How Efficient Are Information Markets? Evidence from an Online Exchange[*]

Paul C. Tetlock[†]

January 2004

Abstract

Do observed inefficiencies in betting markets generalize to real financial markets? Evidence from an online exchange with characteristics of both types of markets suggests they do not. This study replicates some known anomalies in sports wagering markets on the online exchange, such as overreaction to news. The observed inefficiencies do not generalize to other wagering markets, such as financial markets, even though both markets have similar liquidity, volume and contractual structure. These results suggest researchers should proceed with caution before drawing analogies between sports wagering markets and real financial markets. They also show that financial markets can be efficient despite numerous obstacles to arbitrage.

1

1 Introduction

The online exchange studied here is similar in many respects to wagering markets. Somewhat surprisingly, many papers studying sports betting markets find significant market inefficiencies. Several studies (*e.g.*, Ziemba and Hausch, 1986, and Jullien and Salanie, 2000) using horsetrack data have documented the "favorite-longshot bias": expected returns to betting on a horse increase monotonically with the probability of the horse winning. Woodland and Woodland (1994) find a reverse favorite-longshot bias in the baseball point-spread market: expected returns to betting on a baseball team *decrease* monotonically with the probability of the team winning. Woodland and Woodland (2001) replicate this finding in the National Hockey League betting market. And Avery and Chevalier (1999) discover that point spreads in the professional football betting market overreact to information, allowing a marginally profitable contrarian trading strategy.

Do observed inefficiencies in betting markets generalize to real financial markets? Evidence from an online exchange with characteristics of both types of markets suggests they do not. This result contradicts the commonly held belief that inefficiencies in betting markets are analogous to inefficiencies in real financial markets. In fact, some authors of studies on betting markets have argued that "wagering markets have a better chance of being efficient [than securities markets] because the conditions (quick, repeated feedback) are those which usually facilitate learning" (Thaler and Ziemba, 1988).

Studies of betting markets attempt to address three limitations that cloud the interpretation of field studies based on traditional financial markets. First, all field studies must test market efficiency jointly with a model of equilibrium (Fama, 1970; Fama, 1991). Models of asset market equilibrium must specify the appropriate price of risk in the economy and the extent to which different assets are exposed to these risks (Roll, 1977). Wagering markets overcome this difficulty, because most assets traded in these markets are contingent claims with no non-diversifiable (systematic) risk.

Second, the field environment must closely approximate the conditions of the theoretical model being tested. Many models of rational financial markets assume the forces of arbitrage operate unchecked. However, as emphasized by DeLong *et al.* (1990), mispriced assets in real financial markets can become even more mispriced in the short-run and can remain mispriced for years. These factors limit arbitrageurs' ability to correct mispricing, because arbitrageurs often have short time-horizons and limited access to capital. In wagering markets, arbitrageurs know that asset prices will soon reach fundamental value, allowing them to take larger positions and correct mispricing.

Third, studies using new data from wagering markets avoid the problem of data mining that besets studies based on data from real financial markets. It is necessary to test financial theory on data sets other than those used to develop the theory. Because studies of real financial markets perform myriad tests of efficiency using the same data, they are likely to find spurious market inefficiencies.

For these three reasons, authors of studies on wagering markets argue these markets should be quite efficient relative to real financial markets. When previous studies identify inefficiencies in wagering markets, the authors implicitly or explicitly claim that these inefficiencies generalize to real financial markets. Proponents of efficient markets are skeptical of this argument, because betting markets differ from real financial markets in numerous ways that could make them *less* efficient.

First, different sets of contracts are traded on sports wagering markets and real financial markets. There are potentially rational, non-monetary reasons to bet on sporting events. Fans could receive enjoyment from betting on their favorite team or from the excitement of wagering. These non-monetary motives are unlikely to play a large role in the trading of financial event futures contracts. Because most models of market efficiency assume traders act to maximize their expected utility from wealth, it is important to test these models in environments where this assumption is satisfied.

Second, few traders in sports betting markets have experience trading in other realms, such as Wall Street, suggesting trader inexperience could lead to inefficient outcomes. For example, most bettors at horsetracks have no professional trading experience in financial markets. If these bettors are less likely to identify profitable arbitrage opportunities, these opportunities could persist indefinitely on horsetrack markets.

Third, although sports betting markets eliminate some impediments to arbitrage, they introduce new deterrents to arbitrage. Sports betting markets operate on a smaller scale, supply less liquidity, use less transparent and robust trading mechanisms, and have higher transaction and participation costs than most real financial markets. All of these features imply arbitrage will be less appealing and less effective in sports betting markets.

To explore the effect of making betting markets more like real financial markets, I examine an online exchange that differs from most betting markets in four important ways that make it more like real financial exchanges.[1] First, some contracts on the online exchange are based on financial events. Second, many traders on the online exchange have experience

---

[1]Nevertheless, like other betting markets, the online exchange operates on a smaller scale than real financial markets. I will demonstrate later that differences in stakes appear to have very little effect on pricing and trading activity on the online exchange.

trading in other realms, such as Wall Street. Third, contracts on the online exchange are traded via a continuous double auction similar to the mechanisms used on the world's major stock, currency, commodity and derivatives exchanges.[2] Fourth, the online exchange charges lower commissions (fees) on trades and has lower entry and exit costs than typical betting markets.[3]

The results from statistical tests in this study replicate several inefficiencies discovered in earlier studies of sports betting markets, such as overreaction to past performance and the reverse favorite-longshot bias. However, these inefficiencies do not generalize to the financial markets on the online exchange, even though the structure, liquidity and volume of the financial and sports contracts on the exchange are quite similar.[4] The relative efficiency of financial markets as compared to sports markets on the exchange is also surprising, because both markets serve a primarily speculative purpose. Overall, these results show financial markets can be quite efficient despite numerous obstacles to arbitrage.

Finally, it is important to mention that there are larger and more liquid sports and financial exchanges (in Las Vegas and New York, respectively) with nearly identical contracts to those traded on the online exchange examined here. If cross-exchange arbitrage is effective, then the pricing patterns on the online exchange should be similar to pricing patterns in other sports and financial exchanges. Thus, the conclusion that financial markets are efficient relative to sports markets could apply to other exchanges as well.

The next section describes the data set used in this paper. In this section, I also examine the features of the online exchange in greater depth, offering some descriptive statistics. The third section sketches a theoretical framework for modeling the contracts traded on the exchange. The fourth section reports nonparametric and parametric tests of market efficiency. In the final section of the paper, I interpret the results and draw conclusions.

---

[2]Gode and Sunder (1993), Friedman and Ostroy (1995), Cason and Friedman (1996), and Noussair *et al.* (1998) experimentally demonstrate that double auctions are particularly robust mechanisms that promote rapid adjustment towards market equilibrium even in the presence of market frictions and trader irrationality. Gjerstad and Dickhaut (1998) and Satterthwaite and Williams (2002) prove theoretically that double auctions implement efficient allocations even in settings with boundedly rational and finite numbers of traders. By contrast, most betting markets operate under a single market maker who determines contract odds through a parimutuel mechanism. A few betting markets use a fixed odds system in which prices do not respond conventionally to supply and demand.

[3]Horsetrack markets and many other betting markets studied in the past have significant commission costs of 10 to 20 per cent of the amount wagered. Entering these markets can be difficult for prospective traders, because the markets are often held in geographically remote locations with no electronic means for conducting trades.

[4]There is also anecdotal evidence that suggests many traders participate in multiple segments of the online exchange.

4

2 Description of the Exchange Data

The data consist of snapshots of the TradeSports.com web site taken at approximately 30-minute intervals over a six-month period. I developed a program to extract periodically all publicly available information about the contracts listed on the web site.[5] The program then transforms this information into a large contract database, which forms the basis of the statistical analyses performed in this paper. The database includes basic statistics for each contract such as transaction prices, the volume of contract transactions, the quantities and prices of offers to purchase or sell the contract, and the last price at which the contract trades (the expiry price). To reduce the size of the unwieldy database resulting from the extraction program, I select only observations with the greatest relevance for financial theory. The smaller database includes observations on contracts that have recently been traded and contracts that are based on events that are about to occur.

The online exchange market is conducted via a publicly available web site, TradeSports.com. The company TradeSports Exchange Limited hosts this web site and is registered in Ireland. The TradeSports exchange facilitates the trading of event futures contracts by its members; that is, the exchange *does not* conduct transactions for its own account. The owner of a contract receives a pre-specified amount ($10) if and only if a pre-specified, verifiable state of nature occurs. I will refer to the event outcome that determines the state of nature as the underlying event. Expiration is the first time at which the state of nature becomes verifiable.

TradeSports limits the risk that a counterparty to a contract will default by imposing stringent margin requirements for each sale or purchase of a contract by one of its members. In most cases, members must retain sufficient funds in their TradeSports account to guard against the maximum possible loss on a transaction. TradeSports also settles and clears all transactions conducted on its exchange.

It is extremely easy for anyone with access to the Internet to open an account with TradeSports and begin trading almost immediately. I have verified TradeSports' claim that opening an account takes less than 10 minutes. Some traders with accounts on the exchange have trading experience in other realms. Chat forums and informal web site polls hosted on TradeSports.com indicate that many traders are based in New York and Chicago and trade more conventional financial securities in their day jobs.[6] The available evidence suggests these professional traders account for a great deal of volume and market-making activity on TradeSports.com.

---

[5] The TradeSports exchange has granted me permission to run this program.

[6] I have also acquired anecdotal evidence that confirms this conjecture.

The trading interface developed by TradeSports is similar to electronic trading interfaces on financial exchanges. Both members and non-members can view a quoteboard listing all contracts in a user-specified category (*e.g.*, politics) by visiting the web site and clicking through two links. By clicking on a specific contract listing, a user can see the entire order book for that contract. TradeSports lists all buy and sell orders anonymously and aggregates these orders by price.[7] The exchange processes order submissions with a latency of about one second.

For ease of interpretation, the exchange divides most contract prices into 100 ticks, worth $0.10 per tick. Some heavily traded contracts are divided into 1,000 ticks, worth $0.01 per tick. TradeSports charges very low commissions for transactions conducted on its exchange. Commissions equal to 0.4% of the maximum contract price ($10) are levied on a per contract basis whenever a contract is bought or sold. At the time of contract expiration, all outstanding contract positions must be liquidated and incur commissions. Note that the $0.08 round-trip transaction fee is smaller than the value of one tick ($0.10) for most contracts. This implies arbitrageurs have an incentive to push prices back towards fundamental values if they stray by even one tick.

More generally, arbitrageurs face few limitations on their trading behavior. Traders can place unlimited buy or sell orders on any asset at any time, regardless of their current contract position or past pricing patterns. By contrast, many regulated financial exchanges prevent traders from selling certain assets that they do not own or only permit short-sales that follow an "uptick" (a price increase). Although TradeSports does not specifically forbid insider trading, the exchange would presumably report such conduct to Irish law enforcement agencies. The exchange rules do proscribe trades intended to "manipulate prices." Finally, the exchange accepts order submissions and operates for 23 hours out of every day, temporarily closing between 3 a.m. EST and 4 a.m. EST for maintenance.

