry) price is \$47.97 (\$36.94), and the strategy produces a 19.51% rate of return. As an additional check to determine whether limited market depth would prohibit implementation, we measure the frequency at which our contracts trade during the one-minute span in which we sell. We find that market participants execute 1,086 trades of the 157 contracts that we target for our portfolio, and the value of these trades totals \$84,092. Thus, market depth does not appear to prohibit the strategy when employing *Restriction 1*.

Using *Restriction 2*, we sell 516 (600) contracts one (two) minute(s) after each negative (positive) event, and the mean (standard error) *Ticks to Expiry* for this portfolio is -6.42 (2.23). The mean selling (expiry) price is \$53.73 (\$46.51), and the rate of return is 13.76%. During the one-minute span in which we sell, the targeted contracts trade 4,719 times and the combined value of these trades is \$383,217.¹³ Again, market depth would not severely limit the implementation of our strategy.

While it would be useful to quantify the effect of the bid-ask spread on profitability, the required data are unavailable. However, because our strategy restricts trading to those times in which the market is particularly liquid, this issue is unlikely to invalidate our conclusions. Furthermore, because contract prices experience larger declines after negative scores and also within lower price bands, a trader could earn a higher rate of return were she to impose a restriction to sell only if these two conditions were met.

D. Future research

The sports exchange examined here typically lists only one asset per *event*, so participants who believe that an asset is overpriced must sell to capitalize.¹⁴ If there is a shortage of willing sellers, then assets prices can persistently exceed underlying values while events are active. However, an

alternative asset structure could mitigate this potential problem. When there is one asset per *state*, within-exchange arbitrage is possible. For instance, in election markets, an agent may buy a contract with payout dependent upon whether candidate A wins a political seat over candidates B, C, and D. If she believes that contract A is overpriced, she can sell A and buy B, C, and D. Thus, if the supply of A is suppressed, then a within-exchange arbitrage operation can correct any shortfalls in its supply.

This idea may be meaningful when designing contracts to be listed on internal corporate prediction markets. In that setting, the ultimate goal is to obtain precise forecasts of business events. So, it may be preferable to construct assets in a way that facilitates within-exchange arbitrage. Tradesports does offer a set of football contracts for which arbitrage may be possible. For example, the exchange lists contracts for each of the 32 NFL franchises to win the Super Bowl. One potentially fruitful line of research may be to compare the predictive accuracy of event-based contracts to that of state-based assets in the Tradesports market.

IV. CONCLUSION

Financial economists have long studied information processing in traditional casino-type sports betting markets. Here we analyze a relatively new form of sports betting, the within-game financial-market-style of betting conducted on Tradesports.com. In several respects, this setting is well-suited as an alternative domain in which to examine how markets react to new information and to study the effects of contract framing on trader behavior.

We begin with the observation that the prices of NFL contracts are not systematically biased prior to kickoff, a time period in which little new information is flowing and in which arbitrageurs can easily trade across markets. After kickoff, major news events occur (in the form of within-game scores), and traders update their prior estimates of contract value accordingly. However, we find evidence suggesting that unsophisticated NFL bettors avoid selling, and thus force asset prices to rise above underlying values on average. This deviation between contract prices and values is most pronounced immediately after information shocks arrive, especially when the incoming news is negative. While underlying events are active, arbitrage operations are problematic, contracts remain persistently overpriced, and one could profitably exploit the bias by selling contracts shortly after scores.

ACKNOWLEDGEMENTS

I thank David Pennock, Eric Zitzewitz, and an anonymous referee for their insight and meaningful suggestions throughout, Rob Dougherty and Brijesh Patel for assistance with the event data, and the McCoy Endowment for financial support on this project. Any errors are strictly my own.

