THE EFFECT OF MARKET STRUCTURE ON PRICING EFFICIENCY: AN EMPIRICAL STUDY OF SPORTS BETTING MARKETS

TYLER HOFMANN PATLA

Dr. Edward Tower, faculty advisor

Honors thesis submitted in partial fulfillment of the requirements for Graduation with Distinction in Economics in Trinity College of Duke University

Duke University
Durham, North Carolina
2007

Acknowledgements

I could not have completed this thesis without the assistance, suggestions, and advisement of a number of generous people: Dr. Edward Tower, my advisor, for his time, comments and suggestions; Dr. Steven Levitt and Erik Snowberg for their expertise and valuable insight; Dr. Alper Ozgit and Alex Hetherington for sharing their experiences with similar studies; the faculty and my fellow students of the honors thesis seminar for their reviews and suggestions; Dr. Emma Rasiel for her unceasing assistance in times of need. Finally, I'd like to thank Dr. Cathy J. Patla and Raymond S. Patla not only for their proofreading and commentary but for their encouragement, for their dedication, and for providing an academic foundation that has allowed me to achieve my scholastic and personal successes.

Abstract

Many studies have identified systematic pricing inefficiencies in both traditional, dealer-based sportsbooks and in new, order-driven sports betting exchanges. Some studies, most notably Levitt (2004), have found that sportsbooks adjust their lines to further exploit bettor biases in order to maximize profits. In an attempt to contrast the relative efficiency of the different markets, I effectively isolate multiple statistics that predict pricing inefficiencies in an OLS regression model. Initial data from the 2006 Major League Baseball season suggests that sports exchanges are actually *less* efficient than traditional sportsbooks, although high standard errors and data inconsistency restrict definitive conclusions from being made.

I. Introduction

In many ways, sports betting is a simplified analogue to trading in financial markets. Like many financial derivatives markets, sports betting is a zero-sum game in which participants with heterogeneous beliefs seek to profit through trading as uncertainty is resolved over time (Levitt, 2004). Each system's structure has also developed similarly with technological advances. Until the early 1990s, financial markets relied on quote-driven dealer markets. As the decade progressed and technology improved, electronic order-driven systems became more prevalent, and a hybrid market structure developed (Ozgit, 2006). Similarly, online sports betting exchanges have recently emerged in an industry traditionally dominated by market-making sportsbooks to create an analogous dual structure.

A number of studies have investigated sports betting markets as a simplified investment marketplace good for testing market efficiency. Many have confirmed that the market contains otherwise sophisticated bettors with inherent biases. Few, most notably Levitt (2004), comment on the effect of market structure on these inefficiencies, and none to my knowledge have contrasted the results of different structures on tests of market efficiency. The recent emergence of sports betting exchanges that mirror the development of financial markets provides a great opportunity to focus on the effect of market structure on efficiency, and that becomes the focus of this piece.

Sports betting markets provide an excellent opportunity for testing market efficiency. In financial markets, most securities are infinitely lived, with no termination point at which the value of the asset is determined. Sports betting markets, however, have a fixed range of payoffs and outcomes that are decided with certainty at the conclusion of the underlying event (Gray & Gray, 1997). With the maturation of financial markets, significant study has been directed towards the efficient market hypothesis. If market prices fully reflect all available information, every bet and its compliment should have equal inherent expectation. If not, then inefficiencies could be statistically extracted and strategies could be systematically employed that could outperform random selection. Market rationality can be directly tested by comparing observed prices (betting odds) with actual outcomes to see if an asset's price is truly an unbiased predictor of its future value (Gandar, Zuber, O'Brian, and Russo, 1988).

Despite their similar structures, financial markets and gambling markets have a fundamental difference in participant preferences. Financial markets primarily exist as an investment vehicle, whereas sports betting markets exist predominantly because of the entertainment value they provide. While most investors attempt to assess the value of a security objectively, the popularity of sports is due in large part to the emotional attachment fans have to their favorite teams. Vernón (2003) notes that additional utility is derived from aligning personal and financial preferences, and these selection biases by otherwise rational market participants create a fundamental justification for the irrationalities we observe—fans may be willing to forego a few dollars in expectation from their bet if it allows them to cheer for their preferred team more heartily. The entertainment-based function of gambling markets allows a number of financially irrational biases to become more prominent than in financial markets. These inefficiencies will provide means for comparison between differing market structures.

A myriad of economists have discovered systematic inefficiencies in sports betting prices due to public bias. Levitt (2004), in fact, finds that sportsbooks often maximize expected profits by taking a disproportionate amount of bets on one team to exploit public bias. This study attempts to determine whether these documented inefficiencies are solely the result of a biased public, or whether bookmakers' market fading exacerbates the inefficiencies and suboptimal behavior.