All contracts listed on TradeSports.com are designed by the company that runs the exchange, which encourages input on the content and design of its contracts from TradeSports members. Although listed contracts can depend on the outcomes of financial, economic, weather, entertainment, political, legal, and other current events, the most common contracts are based on United States and international sports events. TradeSports.com seems to be unique in the wide variety of contracts offered. Some of the more intriguing contracts from the fall of 2003 are described in Table I.

---

[7]When two users submit orders to buy a contract at the same price, the first order is filled before filling any part of the second order. This is consistent with the rules for most continuous double auctions.

Table I
Examples of Listed Contracts on the Exchange

Contract Description	Price	Volume	Expiry
Dean to be Democratic presidential nominee	35.5	39,609	7/26/04
Palestinian State established before 2006	20.0	377	Unknown
Homeland Security alert system at "Orange"	16.5	500	12/31/03
Dow Jones to close higher on 10/17/03	6.5	1,605	10/17/03
Buccaneers to win SuperBowl XXXVIII	15.9	86,237	Unknown

Source: TradeSports.com as of approximately 14:00 EST October 17, 2003.

The rules describing how the outcome of an event will be verified are extremely thorough. In general, TradeSports requires written statements from at least two major news sources to verify event outcomes. For contracts that depend upon complex events such as Saddam Hussein's capture or neutralization by a specific date, the exchange issues periodic updates and clarifications to the event rules whenever outcome verification is ambiguous. These difficult-to-verify contracts sometimes hinge upon related events, such as "the Pentagon issues a statement declaring Saddam Hussein is captured or neutralized," that are easier to verify. Remarkably, the verification process has worked flawlessly and without controversy in the past.

Table II
Types of Actively Traded Contracts on TradeSports.com

Category	Contracts	Frequency
Sports-related	11,919	83.8%
Financial	1,967	13.8%
Entertainment	175	1.2%
Current Events	148	1.0%
Miscellaneous	16	0.1%
Total	14,225	100.0%

Source: Author's database based on TradeSports.com web sites from March 18, 2003 to October 4, 2003. To be included in this table, a contract must have been traded.

Table II documents the number and relative frequency of different types of contracts actively traded during the sample period studied. The table shows that sports contracts constitute the vast majority of contracts listed on the exchange, but financial contracts are also significantly represented. The duration of contracts is typically quite short as compared to futures contracts traded on other exchanges. Table III reveals that most of the contracts contained in the data sample expired within one day of the time when they were first listed

on the exchange. Still, the sample does contain over eight hundred contracts with durations that exceed two months.

Table III
Duration of Listed Contracts

Duration	Contracts	Frequency	Cumulative Frequency
One day or shorter	8,311	58.4%	58.4%
From one day to one week	4,340	30.7%	89.1%
From one week to two months	739	5.2%	94.3%
From two to five months	490	3.4%	97.8%
Five months or longer	315	2.2%	100.0%
Total	14,225	100.0%	

Source: Author's database based on TradeSports.com web sites from March 18, 2003 to October 4, 2003. To be included in this table, a contract must have been traded.

TradeSports proudly proclaims it has over 20,000 members who have wagered over $380 million through the exchange. This implies an average of $19,000 wagered per member.[8] This amount represents the sum of all capital invested in and withdrawn from contracts on the exchange. Although these stakes are not nearly as large as those on most financial markets, they are orders of magnitude larger than the stakes in economics experiments.

It is difficult to compare the magnitudes on TradeSports to the amounts traded on real financial markets, because participants in real financial markets who own derivative contracts and sell stocks short face unbounded risks. The market for *buying* stock is the most appropriate (albeit imperfect) point of comparison for markets on TradeSports. On the NYSE, roughly 100 million investors have approximately $10 trillion worth of capital at risk in the market for owning stock. Turnover on the exchange is about one stock per year, so the "average" investor on the NYSE trades approximately $100,000 worth of stock. Although investors in real financial markets undoubtedly face larger incentives than TradeSports' members, the stakes on TradeSports are large enough to merit the attention of Wall Street traders.

The scale of the TradeSports exchange is quite large relative to previously studied experimental markets, including the Iowa political futures markets. A typical wager in a *single transaction* on TradeSports is $100; many transactions exceed $1,000.[9] By contrast, participants in the Iowa political futures markets have an average of only $50 of their own money

---

[8]If the distribution of amounts wagered per member is positively skewed (which is likely), the typical TradeSports member has wagered less than $19,000.

[9]The $100 estimate is based on the default order size of 20 contracts on TradeSports times the average amount wagered on each contract. This latter quantity is identically equal to $5 because buyer and seller wagers necessarily add up to the total amount at risk on each contract, which is $10.

at risk *in the entire market*.[10] Subjects in other economics experiments rarely face risks of more than tens of dollars.

In summary, many features of the TradeSports exchange recommend its use as a "natural laboratory" for testing asset pricing theory. It is easy to determine the fundamental value of the contracts on the exchange that are only exposed to idiosyncratic risk. Because events on TradeSports tend to be widely publicized, the assumption that all traders can costlessly acquire information about these fundamental values is probably not a bad approximation.

Moreover, if prices ever stray from fundamental values, experienced traders from real financial markets have excellent opportunities to identify and eliminate arbitrage opportunities on the TradeSports exchange. Prospective arbitrageurs face minimal hurdles to entering the market, near-zero transaction costs, and very few regulations on their trading behavior. If at least some of these arbitrageurs have rational expectations, purely monetary motivations and access to sufficient capital, they will prevent substantial mispricing on the TradeSports exchange. This concept motivates the model in the next section.

3   A Theory of Arbitrage on TradeSports

This section describes a theoretical framework for understanding observed pricing patterns on the online exchange. The model is based on arbitrage pricing theory (APT) from Ross (1976). The intuition behind APT is straightforward: because investors can hold diversified portfolios of contracts, they will require no compensation for bearing idiosyncratic risk. Because many contracts on TradeSports are not exposed to systematic risk, this idea leads to specific pricing predictions. Moreover, these predictions do not require any economic assumptions about trader preferences.

3.1   Setup and Notation

I consider an environment in which traders can invest in N assets (contracts) indexed by $i = 1, ..., N$. For notational simplicity, I assume the exchange allows trading in all contracts for exactly T periods: $t = 1, ..., T$. At time $t = T$, each contract i expires and pays an uncertain dividend $\widetilde{x}_{iT}$, which equals 0 ticks or 100 ticks (*i.e.*, \$0 or \$10). Denote the set of all possible contract price and dividend realizations at time t by S_t, where $s_t \in S_t$. Let S_{iT} be the set of states at time T in which contract i yields a dividend of 100 ticks. In each period

---

[10]This \$50 estimate is based on data from Berg et al. (1997).

prior to expiry, information about the state is revealed publicly. The conditional probability at time t that asset i's dividend equals 100 ticks is $\Pr(s_T \in S_{iT}|s_t)$. The notation for the price of asset i at time t is p_{it}.

3.2   Contract Prices with Perfect Arbitrage

I make the crucial assumption that there are *no* arbitrage opportunities available on the exchange. Ross (1978) proves that this assumption alone implies asset prices can be represented by the expected product of a strictly positive stochastic discount factor and asset payouts. Applying Ross's theorem to the current situation, at each time t, there exists at least one stochastic discount factor m_t that is strictly positive in all states s_t with $\Pr_{t-1}(s_t \in S_{it}|s_{t-1}) > 0$. This stochastic discount factor satisfies the following relationship:

$$\forall i = 1, ..., N, p_{i,t-1} = E_{t-1}(m_t p_{it}) \tag{1}$$

The preceding equation prices any contract i given the contract's prices one period hence. Imposing the no-arbitrage requirement that $p_{iT} = x_{iT}$ and using backward induction, one obtains asset pricing equations for each contract i that relate current prices to prices at all future dates and to dividends at expiry:

$$\forall t_0 = 1, ..., T - t, p_{it} = E_t(m_{t+1}m_{t+2}...m_{t+t_0}p_{i,t+t_0}) \tag{2}$$
$$p_{it} = E_t(m_{t+1}m_{t+2}...m_T x_{iT}) \tag{3}$$

For convenience, define another stochastic discount factor $M_{t,t+t_0} = m_{t+1}m_{t+2}...m_{t+t_0}$ with $M_{t,t} = 1$ In all asset pricing models, $M_{t,t+t_0}$ represents investors' intertemporal marginal rates of substitution, which are complicated functions of many economic variables. I assume that $M_{t,t+t_0} = f(c_{t+t_0}, y_{t+t_0}, t_0)$, where c_{t+t_0} denotes investors' consumption at $t+t_0$ and y_{t+t_0} denotes a vector of state variables that affect investor utility at $t + t_0$. I assume further that $f(\cdot)$ is continuous in all its arguments. Nearly all asset pricing models meet these intuitive requirements.

Under these assumptions, it is straightforward to show that $M_{t,t+t_0}$ converges toward 1 as t_0 approaches 0. Because most contracts on TradeSports have extremely short durations, the stochastic discount factor will be very close to 1. This approximation simplifies equations

(2) and (3) for each contract i:

$$\forall t_0 = 1, ..., T - t, p_{it} = E_t(p_{i,t+t_0}) \tag{4}$$

$$p_{it} = E_t(x_{iT}) \tag{5}$$

The first equation states that all contract prices on TradeSports follow a random walk without drift. The second equation ensures that contract prices equal fundamental values at all times.

3.3 Incorporating Market Frictions

Even on TradeSports, arbitrageurs incur some costs when they trade. Most arbitrage strategies involve purchasing and selling contracts rapidly, suggesting that arbitrageurs must pay the current asking price to buy a contract but receive only the current bid price when they sell a contract. In addition, TradeSports charges a non-trivial commission on each one-way transaction. In this section, I relax the assumption of perfect arbitrage to model these market frictions.