NOTES

1. A 'win' is defined to be expiry at \$100.
2. In a sides (totals) wager, gamblers bet on the difference in (sum of) the number of points that opposing teams will score. An example of a Tradesports NFL sides contract listing is 'NFL.WAS@SFN.WAS-2.5.' The exchange usually lists multiple sides contracts and a single totals contract for each game. In our analysis, we include (exclude) all sides (totals) contracts.
3. We adjust throughout at the contract level. Adjusting instead at the game level does not materially alter results.
4. For instance, the injury status of key players is sometimes revealed several days before kickoff.
5. Prior research suggests that bettors who wager on football are relatively unsophisticated compared with those who wager on other sports (Woodland and Woodland, 1999).
6. We omit field goals from the analysis because they are difficult to interpret; they partly represent bad news because the realization is that a touchdown will not be scored. We also omit extra points from the analysis, thus all scores have an absolute value of six.
7. We do not look past 20 minutes for several reasons. First, we are concerned primarily with mispricing shortly after scores events because we later propose a trading strategy that sells contracts at $t = 1$ or $t = 2$ minutes after such events. Second, as t increases beyond 0, a greater proportion of contracts are affected by subsequent scoring, so *Ticks to Expiry* becomes a less reliable measure of mispricing caused by an earlier event. It is also problematic to examine contract price movements only for those events that are not followed by another score because such circumstances occur almost exclusively near contract expiry. And, as we later demonstrate, the amount of time remaining until expiry is potentially an important determinant of *Ticks to Expiry*.
8. We calculate the total number of units traded throughout the lifetime of each contract, subtract the mean number of units traded per contract across the entire sample, and then take the natural log.
9. On the other hand, Tetlock (2004) finds no difference in returns between featured sports events (those receiving the most television network coverage) and non-featured sports events nor between high- and low-volume contracts. However, the data in that study are gathered by taking price snapshots at 30-minute intervals, thus nearly all observations come from the inactive period.
10. We exclude from this plot all trades made while more than 180 minutes remain to expiry. Games that go into overtime typically last longer than three hours and are more likely to result in an underdog beating the spread than otherwise. Because all events are framed with respect to favorites, contracts with more than 180 minutes remaining to expiry expire at \$0 more often than expected *ex-ante*.
11. $\$27.50 - (\$16.91 + \$0.40) - \$0.40 / (\$72.50 + \$0.40)$
12. Refer to Figures 7A and 7B.
13. The potential dollar value of an implemented strategy would be substantial because the figures above account only for executed trades, and bid volume at slightly less favorable prices is typically large. In addition, each strategy would be effective for longer than just the one-minute spans examined above.
14. An 'event' in this context is defined as a particular sides line.

REFERENCES

- Joyce Berg, Forrest Nelson and Thomas Rietz Results from a Dozen Years of Election Futures Markets Research, Technical Report, University of Iowa (2000).
- Joyce Berg, Forrest Nelson and Thomas Rietz Accuracy and Forecast Standard Error of Prediction Markets, Working Paper, University of Iowa (2003).
- Kay-Yut Chen and Charles Plott Information Aggregation Mechanisms: Concept, Design, and Field Implementation for a Sales Forecasting Problem, Social Science Working Paper No. 1131, California Institute of Technology (2002).
- Wayne Fuller 'Regression analysis for sample survey' *Sankhya* (1975) 37 117–132.
- Steven Levitt 'How do markets function? An empirical analysis of gambling on the national football league' *Economic Journal* (2004) 114 2043–2066.
- Gerhard Ortner Forecasting Markets – An Industrial Application, Working Paper, Technical University of Vienna (1998).
- Pennock, David, Steve Lawrence, C. Lee Giles, and Finn Nielsen The Power of Play: Efficiency and Forecast Accuracy in Web Market Games, NEC Research Institute Technical Report 2000–168 (2001).
- Emile Servan-Schreiber, Justin Wolfers, David Pennock and Brian Galebach 'Prediction markets: does money matter?' *Electronic Markets* (2004) 14 243–251.

PRICE BIASES IN A PREDICTION MARKET

- Paul Tetlock How Efficient Are Information Markets? Evidence from an Online Exchange, Working Paper, University of Texas (2004).
- Justin Wolfers and Eric Zitzewitz Five Open Questions about Prediction Markets, Working Paper, Stanford University and University of Pennsylvania (2005a).
- Wolfers, Justin and Eric Zitzewitz Interpreting Prediction Market Prices as Probabilities, Working Paper, Stanford University and University of Pennsylvania (2005b).
- Justin Wolfers and Eric Zitzewitz Prediction Markets in Theory and Practice, Working Paper, Stanford University and University of Pennsylvania (2005c).
- Bill Woodland and Linda Woodland 'Expected Utility, skewness, and the baseball betting market' *Applied Economics* (1999) 31 337–345.
- Ralph Woodruff 'A simple method for approximating the variance of a complicated estimate' *Journal of the American Statistical Association* (1971) 66 411–414.