To test the effect of market structure on efficiency, I use historical sportsbook betting lines and outcomes of Major League Baseball (MLB) games from 1999 to 2006. I first attempt to quantify market inefficiencies in the traditional sportsbook market using a least squares regression model by regressing the game outcomes on both the closing price of the game and publicly available statistics and information. I then apply this model to data from the sports-betting exchanges to investigate whether exchanges arrive at more accurate prices.

Much of the previous study in this field has hypothesized profitable betting strategies and tested them against historical data. By developing a more robust ordinary least squares (OLS) regression model, I am able to test multiple sources of inefficiency simultaneously and contrast my results between bookmaker markets and exchanges. Are the inefficiencies found using sportsbook data still statistically significant when run on new sports exchange data? Do betting exchange prices incorporate more publicly available information? Does the new order-driven, market-clearing format improve market efficiency, or do speculating investors cause the market to arrive at the same price regardless?

To do this, I describe the different sports betting market structures from the perspective of this study and summarize previous analysis in Section II. Section III incorporates previous studies to develop a theoretical framework in an effort to quantify the present inefficiencies and contrast the results from different betting markets. Section IV describes the data, Section V details the results of my analysis, and Section VI concludes.

II. The Sports Betting Market

2.1 Market Structure

American horse racing and some European sporting events are based on the parimutuel betting system. In parimutuel betting, all the wagering on a certain outcome is combined into a common prize pool. After wagering is closed, the house takes a percentage cut and evenly distributes the remaining money to the winning bettors. As a result, bettors are not guaranteed the odds or price they are receiving when they place a bet—final prices are determined after betting is closed, and bettors instead bet based on the price approximated by the wagers already made. While bias in these markets is well documented, and while it is an interesting study because the zero-sum nature of the event arrives at market-clearing prices, it does not closely resemble the fixed-price system of financial markets and is not the focus of this study.

In American team sports, the fixed-odds wagering system is predominant. With fixed odds, the bettor is guaranteed the spot price when he or she bets. Future price movements do not affect the returns the bettor will receive if the bet wins. This format is more useful within our context because it more closely resembles financial investment markets, in which all transactions settle at the price agreed upon when the transaction is made. Additionally, a bettor can bet against a team by betting for the other team, equivalent to short-selling a security in financial markets.

Some American betting markets, like football and basketball, artificially adjust the score with a “point spread,” where favorites must win by at least the posted margin to win, and payouts are constant on either side. Since baseball’s low scores make that method impractical, most bets on baseball straight bets on a team to win (called "moneylines") with payouts based on the likelihood of that team winning. They are quoted as follows:

New York Mets	+ 120
Atlanta Braves	- 140

When betting on the Mets at +120, one would wager $100 to win $120 and receive $220 (the original bet plus the winning price). To bet on the Braves at -140, one would wager $140 to win $100 and receive $240. Notice that if one wagered one “unit” on each side, one would pay $240 ($100 on the Mets and $140 on the Braves) to receive either $240 or $220. This “round trip” cost represents the commission, or "vig," which is equivalent to bid-offer spreads in financial markets. The bookmaker earns this theoretical charge for the service of providing the bettor with the market and taking the opposite position of the bettor. In fact, the bookmaker-bettor relationship very closely resembles the dealer-broker relationship in financial markets: the bookmaker creates prices and provides liquidity to the bettor in exchange for the theoretical commission contained in the spread. Spreads are also much larger than financial markets, usually 5%

to 20% of the asset's price, because of the increased difficulty in accurately valuing the asset, i.e., the probability of the bet winning.

Sportsbooks create these prices by estimating a "fair price" that represents the actual value of the bet, one which in the long term a bettor would expect to break even. For example, the bookmaker might find the favorite to be -150 and the underdog +150. From here, he will apply the spread so that the customer is paying extra to bet the game: for example, -160/+140. Instead of centering the posted prices around the fair price, the book may choose to "fade" its lines, that is, to adjust the betting line one way or the other in order to achieve a desired bet distribution. For example, if the bookmaker expected most people to bet on the favorite at -160/+140, he might offer -165/+145 instead, so that those who want to bet on the favorite receive a worse price. Sportsbook fading creates an inaccurate implied "fair price," the estimated price without the spread: -155 here instead of -150. Fading is a potential source of market pricing inefficiency—this study attempts to determine whether sportsbooks fade their markets and its affect on overall market efficiency.

With the advent of the internet and online sportsbooks have come online betting exchanges. Like electronic financial exchanges on which investors' bids and offers are matched, bettors can see all the open orders on a game and can choose to match an existing offer or post their own order for the odds they would like to receive. These orders are matched by the exchange, and the bettors receive opposite positions on the game—the exchange does not participate in the action and is not exposed to risk. This market is unique in that it provides a market-clearing structure for fixed-odds betting, and it eliminates any market fading professional oddsmakers may employ. Additionally, because customers can submit a bid for a bet at their desired odds, the orders generate tighter markets with smaller spreads, minimizing the effective "vig" that bettors pay. Whether these tighter markets also more accurately predict future outcomes than traditional sportsbook lines is the question that this study attempts to answer.