The presence of bid-ask spreads and commission costs implies the effective price of a contract is different for contract purchases and sales. In other words, equations (4) and (5) become two inequality constraints that place bounds on the maximum and minimum price of a contract. Let a_{it} and b_{it} be the ask price and bid price for contract i at time t. Denote the commission cost per contract by c. Define the observed price to be the average of the current bid and ask prices, so that $p_{it} = (a_{it} + b_{it})/2$. An arbitrageur wanting to buy a contract at time t and liquidate the contract at time $t + t_0$ must not be able to earn abnormal profits, which places a lower bound on the price of contract i at time t:

$$\forall t_0 = 1, ..., T - t, p_{it} + \frac{a_{it} - b_{it}}{2} + c \geq E_t(p_{i,t+t_0} - \frac{a_{i,t+t_0} - b_{i,t+t_0}}{2} - c) \tag{6}$$

Similarly, no arbitrageur must be able to earn excess profits from selling a contract at time t and liquidating the contract at time $t + t_0$:

$$\forall t_0 = 1, ..., T - t, p_{it} - \frac{a_{it} - b_{it}}{2} - c \leq E_t(p_{i,t+t_0} + \frac{a_{i,t+t_0} - b_{i,t+t_0}}{2} + c) \tag{7}$$

The situation is somewhat different for an arbitrageur using a trading strategy that holds contracts until they expire, because this trader incurs no bid-ask spread cost at expiry. TradeSports still charges traders normal commissions if they hold open positions when contracts

11

expire. Ruling out profitable trading strategies that buy or sell contracts and liquidate them at expiry generates the restrictions:

$$p_{it} + \frac{a_{it} - b_{it}}{2} + c \geq E_t(x_{iT} - c) \tag{8}$$

$$p_{it} - \frac{a_{it} - b_{it}}{2} - c \leq E_t(x_{iT} + c). \tag{9}$$

A careful examination of equations (8) and (9) reveals that they imply that equations (6) and (7) hold. Although equations (6) and (7) are theoretically redundant, it is easier to form accurate empirical measures of the variables in these equations. Accordingly, I will use equations (6) and (7) in the next section to test the model.

The two restrictions from equations (8) and (9) place tight upper and lower bounds on contract prices at all times, but they do not predict the pricing fluctuations that occur within the bounds. Financial theory offers little guidance here, so it becomes necessary to rely on economic insights. I assume that price fluctuates randomly within the arbitrage bounds. The source of uncertainty in the model arises from demands of traders with "complex" motives. These traders can be thought of as liquidity traders with potentially, but not necessarily, rational reasons for buying and selling the contract. For example, fans of a sports team could wish to hedge against their team losing. Or, they might want to splurge on a celebration party if their team wins.

Explicit modeling of these complex trading motives is outside the scope of this paper, but the model summarizes these considerations by allowing contract prices to fluctuate randomly within the arbitrage bounds according to the whims of complex traders. To express this intuition, define a random variable u_{it} that satisfies the relationship:

$$u_{it} \equiv p_{it} - E_t(x_{iT}) \tag{10}$$

I assume u_{it} follows a first-order auto-regressive process that satisfies the arbitrage restrictions in equations (8) and (9). Thus, $u_{it} = \rho u_{i,t-1} + v_{it}$, where $0 \leq \rho < 1$. The v_{it} are distributed independently and identically over time with mean 0 and variance $\sigma_{v_i}^2$. All realizations of v_{it} conform to the restrictions (8) and (9).[11] In the empirical tests in the next section, I allow price innovations to be correlated across contracts. As explained in detail

---

[11] The set of assumptions on v_{it} imply that price innovations have a truncated distribution with a potentially strange shape when price nears the arbitrage bounds. The more realistic assumption that price innovations follow some heteroscedastic process complicates the calculations below without changing the qualitative conclusions of the analysis.

later, the TradeSports exchange lists its contracts in such a way that I am able to identify unrelated contracts and explicitly adjust for interdependencies in contract prices.

As a counterpart to the mispricing variable u_{it} defined above, I define another random variable $\epsilon_{it} \equiv p_{it} - E_{t-1}(p_{it})$ that represents the unexpected change in the fundamental value of contract i. By construction, this variable has a mean of 0 and is independent across contracts and time. I also assume the change in contract i's fundamentals is identically distributed with variance σ_ϵ^2.

These assumptions lead to a statistical characterization of returns on TradeSports contracts. First, it is interesting to look at the (unconditional) expected squared pricing error for different contracts. From equation (10), one can interpret $E\left(u_{it}^2\right)$ as a measure of the degree of mispricing. Straightforward calculations show that the expected squared pricing error is proportional to the variance in mispricing innovations $\sigma_{v_i}^2$. Thus, the model predicts that the empirically observed mean squared pricing error should be higher for contracts in which traders have complex motivations for placing bets.

Second, the model makes predictions concerning the time-series behavior of contracts. A regression of future single-period returns on past single-period returns for contract i yields a coefficient β_{R_i} that satisfies:

$$-0.5 < \beta_{R_i} \approx -\frac{(1-\rho)\,\sigma_{v_i}^2}{(1+\rho)\,\sigma_\epsilon^2 + (2+\rho)\,\sigma_{v_i}^2} < 0 \tag{11}$$

There are slight return reversals within the arbitrage bounds arising from price pressure from complex traders. Because the price pressure is not completely permanent, prices eventually regress toward the middle of the arbitrage bounds defined by fundamentals and transaction costs. Arbitrage ensures that these reversals never become large enough to imply profit opportunities. Empirically, the model predicts larger return reversals (β_{R_i}) for contracts in which complex traders play an important role.

4   Measures of Exchange Efficiency

This section evaluates the efficiency of the TradeSports exchange in light of the predictions of financial theory derived above. Nonparametric tests of market efficiency show that mispricing in sports game contracts is significantly greater, in both statistical and economic terms, than mispricing in financial contracts. This suggests traders with non-monetary motives play a

much more important role in sports game contract pricing than they do in financial contract pricing.

Tests of the time-series behavior of prices reveal that some contract prices overreact in equilibrium. The overreaction phenomenon applies only to the contracts on the exchange that are based on sporting events and does not admit profitable trading strategies. The future returns of financial contracts on the exchange are not significantly predictable from past returns. These findings provide further support for the simple model of arbitrage in the previous section.

Because most contracts on the TradeSports exchange are exposed only to idiosyncratic risk, it is easy to calculate their fundamental values. For an event futures contract with only idiosyncratic risk, fundamental value is equal to the probability that the underlying event occurs conditional on all publicly available information. The current price of a contract is the most obvious piece of public information. The model in the previous section predicts the current price will be approximately equal to the conditional probability that the event will occur. Unfortunately, a researcher cannot directly measure this conditional probability, thereby precluding a perfect assessment of a contract's fundamental value. However, by aggregating contracts into groupings based on their current price, one can estimate the fundamental value of *groups* of contracts as the observed frequency that the underlying *groups* of events occur. Comparing this theoretical portfolio value to the observed portfolio prices yields a test of whether portfolios of "riskless" contracts are mispriced.

4.1   Portfolio Formation

This idea lies behind the following procedure. I form ten portfolios by simulating ten trading strategies that buy specified contracts after publicly known events occur. These strategies create ten portfolios corresponding to ten pricing categories with ranges of 10 ticks each. The kth portfolio purchases a contract (and holds it until expiration) whenever a contract meets all of the following four criteria:

1. the average of the contract's current bid and ask prices is between $10k$ and $10(k-1)$ ticks;

2. the underlying contract event is sports-related or the contract has a duration of less than one week;

3. the spread between the contract bid and ask prices is less than or equal to 10 ticks;

4. and the portfolio does not yet contain the contract.

The first requirement divides contract prices into ten pricing categories with equal size. The first category includes contracts priced between 0 and 10 ticks and the last category includes contracts priced between 90 and 100 ticks. The decision to form categories based on pricing ranges of equal size reflects this study's interest in making inferences about the whole spectrum of agents' probability assessments. To ensure that prices do not reflect stale information, I measure the current price of a contract as the average of the highest bid price and the lowest ask price listed in the active order book. By comparison, the most recent transaction price captures information older than the current price.

The second criterion identifies contracts with only idiosyncratic risk, by selecting those that depend on sporting events and those that expire within one week or less.[12] Clearly, the outcomes of sporting events are unrelated to major macroeconomic shocks; and, in most weeks, the equilibrium price of holding a week's worth of macroeconomic risk is very close to zero.[13]

The third requirement excludes price observations that are potentially contaminated by measurement error. When the bid-ask spread is high, a researcher cannot accurately determine the current price with a reasonable degree of certainty. By contrast, low spreads imply the contract is (or will soon be) actively traded and that the active bid and ask orders are not just cheap talk. In other words, the expected stakes from submitting an order are high.

The fourth and final criterion ensures that no portfolio selects the same contract more than once. The trading algorithm defined by the four criteria generates a portfolio of distinct riskless contracts for each pricing category. Still, portfolios formed from different pricing categories will likely contain overlapping sets of contracts, implying that portfolio return realizations will be positively correlated. I address this issue later.

I consider two subsets of the contracts in the ten portfolios. The first subset contains sports game contracts, which yield payoffs contingent on whether a specific team wins a sporting event. Previous research has focused on these contracts under the implicit assumption that they are analogous to contracts based on financial events. I test the validity of this assumption by examining a second subset of contracts that are based on financial events.

---

[12]I have used many alternative criteria to test the robustness of the results and reached similar conclusions.

[13]The only contracts exposed to macroeconomic risk that comprise a significant fraction of the sample are contracts based on the close of daily and weekly financial index prices. Most of these contracts serve no obvious hedging function and they appear to be priced according to the model from the prior section.

4.2   Nonparametric Estimates of Mispricing

For each subset, I estimate the returns from holding a contract until expiration, conditioning on current contract prices. I use a locally weighted least squares algorithm that estimates the derivative of the conditional expectation at each contract price. The derivative equals the slope of the weighted least squares line through data from contracts with "neighboring" prices. I define neighboring prices to lie within an interval centered around the current price with a bandwidth equal to half the sample. I adopt a standard tricube weighting function, following Cleveland (1979). This procedure provides a nonparametric description of the returns from the sports game contracts and financial contracts in the ten portfolios.

Figure I presents the nonparametrically "smoothed" estimates of expiry returns conditional on current prices and contract type. Two patterns emerge from the table. First, on average, the estimates of expiry returns from sports contracts lie further from the horizontal axis, suggesting there is more mispricing in the market for sports game contracts than the market for financial contracts. Second, the shapes of conditional expiry returns for financial and sports game contracts are quite different, implying that distinct pricing patterns coexist in the two markets.

To measure the degree of mispricing in financial and sports game contracts, I calculate the mean squared pricing error for each set of contracts using the fitted values from the two series of smoothed expiry returns. It is useful to decompose each mean squared error (MSE) into a bias component and a variance component. For sports game contracts, the MSE is 4.06 ticks^2, the squared bias is 2.17 ticks^2, and the variance is 1.89 ticks^2. For financial contracts, the MSE is 1.59 ticks^2, the squared bias is 0.02 ticks^2, and the variance is 1.57 ticks^2 In other words, financial contracts have an MSE that is 60% lower, a squared bias that is 99% lower, and a variance that is 16% lower than sports game contracts. All three comparisons suggest financial contracts are more accurately priced than sports game contracts.

Next, I assess whether the difference in mispricing between the two types of contracts is statistically significant. The idea behind the test is to compute the difference in squared pricing errors between financial and sports game contracts over different pricing intervals. If these differences in squared errors are consistently negative, then it is likely that financial contracts are less mispriced than sports game contracts.

Following Künsch (1989), I employ the moving block bootstrap (MBB) procedure to generate a sampling distribution for the difference in squared pricing errors. The MBB differs from a conventional bootstrap, because it accounts for the dependence in the squared error data by repeatedly sampling observations from consecutive pricing intervals, rather

Figure I
Mispricing in Sports and Finance Contracts
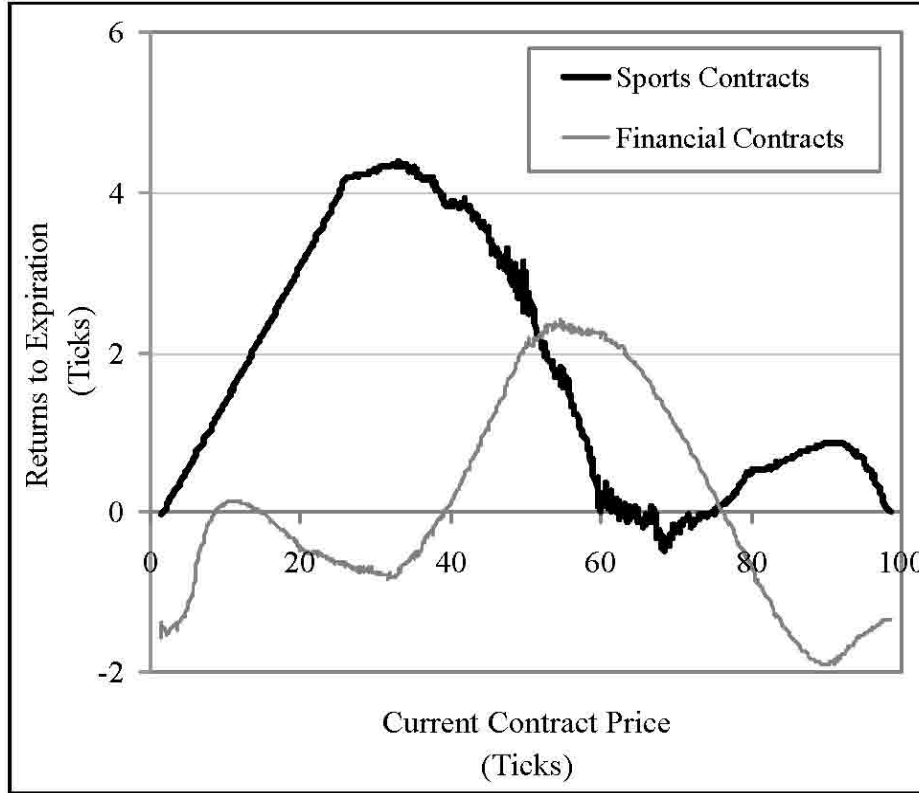
Source: Author's database based on TradeSports.com web sites from March 17, 2003 to October 4, 2003. To be included in this figure, a contract must meet all four of the portfolio criteria described in section 4.1. Each data point represents the fitted value from a locally weighted least squares regression of expiry returns on current prices. In each regression, I apply local weightings from the standard tricube weighting function to only neighboring current prices. The neighborhood for each data point is centered around the data point and includes one half of the sample. The original unsmoothed data come from the ten portfolios described in the text. There are 7,647 observations from sports game contracts and 6,127 observations from financial contracts in the portfolios.

than randomly selecting individual observations. In the MBB, I assume a block size equal to one-tenth of the total financial contract sample size (619 observations). I randomly sample consecutive blocks with this size from the distribution of financial squared pricing errors and sample the sports game squared pricing errors over a matching interval.

The average difference between these two sets of squared errors provides a measure of whether the sports game block or the financial block is more accurately priced. I average this average difference from 10 randomly sampled consecutive pricing blocks to obtain a

bootstrap estimator of the mean difference in squared errors of financial and sports game contracts. I repeat the estimation 1,000 times to find an empirical distribution for the bootstrap estimator. In 953 out of the 1,000 replications of the bootstrap estimator, the difference in squared error between financial and sports game contracts is negative. Thus, a one-tailed test of the hypothesis that the squared errors for sports game contracts and financial contracts on TradeSports are the same rejects the null at the 5% level. I conclude that the difference in mispricing between the two contracts is statistically significant *and* economically large.

Figure I also reveals that expiry returns to sports game contracts fall sharply as price increases within the 45 to 70 tick range. This phenomenon is closely related to the so-called reverse favorite-longshot (RFL) bias. Woodland and Woodland (1994 and 2003) find that the returns to betting on baseball teams to win decline monotonically with the probability that the team will win. Woodland and Woodland (2001) replicate this result for hockey teams. These studies estimate that the gross returns (without commissions) to betting on the teams with lower probabilities of winning are 1.5% for baseball games and 5.2% for hockey games.

To test for this regularity in the TradeSports data, I follow a procedure similar to the one used to create Figure I. Instead of using portfolios of contracts with prices designed to span the full range from 0 to 100 ticks, I focus on the returns to contracts that are purchased before any major news about the event has transpired. Thus, I use only the first observations on sports game contracts and financial contracts that meet criteria (2) through (4) listed above. This procedure yields a very uneven distribution of contract prices with many observations in the middle and very few near the extremes.

I employ a lowess algorithm based on only the observations on the initial contract prices rather than observations over the lifetime of the contract. This lowess algorithm differs from the earlier lowess algorithm in two significant ways. First, I select smoothing bandwidths for the two sets of contracts that contain the same number of observations.[14] The benefit of using bandwidths with equal numbers of observations is that the two smoothed estimates have approximately the same precision. The cost of this choice is that the bandwidths are not based on the same pricing ranges.[15] Second, I do not apply weightings to the observations in each local regression. Figure II displays the smoothed data from this lowess procedure for both types of contracts.

---

[14]Both bandwidths are based on half of the total number of financial contracts (770 contracts), because there are fewer observations on financial contracts.

[15]Adopting a lowess procedure with bandwidths based on the same pricing ranges generates similar qualitative results to those presented here.

Figure II
The (Reverse?) Favorite-Longshot Bias
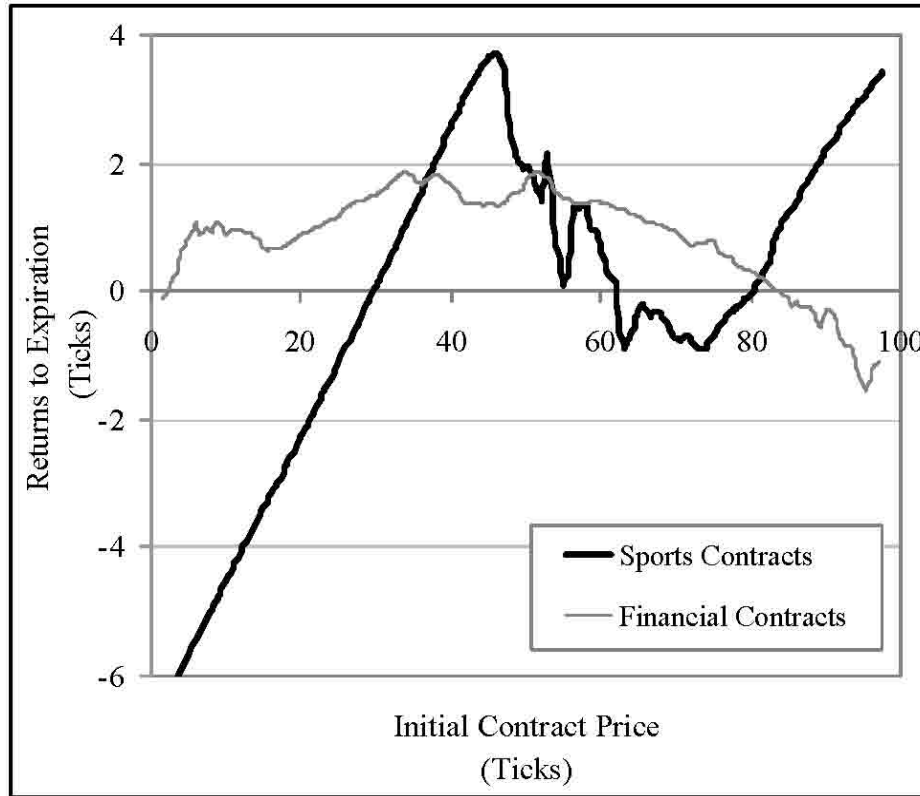
Source: Author's database based on TradeSports.com web sites from March 17, 2003 to October 4, 2003. To be included in this figure, a contract must meet the last three of the four portfolio criteria described in section 4.1. Each data point represents a locally weighted regression of expiry returns conditional on current prices. In each least squares regression, I apply equal weight to all neighboring current prices. The neighborhood for each data point is centered around the data point and includes 770 contracts (half of the total financial contract sample). The original unsmoothed data come from the one large portfolio described in the text. In this portfolio, there are 4,133 observations from sports game contracts and 1,739 observations from financial contracts.

The figure reveals that there is a pronounced RFL bias in sports contracts initially priced between 45 ticks and 75 ticks. The maximum amplitude of the bias from the peak return to the trough return is 4.6 ticks. Most of the data lies in the pricing range where the RFL bias prevails, because the TradeSports exchange lists the majority of sports game contracts from the perspective of the favored team.[16] It is interesting to note that the sports game

---

[16]A few "favorites" become slight underdogs before the initial price has been recorded, which explains why there are some sports game contracts with initial prices less than 50 ticks.

contracts on the web site with prices straying substantially from 50 ticks follow the pattern of the favorite-longshot bias. In other words, the returns to betting on a team increase as the probability of the team winning increases, as long as this probability is not near 50%.[17]

Both the favorite-longshot bias and the reverse favorite-longshot bias in Figure II are consistent with findings from earlier studies. Researchers have found the RFL bias in sports game contracts that are based on one team winning the game. By their nature, these contracts have prices near 50%, because most sports games feature two teams with comparable ability. The favorite-longshot bias has been observed in gambles at horsetracks. Because there are typically ten or so horses in each race, the vast majority of bets at the horsetrack yield payoffs less than 25% of the time. A simple interpretation of prior studies is that the favorite-longshot bias applies only to contracts based on unlikely events, whereas the RFL bias applies to events that are moderately likely to occur. Of course, contracts that depend on unlikely events also implicitly depend upon likely events (the negation of the unlikely event), so it is not surprising that the favorite-longshot bias applies to both unlikely and likely events.

From Figure II, it appears that framing does not have much of an impact on sports game contract pricing. Reframing contracts in terms of the other team winning is equivalent to rotating the plot of the conditional expectation function 180 degrees along an axis that points into the figure and is centered at expiry returns of 0 ticks and an initial contract price of 50 ticks. Because the sports game conditional expectation function has a slanted z-shape, this 180-degree rotation has little impact on the qualitative features of the function. Both the favorite-longshot bias and the reverse favorite-longshot bias remain prominent in similar ranges of the reframed version of the sports game conditional expectation function.