2.2 Literature Review

A number of studies, described below, have detailed the structures of different betting markets and have juxtaposed them to investment markets. Many more have aimed to reject market efficiency in various sports by attempting to find profitable betting strategies. Very few have created regression models or incorporated non-financial utility, and none to date have contrasted the results over different market structures. This paper will aim to bridge that gap in tests of market efficiency over varying market structures.

A majority of gambling market efficiency studies have been focused on horse racing markets. Thaler and Ziemba (1988), Sauer (1998), and Snowberg and Wolfers (2005, 2006) are some of many studies that confirm considerable "favorite/longshot bias," in which returns for bets on favorites are significantly higher than returns for bets on longshots. Snowberg and Wolfers (2006) determine that this bias is a result of bettors that systematically overestimate the probability that a longshot will win the race because of an inability to discern between small and tiny probabilities. This results in a disproportionate number of bets on longshot horses and a resultant poorer return due to the parimutuel betting format. While these results cannot be directly applied to fixed-odds betting and financial markets because of the differing structures, it is the most prominent empirically-shown example of irrational bettor behavior that gives credence to future study and analysis.

Additionally, Gandar, Dare, Brown, and Zuber (1998) determine that the closing line on an event is a significantly better predictor than the opening line of the future outcome, implying that market activity can serve to increase the efficiency of the market by aiding price discovery. Price takers can then be seen as sophisticated market participants whose preferences deserve further study.

In team-based, fixed-odds sports betting markets, many studies have found a reverse of the favorite-longshot bias of horse racing. Previous studies of the MLB market by Woodland and Woodland (1994) and the National Hockey League by Woodland and Woodland (2001) state that favorites tend to be overvalued and have less expectation than playing underdogs. Borghesi (2004) concludes that due to the effects of weather late in the National Football League season, home underdogs do significantly better than betting odds would predict. Additionally, they found that the markets did not increase in efficiency over the course of their study, implying that these systematic biases are not being exploited to the extent that the inefficiency disappears from the market.

The deluge of "expert" opinions offered in the media, public overreactions to news events or pricing moves, and systematic bettor biases could all cause the actual market price to differ from the efficient one (Gandar, et al., 1988). For example, Gilovich, Vallone, and Tversky (1985) and Camerer (1989) realize that the public significantly overestimates performance trends and will tend to overvalue teams on winning streaks (called the “hot hand” theory). Vernón (2003) observes that in the NCAA basketball tournament, bettors systematically overvalued underdogs in the tournament because of the exaggeration by the media of underdogs’ victories and overall unpredictability of the tournament, epitomized by the nickname "March Madness." Also, Dixon and Coles (1997) are able to use Poisson distributions to more accurately predict European soccer games than local betting markets.

Nearly all studies to date have openly assumed that the mid-market odds or point spread was a market clearing price. Levitt (2004) finds that traditional bookmakers, however, do not attempt to balance their books and eliminate their risk, but instead exploit the view of the betting public to "fade" the markets on these games. A bookmaker fades his market by adjusting his spread around the fair value to maximize his expected profit based on what he knows about his customer. Creating a fair price or market clearing price would each produce an equal expected profit on the game, but, depending on bettor preferences, a price that achieves greater maximum profit may exist between them. Levitt found that in over half the games in his data set, two thirds or more of the wagers were placed on one team. He concludes that bookmakers are systematically better than the public at both predicting game outcomes and predicting public betting patterns, and they adjust their markets accordingly to maximize profit. In his sample, line fading was able to increase profit by almost 25% over using equilibrium prices.

This suggests that any systematic biases in bettor preferences are *exacerbated* by the bookmaker structure, and that the structure of these markets furthers the inefficiency. Additionally, Strumpf (2003) details team-specific biases in which local bookmakers would consistently give worse odds on popular teams. Specifically, he found a group of New York bookies that would offer substantially worse odds to bettors they knew were Yankee fans.

Ozgit (2006) was the first to develop an analysis of the structure of online betting exchanges. He concludes that although their structure results in smaller spreads and greater returns than sportsbooks for a small bettor, current markets the lack the depth to provide the liquidity necessary to replace sportsbooks. Hetherington (2006) is able to isolate arbitrage opportunities in the NFL betting markets of the exchange TradeSports.com using backtesting methods similar to those used by Woodland and Woodland (1994, 2001). Neither investigates the accuracy and efficiency in pricing on the exchange *relative* to the sportsbooks, however, and I hope to build my results on their empirical foundation.

Using Levitt's and Strumpf's findings that bookmakers often don't balance their books and are consistently able to assess outcome probabilities more accurately than the betting public, I attempt to determine whether this type of market fading increases the effect of selection bias inherent in the betting public. I contrast this traditional format to sports betting exchanges, which balances order flow, to determine the effect of market structure on efficiency.