The financial contracts *do not* show the same regularities as sports game contracts. There is no reverse favorite-longshot bias for financial contracts priced near 50 ticks. And the conditional expectation function for financial contracts does not have any notable symmetry. With the exception of a minor positive bias in returns, financial contracts are approximately correctly priced. Although most of the financial contracts appear to be priced below fundamental value, this bias is not statistically or economically significant. A comparison of the results in Figure I and Figure II suggests that any positive return bias in financial contracts goes away over time.

To summarize the results from Figures I and II, the prices of sports game contracts

---

[17]This regularity has been noted, but largely ignored, by other authors such as Woodland and Woodland (1994).

appear more biased and less precise than the prices of financial contracts on the TradeSports exchange. There are also more psychologically motivated patterns in the sports game data. These patterns include both the reverse favorite-longshot bias *and* some evidence of the favorite-longshot bias. The next section adopts a parametric framework to assess the ability of behavioral theory to explain the sports game and financial contract pricing patterns. The functional form estimated embeds both rational and behavioral models.

4.3   Estimating a Parametric Probability Weighting Function

One of the most prominent and well-respected behavioral theories of probability assessment comes directly from Kahneman and Tversky's (1979) prospect theory. Based on a substantial body of experimental evidence, Kahneman and Tversky hypothesize that agents tend to overweight the likelihood of low probability events and underweight the likelihood of high probability events. Prelec (1998) posits a specific functional form that describes the relationship between actual probabilities and agents' perceptions of these probabilities. He finds that the one-parameter version of this probability weighting function captures the most important features of extant experimental evidence on agents' beliefs. Prelec (2000) provides an axiomatic justification for the weighting function:

$$w(p) = \exp\left[-\left(-\ln p\right)^{\alpha}\right], \alpha > 0 \tag{12}$$

The α parameter controls the degree of curvature in the function. Except when $\alpha = 1$, the Prelec probability weighting function has a unique fixed point at $p = 1/e \approx 0.37$. It is also regressive when $\alpha < 1$, meaning it overweights low likelihood events and underweights high likelihood events as in prospect theory. These features are all consistent with experimental empirical estimates. Also, note that the Prelec probability weighting function nests the expected utility model as a special case ($\alpha = 1$). Gonzalez and Wu (1996) and Prelec (2000) estimate $\alpha = 0.65$ and $\alpha = 0.74$, respectively, from their experimental data.[18] Both studies

---

[18]Other studies have estimated a probability weighting function that is linear in log odds (Lattimore, Baker and Witte, 1992), which possesses the same qualitative properties as the Prelec weighting function:

$$w(p) = \frac{p^{\gamma}}{[p^{\gamma} + (1-p)^{\gamma}]^{1/\gamma}}, \gamma > 0 \tag{13}$$

In fact, this function's curvature parameter γ is quantitatively similar to the curvature parameter α in Prelec's specification. Gonzalez and Wu (1999) argue: "For typical values of probability used in most empirical studies..., it will not be possible to distinguish the linear in log odds function from the Prelec function because both functions can be linearized, and these linearized forms themselves are closely linear

statistically reject the EU restriction on α, finding it is significantly less than 1.

Of course, experimental investigations of the probability weighting function use individual choice data whereas each data point in this study represents a market equilibrium resulting from the interaction of many traders. Under the assumption that all market participants behave according to the same probability weighting function, analyzing the two different data sources will produce identical conclusions. On the other hand, if certain traders behave rationally and exert a disproportionate influence on market prices, market-based field estimates of α will be much closer to 1 than will individual-choice-based experimental estimates.

The predictions of prospect theory for the TradeSports data can only be tested jointly with a theory of how many behavioral traders are present in the market and how those traders affect prices. Implicitly, this study is testing prospect theory under the assumption that arbitrageurs use the theory's probability weighting function and that the model in the previous section accurately describes their influence on prices. All the equations in the previous section still apply when expected dividends are calculated using traders' perceived probabilities rather than objective probabilities.

One paper that has used real market-based data to estimate α from the Prelec weighting function is Jullien and Salanie (2000). The authors arrive at mixed conclusions from their data set, which contains horse track odds and results. They statistically reject the expected utility hypothesis in favor of the two-parameter Prelec specification. But the authors estimate $\alpha = 0.879$, which they consider to be a small departure in economic terms from the EU restriction of $\alpha = 1$.[19] A simple interpretation of this result is that some, but not all, traders in the Jullien and Salanie data use the prospect theory weighting function.

Thus, the jury is still out on the ability of behavioral theories to predict probability weighting functions derived from real market data. Using the data from Figure II, this study attempts to estimate α in the Prelec weighting function. Whereas Prelec (2000) and most other experimental studies that estimate α are able to observe (or even manipulate) the true event probability (p) agents' face, this study must infer true probabilities from data on observed event frequencies. To estimate α, Prelec (1998, 2000) suggests transforming the equation for the weighting function by taking logarithms of each side twice. In double-

---

for the probability range (0.01, 0.99)" (p.140). With this similarity in mind, one can infer estimates of α from studies that have estimated the linear in log odds weighting function. Tversky and Kahneman (1992), Camerer and Ho (1994), and Gonzalez and Wu (1996) find estimates of between 0.56 and 0.71 for the curvature parameter in the linear in log odds specification, which is quantitatively similar to the values of α estimated from other studies.

[19]The data in Jullien and Salanie (2000) follow the well-known pattern of the favorite-longshot bias. The expected return to betting on a horse increases monotonically with the probability of the horse winning.

log space, the relationship between probability perceptions and true probabilities is linear, allowing Prelec estimate the weighting function with classic OLS techniques.

Unfortunately, simply tweaking Prelec's estimation methodology to accommodate frequency data by replacing true probabilities with event frequencies is fraught with difficulties. First, events must be aggregated into pricing categories to use frequencies as estimates of probabilities.[20] This categorization process does not efficiently use all the information about contract prices available in this data set. Second, category frequencies are imprecise measures of true category probabilities. This measurement error in the independent variable causes attenuation bias in the OLS estimates.[21]

Maximum likelihood estimates of α based on contract prices and event dummy variables (indicating event success or failure) overcome these limitations. I base the maximum likelihood estimates on the data from Figure II, which contains $3,538$ independent observations on contract prices and the success or failure of underlying contract events. This number is substantially smaller than the $5,818$ raw observations included in Figure II, because the figure contains some related contracts.[22] The TradeSports exchange groups these related contracts into sets called subcategories.[23] To ensure independence across contracts, I use only one contract per sports game and one financial contract per subcategory in the maximum likelihood estimates.

I separately test whether contract prices equal fundamental values for sports games and financial contracts traded on TradeSports. Please refer to the appendix for a complete description of the estimation procedure and the algorithm use to solve for the maximum likelihood estimators.

---

[20]Specifically, events must be grouped into categories in which all category frequencies are positive to avoid taking the logarithm of zero.

[21]I have performed an OLS regression of double log portfolio prices on double log portfolio event frequencies. Despite the attenuation bias, estimates of α from this regression closely match the maximum likelihood estimates described below.

[22]Specifically, Figure II contains multiple financial contracts whose outcomes depend on the same or related financial indices. The figure also contains multiple sports contracts based on one team winning by different margins of victory.

[23]The top layer of the TradeSports grouping hierarchy is the contract type (*e.g.*, financial), the next layer is the contract category (*e.g.*, indices), the next layer is the contract subcategory (*e.g.*, daily indices close for March 19), and the final layer is the contract itself (*e.g.*, Dow Jones index to close higher than the previous day on March 19). For example, the following contracts are listed in the subcategory "daily indices close for March 19": Dow Jones to close higher; Dow Jones to close more than 100 points higher; Dow Jones to close more than 100 points lower; S&P to close higher; and Nasdaq to close higher.

The event outcomes underlying contracts from the same subcategory are obviously related. But event outcomes from different subcategories, such as "daily indices close for March 19" and "daily indices close for March 20," are approximately uncorrelated. The weighting procedure is designed to account for this correlation.

Table IV
Maximum Likelihood Estimates of the Prelec Probability Weighting Function

Parameter	Games Contracts	Financial Contracts
$\widehat{\alpha}$	0.921	0.960
Std Error	0.052	0.085
ln L	-641.3	-232.3
ln L (EU)	-642.3	-232.4
LR p-value	0.164	0.656

Source: All estimates are based on the independent observations from Figure II. The one-parameter Prelec weighting function is given by equation (12). ln L denotes the unrestricted maximum of the log likelihood function; and ln L (EU) denotes the maximum of the log likelihood function under the expected utility parameter restriction on the Prelec weighting function ($\alpha = 1$). The LR p-value is the likelihood ratio p-value from testing the expected utility restriction.

By themselves, the estimates in Table IV suggest that both types of contracts are quite efficient. Of the two estimates, the estimates of $\widehat{\alpha}$ based on sports game contracts most closely resemble experimental evidence and the Jullien and Salanie (2000) estimates. But drawing inferences from Table IV alone can be somewhat misleading. The nonparametric analysis of sports game contracts in section 4.2 suggests that the parametric weighting function has a difficult task in trying to fit both the reverse favorite-longshot bias and the favorite-longshot bias. The appropriate interpretation of the estimated value of $\widehat{\alpha} < 1$ for sports game contracts is that the favorite-longshot bias slightly dominates the reverse favorite-longshot bias.

Because the parametric tests impose an inappropriate functional form on the mispricing relationship, they do not reveal significant differences in mispricing between sports game and financial contracts. By contrast, the nonparametric tests that let the data speak show that financial contracts are more accurately priced than sports game contracts by a statistically and economically significant margin. In cross-sectional estimates, the mean squared pricing error is over 60% lower for financial contracts. It is still possible, however, that time-series tests of efficiency tell a different story. The next subsection tests whether the prices of sports games contracts and financial contracts on TradeSports follow a random walk over time.

4.4 The Market Response to Information

Many financial models predict asset market prices will follow a random walk—*i.e.*, they will not be predictable based on past information. As explained section 3, however, a reasonable

model of trading behavior on TradeSports implies a small degree of negative serial correlation between returns. If some traders are motivated by non-monetary concerns, such as devotion to their favorite sports team or direct utility from gambling, they can exert pressure on contract prices. Still, the rational traders in the model ensure that prices do not stray beyond the arbitrage bounds from equations (8) and (9). Despite the negative serial correlation in equilibrium prices, the model in section 3 predicts no trading strategies based on past returns will generate excess returns in equilibrium.

Equation (11) predicts negative serial correlation in contracts in which there is a lot of "complex" trading—*i.e.*, trading motivated by non-monetary concerns. Presumably, sports games meet this criterion, because some sports fans prefer to bet on their favorite team or derive utility from the excitement of gambling on sporting events. Based on the assumption that these motives do not generalize to other contexts, such as financial markets, equation (11) predicts smaller return reversals in financial contracts.

This theory of overreaction differs from behavioral theories of overreaction that are based on mistaken beliefs. For example, Avery and Chevalier (1999) argue that investor sentiment leads to overreaction of betting lines in the professional football market. Most psychological theories of overreaction posit that salient news provokes overreaction from traders. Some events underlying contracts on the TradeSports exchange attract a great deal of media coverage, suggesting they are salient in the minds of traders that receive their information from the media. According to behavioral theory, traders overreact to positive contract returns when these returns are accompanied by salient news. This phenomenon generates excessive buying activity in salient contracts with positive past returns, thereby pushing the prices of these contracts above their fundamental values and leading to negative expected returns for these contracts.[24] Thus, the behavioral theory of overreaction suggests return reversals should be greatest for contracts based on events with a lot of media coverage.

First, I test the central prediction from the rational model of return reversals. Figure III compares the magnitude of reversals in contracts based on sports games and financial events. All of the observations in Figure III describe contracts that are within one day of expiration and have been traded in the past 30-minutes.

Each sports game contract yields payouts when one team wins the game. Each financial contract yields payouts when a financial measure, such as the Dow Jones Industrial Average, falls within a certain range of values at the end of the trading day. This is directly analogous

---

[24]Because all contracts on TradeSports expire at fundamental value, contracts with prices greater than fundamental value have negative expected future returns.

to a sports contract whose outcome depends on whether the difference between two teams' scores falls within a certain range at the end of a sporting event. For each of the two contract types, I perform a locally weighted regression of future 30-minute returns on past 30-minute returns. The local regressions are linear least squares estimates of the relationship between future and past returns in the neighborhood of each past return. Again, I employ a smoothing bandwidth equal to half the sample. I only count the contracts that could be sold 30 minutes after portfolio formation.[25]

Figure III suggests reversals on the TradeSports exchange are largely confined to contracts based on sporting events. For example, for sports contracts, the difference between future returns following past returns of -20 and +20 ticks is 6.29 ticks. The difference in future returns is 3.14 ticks after past returns of -10 and +10 ticks. For financial contracts, the corresponding differences in future returns are just 0.61 ticks and 1.53 ticks; both of these values are statistically and economically insignificant. The sports game contract reversals are four to five times greater in magnitude.

The rational theory of overreaction advanced in section 3 accurately predicts this effect. Members of the TradeSports exchange that trade contracts based on sports games probably have non-monetary motivations. However, the salience-based theory of overreaction also predicts this effect under the assumption that sports games capture trader attention better than financial events. Without additional information, one cannot distinguish between the rational and behavioral interpretations of the evidence in Figure III.

To better understand the cross-section of return reversals, I examine linear regressions of future returns on past returns. Because the nonparametric estimates in Figure III approximately lie in a straight line, I prefer to use classic linear regression methods for their ease of interpretation. By interacting past returns with environmental variables, I measure how the magnitude of overreaction varies across different situations. The rational theory of overreaction suggests it should be strongest when traders with non-monetary motives play a large role in price determination, whereas the behavioral theory of overreaction predicts return reversals will be strongest when events are most salient.

Contracts based on sports games provide an opportunity to distinguish the two theories, because they provide separate measures of traders' non-monetary motives and the salience

---

[25]In practical terms, it might be impossible to sell a contract 30 minutes after portfolio formation if either no traders bid on the contract or the contract expires within 30 minutes of portfolio formation. I censor such observations for the purpose of brevity, but including these observations does not substantially alter the results. In fact, under the alternative assumption that each portfolio includes contracts that cannot be sold until they expire, the results become slightly stronger.

Figure III
Return Reversals in Different Types of Contracts
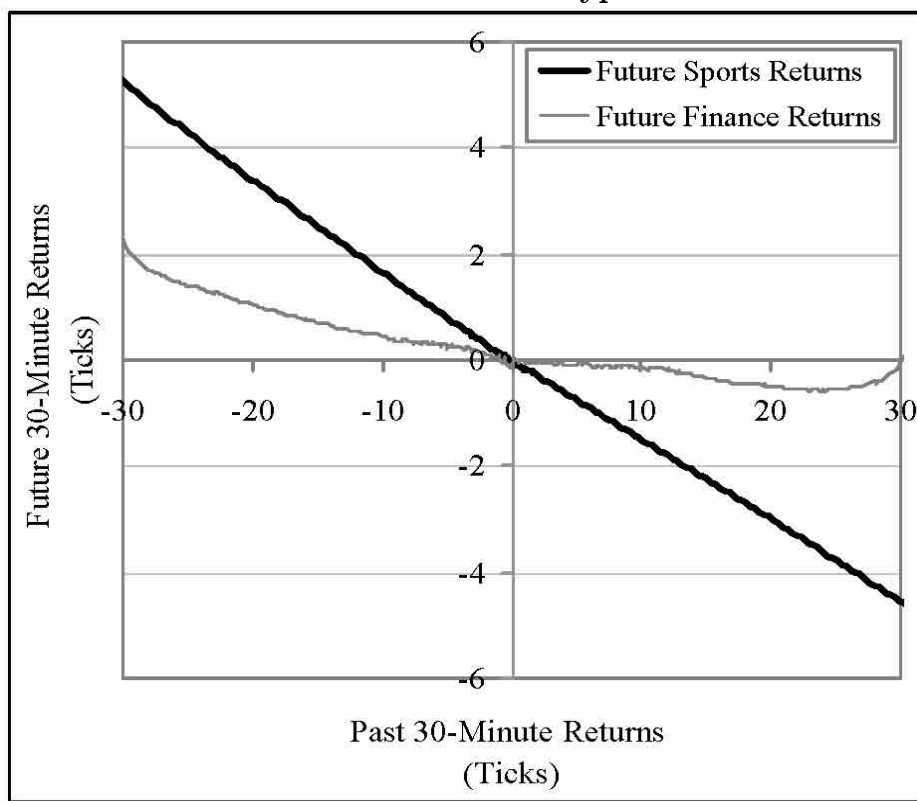
Source: Author's database based on TradeSports.com web sites from March 17, 2003 to October 4, 2003. To be included in this figure, a contract must have been traded in the past 30-minutes. Each data point represents a locally weighted regression of expiry returns conditional on current prices. In each least squares regression, I apply local weightings from the standard tricube weighting function to only neighboring current prices. The neighborhood for each data point is centered around the data point and includes one half of the sample. There are 25,602 observations from sports game contracts and 4,222 observations from financial contracts in the original unsmoothed data.

of contract events. Because fans of specific sports teams have compelling non-monetary motives to trade, the rational theory predicts contracts based on sports games will exhibit greater return reversals than other contracts. By contrast, the salience-based theory predicts overreaction will be most pronounced for well-publicized sporting events, such as "featured" sports games. The rational theory makes no obvious prediction about featured games as opposed to non-featured games. If anything, fans probably represent a smaller proportion of viewers in featured games, because truly devoted fans are more likely to watch non-featured games than are casual fans.

To measure the magnitude of return reversals in different environments, I regress future 30-minute returns (FUT.RET) on past 30-minute returns (LAG.RET), four interaction terms with past returns, and control variables. I separately measure reversals for contracts based on "featured" events in the media, contracts based on sports games, contracts with greater volume, and contracts with greater bid-ask spreads.

The TradeSports exchange categorizes its sports games by whether or not they are featured on a major television network, such as ESPN or Fox. I create a dummy variable equal to 1 for observations on featured contracts and equal to 0 for observations on non-featured contracts. Multiplying this dummy variable by past returns creates an interaction term (FEAT.INT) that equals past returns only for observations on contracts based on featured games and equals zero otherwise. I create an analogous interaction term (GAME.INT) by multiplying past returns by a dummy variable that equals 1 only for observations on contracts based on sports games.

Because differences in liquidity across contracts could affect the magnitude of return reversals, I create an additional interaction term to capture this effect. The variable SPRD.INT is the interaction of the contract's bid-ask spread and past 30-minute returns (LAG.RET). I also attempt to measure the impact of increases in stakes by measuring how overreaction varies with the volume of transactions. The interaction variable VOL.INT is the product of the log of demeaned contract volume and past 30-minute returns (LAG.RET).

To control for the influence of temporary price pressure, I construct a variable called BID.RAT, which equals the contract's bid ratio (or net bids divided by the total interest for a given contract).[26] This variable ranges from -1 to 1. When BID.RAT equals zero, the number of bids and asks are equal, indicating the contract price is not influenced by temporary price pressure. If BID.RAT is positive, the contract price is likely to increase in the short-term, because there are more offers to buy than offers to sell at the current prevailing price.

In the final specification, I regress future 30-minute returns (FUT.RET) on past 30-minute returns (LAG.RET), the four interaction terms described above (GAME.INT, FEAT.INT, VOL.INT, and SPRD.INT), and the contract bid ratio (BID.RAT):

$$
\begin{aligned}
FUT.RET_{it} &= \alpha + \beta_1 LAG.RET_{it} + \beta_2 GAME.INT_{it} + \beta_3 FEAT.INT_{it} \quad (14) \\
&+ \beta_4 VOL.INT_{it} + \beta_5 SPRD.INT_{it} + \beta_6 BID.RAT_{it} + \epsilon_{it}
\end{aligned}
$$

---

[26]The net bids are equal to the difference between the total quantity of offers to buy and the total quantity of offers to sell. The total interest is the sum of these two quantities.

I estimate equation (14) using standard OLS techniques. I employ the Huber-White formula for estimating standard errors in the presence of heteroscedasticity, because future returns are typically more volatile after past returns with a large absolute value. The Huber-White robust standard errors are much larger than the traditional standard error estimators, confirming the presence of heteroscedasticity.

Table V reports the results from the OLS regression based on equation (14). The coefficient on past returns (LAG.RET) is almost exactly equal to zero, suggesting that there is no return reversal in contracts based on events other than sports games. The negative coefficient on the interaction term GAME.INT suggests the presence of fans with non-monetary motives that trade contracts based on sports games, as hypothesized in the rational model The magnitude of this coefficient shows that, on average, price movements in sports games contracts are followed by return reversals equal to 13.8% of the original movement. On the other hand, the approximately zero (and statistically insignificant) coefficient on the featured games interaction term (FEAT.INT) contradicts the salience-based overreaction model's central prediction. There is almost no additional return reversal in well-publicized, featured sports games.