Although many studies have investigated the efficiency of sportsbooks and others the efficiency of sports exchanges, none have contrasted them in the context of the same model and data set. Most studies have hypothesized strategies that might create excess returns and tested them, using the discovery of a profitable strategy as evidence for inefficiency. A few, like Golec and Tamarkin (1991), use an OLS regression to determine whether prices include all publicly available information. I aim to find a regression model that can then be used to compare disparate data sets in an effort to assess the efficiency of the new exchanges within the context of traditional sportsbooks.

III. Theoretical Framework

Eugene Fama (1970) defines an efficient market as one in which all publicly available information is incorporated in security prices. In a weak-form efficient market, the current price is the best, unbiased estimate of the value of a security; in a semi-strong form efficient market, all publicly available information is incorporated in the price, and there must be no fundamental analysis that can create excess returns. This study will test for both weak-form and semi-strong form efficiency in the market, using historical prices and outcomes but also incorporating publicly-available statistical data.

In a perfectly efficient market, the price will be an unbiased predictor of the value of a bet. This price will not necessarily be perfectly accurate in a given situation, but there will be no systematic inconsistencies between the "true" probability and the consensus market price:

$$\pi = \beta_0 + \beta_1 p + \varepsilon$$

where π is the true probability of the event occurring, p is the probability implied by the price of the security, and ε incorporates all additional information not contained in the price.

Under the assumption that the price is an unbiased predictor of the asset's value, the null hypothesis for efficiency is $\beta_0 = 0$, $\beta_1 = 1$, and ε cannot be present systematically. In an environment with weak-form efficiency, all information regarding the value of the security is contained in the price. No other available information should have a statistically significant effect, i.e., the coefficient of any additional variables should be zero as well.

If we assume market participants are rational, there must be justification for the inefficiencies that have been statistically identified in previous studies. Vernón (2003) identifies personal preference as a driver for market prices, where the market price is a sum of the true value of the asset as well as the net utility the bettor receives from betting that asset:

$$p = \pi + ne$$

where *ne* represents the net effects of public bias on the price of the asset. By incorporating this into the efficiency model, we can attempt to quantify these net effects:

$$\pi = \beta_0 + \beta_1 p + ne + \varepsilon$$

To test for efficiency in the different markets, we will first assume them to be efficient by assuming $\pi = p$, so creating a new variable for the residual, *r:*

$$r = \pi - p$$

For any given game, *r* represents the net return of a bet on that game—if the bet loses, the bettor's return is -*p*, if the bet wins, the net return is 1-*p*. If any publicly available statistics are correlated with *r* and can be used to predict *r*, they can be used to predict excess profits are losses and to reject semi-strong form efficiency. We can incorporate the various behavioral biases found by previous studies to incorporate into this model. If the market is truly efficient, all these additional parameters will have coefficients of zero.

I use the following framework to identify and quantify sources of inefficiency:

$$r = \beta_0 + \alpha \mathrm{P} + \gamma \mathrm{H} + \phi \mathrm{T}$$

Price-based effects, **P**, will attempt to find sources of inefficiency that can be determined by the market price of the security. Snowberg and Wolfers (2006) discover that bettors have difficulty accurately predicting small probabilities, so the effect of odds on returns could be nonlinear—I add a p^2 term to the regression to capture this effect. Woodland and Woodland (1994, 2001) and Levitt (2004) find that favorites, particularly home favorites, are overvalued. Since all games are from the home team's perspective, I add a dummy variable equal to 1 if the team is a favorite could capture this effect.

The effect of historical results, **H**, could cause markets to arrive at inefficient prices. Gilovich, et al. (1985) and Camerer (1989) find that the public overvalues recent performance, believing in the "hot hand," so I will include the current winning or losing streak as well as the results of the previous game, ten games, and season will attempt to capture these effects.

Finally, team specific and popularity-based effects, **T**, could affect market prices. Levitt (2004) and Strumpf (2003) both find that bookmakers tend to exploit fans of specific popular teams, so incorporating quantifiable measures of team popularity, namely *Forbes Magazine* franchise value and city population, may yield significant results.

The final function may be presented as:

$$r = \beta_0 + \beta_1 fav + \beta_2 p^2 + \beta_3 pctdif + \beta_4 last + \beta_5 ws + \beta_6 forbes + \beta_7 pop + \varepsilon$$

where *fav* is a dummy variable if the team is a favorite. To attempt to center most variables around zero (so a test of their significance is more intuitive), other variables will be determined as the difference between the home team and away team in that statistic. *Pctdif* is the difference between the teams' winning percentages from the previous season, *last* represents the result of the teams' last game and is equal to 1 for a win and -1 for a loss, *ws* represents current the winning streak (or negative for losing streak), and *forbes* and *pop* represent the difference between the home team and away team in franchise value and city population, respectively.