Table V
OLS Regressions of Future Returns on Past Information

	LAG.RET	GAME.INT	FEAT.INT	VOL.INT	SPRD.INT	BID.RAT
Coefficient	0.001	-0.138	-0.003	-0.011	-0.002	0.419
Std Error	(0.027)	(0.068)	(0.080)	(0.009)	(0.006)	(0.183)
t-statistic	0.03	-2.04	0.55	-1.15	-0.33	2.29

Source: Author's database based on TradeSports.com web site. The regression is based on 51,650 observations of future and past returns at 30-minute horizons. The regression only includes observations in which the contract was actively traded over the past 30 minutes. The R-squared statistic is 0.0100. The dependent variable is FUT.RET (the future return for a contract). LAG.RET is the past return. The coefficients on all interaction effects with LAG.RET are additive; GAME.INT is the interaction of a dummy variable indicating a sports game contract and LAG.RET; FEAT.INT is the interaction between a dummy variable indicating a featured game contract and LAG.RET; VOL.INT is the interaction of the log of demeaned contract volume and LAG.RET; and SPRD.INT is the interaction of the contract's bid-ask spread and LAG.RET. BID.RAT is the ratio of net bids to total interest for a contract. "Std Error" is the Huber-White heteroscedasticity consistent estimate of each coefficient's standard error.

Interestingly, there appears to be no difference in the magnitude of return reversals for games with greater liquidity or greater stakes. Both the spread and volume interaction terms

are statistically and economically insignificant. Finally, the positive coefficient on BID.RAT implies that temporary price pressure exerts a strong effect on future returns.

Statistical tests based on Table V *cannot* reject the hypothesis that the coefficient on the LAG.RET variable is equal to zero. In other words, prices in non-sports contracts follow a time-series process that is statistically indistinguishable from a random walk, as predicted by the rational model. The preponderance of evidence from cross-sectional and time-series tests points toward the rational model of the TradeSports exchange.

4.5   Testing the Profitability of Trading Strategies on TradeSports

Nevertheless, if overreaction in *sports games* presents profit opportunities in equilibrium, then the rational model requires some modification. To look for these opportunities, I assume that traders have access to the results from Table V and use these results to predict future returns. In other words, I use 20-20 hindsight to look for abnormal returns to trading strategies in contracts based on sporting events "near" their expiry (when overreaction is strongest).

I consider trading strategies that hold contracts based on sports games until the outcome of the game is known, because other strategies must incur bid-ask spread costs twice and require that markets for contracts remain liquid until expiry (a condition not always met in practice). Because the precise ending time of a sporting event is rarely known in advance, traders cannot implement trading strategies that condition on the exact amount of time left to expiry. In practice, however, I must use the information on time left until expiry to identify when games start. The trading strategy buys (sells) a contract at the lowest current ask price (highest current bid price) whenever it meets the following criteria:

1. the contract is based on a sports game;

2. the contract is within 4 hours of expiry;

3. the contract has experienced bad (good) news within the last data recording period, defined as a past 30-minute return of less than 5 ticks (greater than 5 ticks);

4. and the current bid-ask spread is less than or equal to 3 ticks.

Over the 6-month data recording period, this strategy buys 340 contracts and sells 412 contracts and closes each contract position when the contract expires. The average gross return from the strategy is only 1.29 ticks, while the standard deviation of the strategy is 40.79

ticks.[27] After subtracting the round-trip commission cost of 0.8 ticks, this trading strategy yields 0.49 ticks per trade. This small expected profit is somewhat surprising considering the data mining and research required to implement such a strategy. The slightly positive return to betting on return reversals in sports game contracts could be ascribed to chance.

Perhaps combining the results from both the time-series and cross-section tests of efficiency allows an improvement upon the simple overreaction-based trading strategy. Specifically, I consider a trading strategy based on both the reverse favorite-longshot bias (RFL) and overreaction.[28] The RFL bias suggests betting on underdogs should be more profitable than betting on favorites. The overreaction bias suggests it should be even more profitable to bet on underdogs that were, at some point in the past, favorites.

I search for this bias in contracts based on sports games on the TradeSports exchange. I identify favorites as those contracts that yield payouts contingent upon one team winning and have a price equal to or exceeding 50 ticks; underdogs are contracts with prices below 50 ticks. I form a portfolio of *new* underdogs that selects contracts as soon as their prices fall below 50 ticks and they have positive volume. I also create an analogous portfolio of contracts that are *new* favorites.

Note that these portfolio definitions could apply to any contract, including financial contracts. Because the authors of previous studies on sports wagering markets have emphasized their similarities to financial markets, one could expect that the reverse favorite-longshot bias and overreaction generalize to financial contracts when these contracts have a similar structure to sports contracts. Just like sports contracts, all the financial contracts on TradeSports in my sample yield payouts of 100 ticks or 0 ticks depending on the outcome of the event. Thus, I create two portfolios consisting of contracts based financial events that have become underdogs (price fell below 50 ticks) and financial events that have become favorites (price rose above or is equal to 50 ticks) to see whether these portfolios differ from the portfolios of sports game new favorites and new underdogs.

---

[27] The standard deviation of the average gross return to the strategy is 1.49 ticks, implying that the 95% confidence interval for the average gross return is -1.69 ticks to 4.27 ticks.

[28] Even though the favorite-longshot bias appears in certain regions of the TradeSports sports game contract data, I consider a trading strategy that does not account for this third bias. This decision is based on many reasons. First, the favorite-longshot bias is less robust and based on less data than the RFL and overreaction biases. Second, the favorite-longshot bias is strongest when contracts are first listed on the exchange (see Figures I and II). But I cannot apply the overreaction bias to contracts that have no past trading history. Third, I did not expect to find the favorite-longshot bias in sports games, so using this bias in a trading strategy would be data mining. Fourth, conditioning the trading strategy on a third bias reduces the number of observations in each portfolio created by the trading strategy, thereby decreasing the statistical power of the tests that follow.

Table VI reports all four portfolio returns under the assumption that each portfolio holds its contracts until they expire (*i.e.*, the underlying events occur). The table shows that the RFL and overreaction biases are present for contracts based on sports games. This result is consistent with previous studies and the results from earlier in this paper.

<div align="center">

Table VI
The RFL and Overreaction Biases in Different Types of Contracts

	Sports Games	Finance
New Favorites Return	0.45	2.00
New Underdogs Return	5.36	1.05
Difference in Returns	-4.91	0.95
t-statistic	-3.71	0.37

</div>

Source: Author's database based on TradeSports.com web site. The portfolios of new favorites for sports game and financial contracts include 3,017 and 1,023 observations respectively. The new underdog portfolios for sports games and financial contracts include 790 and 1,366 contracts respectively. The standard errors on portfolio returns assume independence across sports games and correlations of 0.5 for financial contracts.

However, there is no analog to these results in financial contracts. If anything, the opposite effect prevails: financial "bets" that have become likely to occur (new favorites) offer slightly higher returns than financial bets that have become unlikely (new underdogs). Also, both groups of financial contracts have positive returns on average. Neither the slight favorite-longshot bias nor the positive returns effect is statistically or economically significant.[29] This evidence lends additional support to the hypothesis that the prices of financial contracts do not suffer from the same biases as the prices of sports game contracts on Trade-Sports. In fact, in all of the parametric and nonparametric tests in this study, I cannot reject the hypothesis that financial contracts are priced according to their fundamental values.

The reverse favorite-longshot bias in sports contracts potentially admits profitable trading strategies. To investigate this possibility, I simulate a trading strategy that purchases the 731 contracts in the new underdog portfolio for sports games that have bid-ask spreads less than or equal to 5 ticks. After incurring the bid-ask spread cost and round-trip transaction costs, this strategy yields average returns of 3.79 ticks on invested capital of only 37.13 ticks. Moreover, a trader obtains this 10.2% expected return to betting on new underdogs over a horizon of only one day for the typical sports game.

---

[29]The returns from many of the financial and other contracts included in this test are correlated with one another, because they depend on related underlying events. The results in Table VI are robust to any reasonable estimate of this correlation, including the assumption of independence. I calculate the standard errors in the table under the assumption that this correlation is equal to 0.5.

The *statistical* inefficiency in sports games is extremely unlikely to arise from data mining, because the p-value for the t-statistic on new underdog returns is only 0.0002. To assess the impact of data mining, I employ the Bonferroni adjustment to the p-value. I would have had to consider 250 trading strategies in this paper to inflate the original p-value of 0.0002 to 0.05. A truly upper bound on the number of trading strategies considered in this paper would be 128. This number equals 4 ways to partition the data into pricing groups times 4 potential time horizons for trading strategies times 2 categories of contracts times 2 potential cross-section biases (reverse and normal favorite-longshot bias) times 2 potential time-series biases (overreaction and underreaction). Of course, many of the strategies just described have positively correlated returns and I had strong priors about which way to partition the data, so 128 is indeed an upper bound on the appropriate Bonferroni adjustment factor. Even this extreme upward adjustment of the p-value does not alter the conclusion that sports game contracts are mispriced by a statistically significant margin.

Nevertheless, it is important to mention the possibility that the *economic* inefficiency (the profitable trading strategy) in sports game contracts *did* arise from chance. The t-statistic for the estimated expected *net* return of 3.79 ticks (after commission and spread costs) could be attributable to chance. The standard deviation of the estimated new underdog return is 1.66 ticks, so the true expected return from this trading strategy could turn out to be as low as 0.54 ticks (the lower end of the 95% confidence interval). Very little data mining is required to find an economic return of this magnitude. Perhaps it is not surprising that one trading strategy out of all those implicitly considered in this paper yields "significantly" profitable returns.

Even if one supposes betting on new underdogs in sports game contracts yields positive profits, this finding is not necessarily inconsistent with a broad interpretation of the efficient markets doctrine. In the Grossman and Stiglitz (1980) model of costly information acquisition, traders with private signals about fundamental value earn positive economic rents in equilibrium. A trader implementing the new underdog strategy on TradeSports could only buy an average of 180 underdog contracts per transaction, because TradeSports is not a perfectly liquid market. Over the sample period of 6 months, the new underdog trading strategy would have produced *at most* $700 per day. This is not an outrageous amount of money for a Wall Street professional, because implementing the new underdog strategy requires constant monitoring and research.

5 Conclusions

A large body of research on betting markets attempts to measure the implied expectations of economic agents in the hope that these measurements generalize to real financial markets. The results above suggest future researchers should proceed with caution before drawing analogies between wagering markets and real financial markets. Inefficiencies that are present in wagering markets for sports games do not necessarily have counterparts in financial markets.

The nonparametric tests of market efficiency reveal that mispricing in sports game contracts is significantly greater, in both statistical and economic terms, than mispricing in financial contracts. The systematic patterns of mispricing in sports game contracts are consistent with findings from earlier studies, but these patterns do not appear in financial contracts on the exchange. For example, there is a significant reverse favorite-longshot bias in contracts based on sports games on TradeSports; but there is no such bias in financial contracts. I also find some support for the finding that sports bettors overweight the likelihood of longshots winning and underweight the likelihood of favorites winning (Thaler and Ziemba, 1986; and Jullien and Salanie, 2000), but this bias does not apply to the financial markets on TradeSports.

Furthermore, nonparametric and parametric time series tests of efficiency show overreaction in contracts based on sports games listed on the TradeSports exchange, but not in contracts based on financial events. This evidence suggests overreaction in wagering markets (Avery and Chevalier, 1999) is a different phenomenon from overreaction in financial markets (*e.g.*, De Bondt and Thaler, 1985; and Huberman and Regev, 2002). Taken collectively, this set of results reveals that conclusions from sports wagering market data may not apply to real financial markets.

The presence of inefficiencies in sports game markets and the absence of these same inefficiencies in financial markets is somewhat surprising. Financial and sports contracts are traded using the exact same mechanism. The two markets have similar liquidity and volume; and the contracts themselves are structured almost identically. There is even anecdotal evidence that suggests many of the same traders participate in both markets. Yet one market appears efficient while the other does not.

The efficiency of the financial contracts is perhaps more remarkable than the inefficiency of the sports contracts. The financial contracts have very low liquidity and low stakes (volume) relative to traditional financial markets, such as those conducted on the NYSE. But the results from time series tests show that return reversals do not occur even in the financial

contracts with the lowest liquidity and lowest stakes (volume). Furthermore, the financial contracts on TradeSports are fundamentally speculative–*i.e.*, these contracts do not finance capital investments and they serve no obvious hedging purpose. Such speculative contracts could attract traders with non-standard preferences or expectations, thereby diminishing market efficiency.

Finally, based on the observed inefficiencies in sports game markets, it is reasonable to assume that at least some traders on the TradeSports exchange conduct trades based on non-standard preferences or expectations. Considering all these potential pitfalls, it is quite surprising that the financial markets on TradeSports are so robust. The TradeSports data demonstrate that financial markets that merely aggregate information can be quite efficient even at small scales and in the presence of irrational behavior. Specifically, financial markets for *short-term contracts based on objective fundamentals* (such as the occurrence of a verifiable event–*e.g.*, a particular dividend payment) are relatively efficient despite numerous obstacles to arbitrage.

Appendix
Maximum Likelihood Estimates of the Prelec Weighting Function

Recall that the maximum likelihood estimates of the Prelec function are based on contracts with independent underlying events. Let p_i be the unobserved true probability that the underlying event i occurs. Let s_i denote a dummy variable for event i; so $s_i = 1$ when event i "succeeds" and $s_i = 0$ when event i "fails" to occur. The total number of contracts is given by n.

Assuming contract events are independent, the likelihood function L for contract realizations is proportional to:

$$L \propto \prod_{i=1}^{n} p_i^{s_i}(1 - p_i)^{1-s_i} \tag{15}$$

Note that market equilibrium predicts specific values for p_i in the equation above. As argued in section 3, the true probabilities p_i should approximately satisfy:

$$w_i \equiv w(p_i) \Rightarrow p_i = w^{-1}(w_i) \tag{16}$$

The form of the probability weighting function is determined by the theory of interest to the researcher. In this study, I estimate the Prelec weighting function in equation (12), which nests the EU and prospect theory models. Inverting the Prelec function to solve for p_i gives:

$$p_i = \exp[-(-\ln w_i)^{\frac{1}{\alpha}}] = w^{-1}(w_i, \alpha) \tag{17}$$

Next, I take the logarithm of the likelihood function and substitute the inverted Prelec weighting function above to obtain the maximization problem:

$$\max_{\alpha,\beta} \ln L(\alpha, s_1, ..., s_n)$$

$$\Rightarrow \max_{\alpha} \sum_{i=1}^{n} s_i[-(-\ln w_i)^{\frac{1}{\alpha}}] + (1 - s_i)\ln[1 - \exp[-(-\ln w_i)^{\frac{1}{\alpha}}]] \tag{18}$$

The likelihood function satisfies all the regularity conditions required for the maximum likelihood estimators to be consistent, asymptotically normal and asymptotically efficient. Although the first-order conditions for the maximization problem in (18) have no analytic solution for α, standard numerical optimization techniques approximate the solution with arbitrary precision. Let $\widehat{\alpha}$ be approximate solution to the maximization problem in (18).

To obtain an estimate of the standard errors for $\widehat{\alpha}$, I numerically evaluate the observed second derivative of the log likelihood function at the maximum likelihood estimator. This computation provides a consistent estimate of the (negative) information matrix, which can be inverted to yield an estimate of the asymptotic variance of the maximum likelihood estimators:

$$I(\widehat{\alpha}) \equiv -\frac{\partial^2 L}{\partial \alpha^2}(\widehat{\alpha}) \tag{19}$$

References

[1] Avery, Christopher, and Judith Chevalier, 1999, Identifying Investor Sentiment from Price Paths: The Case of Football Betting, *Journal of Business* 72, 493-521.

[2] Berg, Joyce, Robert Forsythe, and Thomas Rietz, 1997, What Makes Markets Predict Well? Evidence from the Iowa Electronic Markets, in Wulf Albers and Martin Strobel eds.: *Understanding Strategic Interaction: Essays in Honor of Reinhard Selten* (Springer: Heidelberg and New York).

[3] Black, Fischer, 1972, Capital Market Equilibrium with Restricted Borrowing, *Journal of Business* 45, 444-55.

[4] Brav, Alon, George M. Constantinides and Christopher C. Geczy, 2002, Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence, *Journal of Political Economy* 110, 793-824.

[5] Camerer, Colin F. and Teck-Hua Ho, 1994, Violations of the Betweenness Axiom and Nonlinearity in Probability, *Journal of Risk and Uncertainty* 8, 167-96.

[6] Cason, Timothy N. and Daniel Friedman, 1996, Price Formation in Double Auction Markets, *Journal of Economic Dynamics and Control* 20, 1307-37.

[7] Cleveland, William S., 1979, Robust Locally Weighted Regression and Smoothing Scatterplots, *Journal of the American Statistical Association* 74, 829-36.

[8] De Bondt, Werner F.M., Richard H. Thaler, 1985. Does the Stock Market Overreact? *Journal of Finance* 40, 793-805.

[9] DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990, Noise Trader Risk in Financial Markets, *Journal of Political Economy* 98, 703-38.

[10] Gode, Dhananjay K. and Shyam Sunder, 1993, Allocative Efficiency of Markets with Zero-Intelligence Traders: Markets as a Partial Substitute for Individual Rationality, *Journal of Political Economy* 101, 119-37.

[11] Fama, Eugene F., 1970, Efficient Capital Markets: A Review of Theory and Empirical Work, *Journal of Finance* 25, 383-417.

[12] Fama, Eugene F., 1991, Efficient Capital Markets: II, *Journal of Finance* 46, 1575-617.

[13] Friedman, Daniel and Joseph Ostroy, 1995, Competitivity in Auction Markets: An Experimental and Theoretical Investigation, *Economic Journal* 105, 22-53.

[14] Gjerstad, Steven and John Dickhaut, 1998, Price Formation in Double Auctions, *Games & Economic Behavior* 22, 1-29.

[15] Gonzalez, Richard and George Wu, 1999, On the Shape of the Probability Weighting Function, *Cognitive Psychology* 38, 129-166.

[16] Grossman, Sanford J. and Joseph E. Stiglitz, 1980, On the Impossibility of Informationally Efficient Markets, *American Economic Review* 70, 393-408.

[17] Hausch, Donald B. and William T. Ziemba, 1990, Arbitrage Strategies for Cross-Track Betting on Major Horse Races, *Journal of Business* 63, 61-78.

[18] Huberman, Gur and Tomer Regev, 2002, Contagious Speculation and a Cure for Cancer: A Nonevent that Made Stock Prices Soar, *Journal of Finance* 56, 387-96.

[19] Jullien, Bruno and Bernard Salanie, 2000, Estimating Preferences under Risk: The Case of Racetrack Bettors, *Journal of Political Economy*, 108, 503-30.

[20] Kahneman, Daniel and Amos Tversky, 1979,. Prospect Theory: An Analysis of Decision under Risk, *Econometrica* 47, 263-91.

[21] Künsch, Hans R., 1989, The Jackknife and the Bootstrap for General Stationary Observations, *The Annals of Statistics* 17, 1217-41.

[22] Lattimore, Pamela K.,Joanna R. Baker, and Ann D. Witte, 1992, The Influence of Probability on Risky Choice: A Parametric Examination, *Journal of Economic Behavior and Organization* 17, 377-400.

[23] Noussair, Charles, Stephane Robin, and Bernard Ruffieux, 1998, The Effect of Transactions Costs on Double Auction Markets, *Journal of Economic Behavior and Organization* 36, 221-33.

[24] Prelec, Drazen, 1998, The Probability Weighting Function, *Econometrica* 66, 497-527.

[25] Prelec, Drazen, 2000, Compound Invariant Weighting Functions in Prospect Theory, in Daniel Kahneman and Amos Tversky eds.: *Choices, Values, and Frames* (Cambridge University Press, Cambridge).

[26] Roll, Richard, 1977, A Critique of the Asset Pricing Theory's Tests: Part I: On Past and Potential Testability of the Theory, *Journal of Financial Economics* 4, 129-76.

[27] Ross, Stephen A., 1976, The Arbitrage Theory of Capital Asset Pricing, *Journal of Economic Theory* 13, 341-60.

[28] Ross, Stephen A., 1978, Mutual Fund Separation in Financial Theory—The Separating Distributions, *Journal of Economic Theory* 17, 254-86.

[29] Satterthwaite, Mark A. and Steven R. Williams, 2002, The Optimality of a Simple Market Mechanism, *Econometrica* 70, 1841-63.

[30] Shleifer, Andrei and Robert W. Vishny, 1997, The Limits of Arbitrage, *Journal of Finance* 52, 35-55.

[31] Thaler, Richard H. and William T. Ziemba, 1988, Parimutuel Betting Markets: Racetracks and Lotteries, *Journal of Economic Perspectives* 2, 161-74.

[32] Tversky, Amos and Daniel Kahneman, 1992, Advances in Prospect Theory: Cumulative Representation of Uncertainty, *Journal of Risk and Uncertainty* 5, 297-323.

[33] Woodland, Linda M. and Bill M. Woodland, 1994, Market Efficiency and the Favorite-Longshot Bias: The Baseball Betting Market, *Journal of Finance* 49, 269-79.

[34] Woodland, Linda M. and Bill M. Woodland, 2001, Market Efficiency and Profitable Wagering in the National Hockey League: Can Bettors Score on Longshots? *Southern Economic Journal* 67, 983-95.

[35] Woodland, Linda M. and Bill M. Woodland, 2003, The Reverse Favourite-Longshot Bias and Market Efficiency in Major League Baseball: An Update, *Bulletin of Economic Research* 55, 113-23.

[36] Wu, George and Richard Gonzalez, 1996, Curvature of the Probability Weighting Function, *Management Science* 42, 1676-90.

[37] Ziemba, William T. and Donald B. Hausch, 1986, *Betting at the Racetrack*, Vancouver and Los Angeles: Dr. Z. Investments, Inc.