Section V will detail the construction of a reliable model to incorporate the systematic inefficiencies. Then, the model will contrast the results from sportsbook prices and sports exchange prices from the same set of games. Doing so will help to discern whether the sports exchange format is susceptible to the same biases and inefficiencies that sportsbook markets imply.

One must realize, however, that the markets being contrasted are not independent—market participants on the sports betting exchange likely have accounts with at least one traditional sportsbook. As a result, markets are bound by "no arbitrage" restrictions—if the markets misalign to the extent that one can bet each side in a different market to ensure risk-free profit, arbitrageurs will execute the trades, pushing prices back toward a common value. Also, because the top linesmakers still work for the most popular sportsbooks, exchanges market-makers tend to mimic any moves made in the corresponding sportsbooks. These factors add to the inter-market price interactions from which I am trying to extract structure-induced differences. I hypothesize, then, that if we consider a spectrum of price inefficiency, this situation will form an upper bound of inefficiency on the sportsbook price and a lower bound on the fair price. Sports exchange prices will converge somewhere in between, with market interactions putting pressure on the sports exchange price toward the upper bound. I will attempt to find whether there is an appreciable difference between the sports exchange levels of inefficiency and the sportsbook's upper bound.

IV. Data

4.1 Sportsbook Data

This study uses moneylines—odds relying solely on the team's ability to win, not the winning margin—for Major League Baseball from 1999-2006. Major League Baseball was the strongest choice of a dataset because of the high frequency of games and the relative consistency of market participants (as opposed to more popular betting sports like football, which have a fluctuating number of casual bettors). In fact, Michael "Roxy" Roxborough, president of Las Vegas Sports Consultants and a premier linemaker, states that those that wager on baseball are widely considered the most knowledgeable and sophisticated of all sports bettors (Woodland & Woodland, 1994). Studying the actions of the most sophisticated market participants would most closely parallel the participants in financial markets. Additionally, the moneyline bets offered on baseball games provide a wide range of prices, creating an additional degree of freedom not contained in point spreads, which are usually traded at a constant price.

To obtain the markets created by bookmakers, I compiled historical closing prices for the moneylines of MLB games from Covers.com, a third-party database for sports gamblers. By using a JavaScript code to retrieve the data from the website, I extracted the home team, visiting team, date, starting pitchers for each team, moneyline, final score, and "over/under" line, which is a bet on whether the total runs scored by both teams will be greater or less than the given number. The data set includes 19,418 regular season games from the 1997 through 2006 seasons. This figure is slightly less than 19,440, the total if all 30 teams played 162 games for each of the eight years. This is because 15 games over the past eight seasons were rained out and never rescheduled, and seven more ended in ties and were not completed.

All games are analyzed from the home team's perspective. The prices are quoted as the odds one would get to bet on either team (e.g., +140 for the visiting team and -160 for the home team). From this I extract a mid-market, "implied fair" price at which a wager on either team would have equal expectation (-150 for the home team). I convert these odds to the implied probability of the home team's chance of winning, or the percentage of time the team would have to win for the bet to break even, which I will refer to as "probabilistic form" (60)[1].

4.2 Sports Exchange Data

In order to contrast results from sportsbooks with sports exchanges, I purchased trading logs from TradeSports.com, one of the internet's premier and most active online betting exchanges. On the website, customers can post orders (or accept a posted order) on the outcome of a game by buying or selling contracts. (Buying the contract is the effect of betting on the posted team; selling the contract is betting against the posted team and for the opponent). The exchange provided me with each trade made on a MLB game during the 2006 season and consists of the trade date, trade time, contract (game) traded, contract size, and price. The contracts traded on TradeSports.com mature to 100 for the winning side, so it is already in probabilistic form and analogous to the prices I have calculated for the sportsbook format. There were 190,615 moneyline transactions on the 2430 games in the 2006 season, for an average of 78 trades per game.

In order to find a "closing" price implied by the exchange that is analogous to the closing market of the sportsbook, I isolated trades that occurred in the last half-hour

[1] Please see the appendix for the derivation of this calculation.

before a game. By using game start times obtained from *Doc's Sports Service*, a sports information website, a VBA script filtered the trades based on their transaction time relative to the game start. In my opinion, using the last half-hour of data provided the greatest compromise between having enough volume to ensure accuracy and assuring that the price closely represented the value of the security at market close, minimizing the effect of price movement by speculators (or news events) over the course of the day. If there Of the 190,615 moneyline transactions, 17,565 took place in the half hour before the game, or 7.2 per game. In 15 cases, there were no trades in the last half hour before the first pitch, so the last transaction was used as long as the price did not differ greatly from the sportsbook market. Comparison to the sportsbook market was used to check for errors—in the 96 games in which the price differed by more than two points from the sportsbook market, I manually investigated the transaction log to interpret the probable start times. In most cases, the start time had been moved up from the posted time, so the script was capturing inter-game betting, not pre-game betting. These times were fixed manually.

Table 1 summarizes the data retrieved from the sportsbooks and sports exchange. It indicates the number of observations and average prices recorded for home teams. Both the sportsbook and exchange prices, when aggregated over the entire data set, are very close to observed results in recent years. When directly comparing the sportsbook and exchange data sets, only games from the 2006 season will be used so that each data set is based upon the same set of games.

Table 1
Data Summary

	Sportsbook	Exchange
Years of data	8	1
Games	19418	2430
Total prices (or transactions)	19418	190615
Average price (for home team)	0.5321	0.5294
Winning Percentage	0.5320	0.5290

4.3 External Data Sources

In addition to the primary data relevant to sports betting platforms, I obtained a number of exogenous data sources to help quantify the biases found by Levitt (2004) and Strumpf (2003). Because isolating team fixed effects may provide too many degrees of freedom for the regression to remain meaningful, I attempt to find a proxy that numerically represents the popularity of a team. *Forbes Magazine* regularly publishes Forbes Franchise Values for major sports organizations, and is publicly available on Forbes.com. This report is an aggregate measure of a team's metropolitan population, attendance, television sponsorship, and merchandise sales, and is a potentially robust numerical measure of a team's popularity. Since a team's valuation for a given year is calculated from statistics for the previous year, it is an accurate measure of current popularity and expectation for team performance without fear of reverse causality, i.e., unexpectedly high performance causing a team's value to increase. In addition, I have added team attendance, payroll, and team revenue to attempt to isolate measures of popularity. All of the above were compiled on *RodneyFort.com*, a sports economics website.

4.4 Dataset Weaknesses

Unfortunately, the secretive world of sportsbook operations and the relative infancy of betting exchanges does not allow for perfect data acquisition. I must trust Covers.com's data to be representative of real sportsbooks, although I have spot-checked over 50 games by comparing the price posted on Covers.com to a number of major online sportsbooks. I found Covers.com's data to represent the consensus market price on all occasions, and other studies, namely Wolfers (2006), have used this data source in the past with success.

Additionally, the differing format of the data could weaken any conclusions. While sportsbook data provides a snapshot of the closing market, TradeSports.com provides trading data that I must aggregate over the final period of trading, which makes it subject to capture price movements close to gametime instead of the final price. Also, since Tradesports.com provided matched transactions, not posted markets, it is not perfectly analogous to sportsbook data and may not be representative of the true implied price of the posted market. Because the order-driven format is market clearing, however, this source of inaccuracy should be minimal in practice, but there is no practical way to test this under current market conditions. Also, the relative infancy of exchanges and high cost of obtaining transaction logs prevented me from obtaining more than one year of data, and, compared with many years of sportsbook data, may cause the results to be highly sensitive.

Other weaknesses arise because, as discussed in Section III, the markets are neither independent nor isolated. The market participants are not exclusive to one site nor randomly distributed; since sports exchanges tend to be slightly less intuitive to the mainstream user than more familiar traditional sportsbooks, it is possible that exchange users, on the whole, are more sophisticated, discriminating customers. This may make it difficult to discern whether empirical differences are a result of market structure or of participant demography. A controlled, isolated experiment in which randomly assigned bettors act only in one market and do not see the other market would be ideal in minimizing potential exogenous effects. The dual, interactive market structure, however, provides a more complex framework that is more representative of financial market structure.

The imperfections in the data, the secretive nature of sports odds creation, and the relative infancy of sports exchanges, cause the best dataset available at this time to be far from optimal. Hopefully the methods developed in this study can be later applied to a more robust and consistent dataset once exchanges mature and posted markets can be observed over time. By contrasting traditional sportsbook prices obtained from Covers.com and sports exchange trading data from TradesSports.com, I will examine the effect of additional, external data from various publicly available sources to identify market inefficiencies.

V. Empirical Analysis

5.1. Betting Strategies: The Reverse Favorite-Longshot Bias

Many studies, including Woodland and Woodland (1994, 2001) and Borghesi (2004) discover that favorites tend to be overvalued, particularly away favorites, making home underdogs disproportionately good bets. Levitt (2004) finds that bookmakers exploit this bettor tendency, fading their markets to offer poorer odds to those bettors. Since the sportsbook dataset was created by bookmakers, this relationship should persist, but how does it compare to the sports exchange data?

Before creating a more complex regression model, I will first repeat the works of previous authors by finding a simple system that can be systematically applied to achieve excess returns in the market. Levitt (2004) includes a table analyzing net returns regarding the home underdog bias that inspired my creation of Table 2, which shows the net return on investment of purchasing one contract for each of the home/away

Table 2

The Reverse Favorite-Longshot Bias over various datasets: the net returns of simple betting strategies

	Home Underdog	Home Underdog	Away Favorite	Away Underdog
Sportsbook (all seasons)	-0.67%	1.81%	-1.36%	0.98%
Sportsbook (2006)	-0.03%	5.91%	-4.72%	0.04%
Sports Exchange (2006)	-0.09%	7.72%	-6.12%	0.13%

favorite/underdog combinations.[2] The results here are consistent with Woodland and Woodland's (1994, 2001) and Levitt's findings—betting exclusively on home underdogs with no other restrictions on every game from 1999-2006 would have yielded a 1.8% return.[3, 4]

The results below the full dataset show the returns for the strategies for sportsbooks in 2006. It was an above-average year for the inefficiencies and the profitability of the home underdog trading strategy (although not an outlier in light of the year-to-year variance these returns have). The sports exchange data, however, showed a *higher* return than sportsbooks, indicating *more* bias in the market and *less* efficient prices. Despite a market-clearing structure and the absence of official sportsbook dealers, sports exchange prices seem to be *more* susceptible to reverse favorite-longshot bias than sportsbook prices, at least for the 2006 dataset. Does this hold, however, for a more complex data regression?

5.2. Building a model

With an original, weak-form efficient model of:

$$\pi = \beta_0 + \beta_1 p + \varepsilon$$

[2] Levitt (2004) incorporated the number of participants who made each bet into his analysis. I treat each game equally, but incorporate the differing prices bettors receive in moneyline transactions.
[3] The return for bets on home underdogs is not the additive inverse of the return on away favorites because the larger nominal value needed to purchase a contract on a favorite causes the magnitude of the net return to be lower for the same nominal return.
[4] The 1.8% return, however, is before commissions, which are usually 2.25% to 4.5% for baseball, depending on the sportsbook. Therefore, this strategy would not have been inherently profitable at a traditional sportsbook over the dataset.

where p, the price, is the independent variable that should estimate π, the outcome of the event. Table 3 shows the results of this regression. As expected, β_0 is estimated to be greater than 0 and β_1 less than 1, although the joint F-test on $\beta_0 = 0$ and $\beta_1 = 1$ cannot be rejected with high confidence. While not statistically significant, the estimated results are consistent with previous studies, as Woodland and Woodland (1994) and others have found that favorites will win less often than implied by the price, and the regression gives credence to further investigation.

Table 3
Test of weak-form efficiency: 1999-2006 sportsbook data

	Base Model
price	0.9354
	(0.0375)
constant	0.0346
	(0.0205)
F-test	1.48
p-value	0.22

Starting with the model of:

$$r = \beta_0 + \alpha P + \gamma H + \phi T$$

I isolated the most effective statistics for each family of effects. After regressing r against all the data collected, I find the model is best represented by:

$$r = \beta_0 + \beta_1 fav + \beta_2 p^2 + \beta_3 pctdif + \beta_4 forbes + \beta_5 pop$$

The original population and Forbes franchise value figures, however, were not statistically significant. By creating "tiers" that isolated the most and least popular teams, however, this data showed surprising relevance. The *forbes* value is 1 if the home team is in the tip five Forbes franchise values in a given season and -1 if the home team is in the bottom five, with the other twenty teams 0. The *pop* variable is 1 if the home team is in a top ten metropolitan area—New York City (two teams), Los Angeles (two), Chicago (two), Arizona (Phoenix), Houston, Philadelphia, and Toronto—and -1 if the team is in the bottom ten—Cleveland, Kansas City, Atlanta, Cincinnati, Florida (Miami), Minnesota (Minneapolis/St. Paul), Oakland, Pittsburgh, St. Louis, and Tampa Bay (Tampa). The middle ten teams' values are zero. Both variables are then centered at zero, and high popularity is represented with positive observations. Note that both these variables are only used for the *home* team—including the statistic for the visiting team has a relatively insignificant effect.

The results are summarized in Table 4. When looking at each group of price-based factors, historical results-based factors, or team popularity-based factors, we can reject with relatively high confidence the null hypothesis that the coefficient on each group of factors is equal to zero.

Table 4

Base regression model, 1999-2006 sportsbook data

Predictors of residual	
p^2	-0.0819
	(0.0605)
fav	-0.00556
	(0.01187)
pctdif	0.107
	(0.039)
forbes	-0.0120
	(0.0068)
pop	-0.00686
	(0.00470)
Constant	0.0279
	(0.0143)
Test on price-based factors (p^2, *fav*)	
F-statistic	2.87
p-value	0.0567
Test on previous results (*pctdif*)	
t-statistic	2.77
p-value	0.0060
Test on popularity factors (*forbes*, *pop*)	
F-statistic	4.17
p-value	0.0155
Test for efficiency (all factors)	
F-statistic	2.95
p-value	0.0070

The price-based factors, p^2 and *fav*, both have negative coefficients, illustrating the reverse favorite-longshot bias in the dataset that home underdogs are disproportionately undervalued. The F-test on both coefficients being statistically insignificant yields a p-value of 0.0567, allowing us to reject efficiency with greater than 90% confidence. Each coefficient is not significant on its own because of high multicolinearity, but together they illustrate the overvaluation of favorites in the betting markets. Creating dummy variables for price "ranges" (e.g., 45 to 50, 50 to 55) to determine whether specific prices were under- or overvalued yielded some significant results (prices in the 42 to 46 range were particularly undervalued), but they did not illustrate the data as simply and effectively as the p^2 and *fav* combination.

Many derivatives of recent team performance were, surprisingly, statistically insignificant for the dataset, despite their significance in previous studies like Gilovich, et al. (1985) and Camerer (1989). Winning streaks, wins in the last ten games, and a previous game result all showed little to no statistical significance for either the home or away team. However, the *pctdif* statistic, which represents the difference in winning percentage between the home and away team from the previous season, has a coefficient of 0.107, which is statistically significant with 99% confidence. This sign is opposite from what is intuitive—one would expect bettors to overweight a team's lagged performance from the previous season, causing them to overvalue a team that did well the previous year (and causing a negative coefficient). This variable, however, must be indirectly capturing the more recent performance. If the previous *season's* performance is a better indicator of team quality than the previous *week's* performance, bettors may be underweighting team quality relative to recent performance, which is estimated by the *pctdif* metric.

The proxies created to capture team popularity—*Forbes Magazine* franchise value and metropolitan area population—were insignificant for the home team, the visiting team, or the difference between them. Simplifying these statistics into "tiers" and looking solely at the home team, however, produces strikingly strong results. Isolating the top and bottom five teams in Forbes franchise or top and bottom ten metropolitan areas each produced the most statistically significant results. While high multicolinearity causes both coefficients to be statistically insignificant separately, a test on both shows that efficiency can be rejected with over 95% confidence. A team that is top five in franchise value and top ten in metropolitan population, like the Yankees, is predicted to be overvalued by 1.89%, an economically significant figure that could possibly be combined with other advantages to find theoretically profitable wagers, and confirms the biases found in Strumpf (2003) and others.

Finally, a test of overall market efficiency—all coefficients are equal to zero—is rejected with 99% confidence. The model constructed efficiently captures market inefficiencies with the simplicity of a five factor model (as opposed to incorporating team fixed effects or price ranges, which would add many more degrees of freedom). Now, we use the model to contrast sportsbook and sports exchange data from the 2006 season.

4.3. Contrasting Sportsbook and Sports Exchange Results

Using the model constructed in section 4.2:

$$r = \beta_0 + \beta_1 fav + \beta_2 p^2 + \beta_3 pctdif + \beta_4 forbes + \beta_5 pop$$

I regress on the sportsbook and sports exchange data from the 2006 season, and the results are summarized in Table 5.

The coefficients in the two new regressions seem to imply that sports exchanges, despite their market-clearing format and without explicit market-fading dealers, are actually *less* efficient than traditional sportsbooks over the dataset. The coefficients on price-based statistics, p^2 and *fav*, are substantially larger for the sports exchange than for the sportsbook, which agrees with the "home underdog" betting strategy results from section 4.1. Additionally, the coefficient on *pctdif* is larger in the sports exchange data than the sportsbook data, implying that sports exchange participants underweight team quality or previous season's results more than sportsbook participants. The level of price inefficiency, therefore, seems to exceed the sportsbook-based upper bound discussed in

Section III. These differences, however, are not statistically significant, so no definitive conclusions can be made regarding the markets' relative efficiency.

Despite the consistently larger coefficients in the sports exchange sample, the standard errors are too large for both 2006 datasets for confident conclusions to be made regarding the efficiency of the markets. None of the tests that were strongly rejected in the eight-year sportsbook sample can be rejected with any level of confidence in the 2006 dataset.

Table 5

Base regression model on data from 2006 season

	Sportsbook	Exchange
p^2	-0.0186	-0.0440
	(0.1823)	(0.1808)
fav	-0.01987	-0.02689
	(0.03197)	(0.03233)
pctdif	0.0343	0.0518
	(0.1273)	(0.1272)
forbes	-0.0229	-0.0221
	(0.0210)	(0.0210)
pop	-0.00361	-0.00381
	(0.01497)	(0.01410)
Constant	0.0278	0.0414
	(0.0443)	(0.0434)
Test on price-based factors (p^2, *fav*)		
F-statistic	0.43	0.95
p-value	0.6486	0.3882
Test on previous results (*pctdif*)		
t-statistic	0.27	0.41
p-value	0.7870	0.6840
Test on popularity factors (*forbes*, *pop*)		
F-statistic	1.11	1.06
p-value	0.3299	0.3465
Test for efficiency (all factors)		
F-statistic	0.73	1.00
p-value	0.6247	0.4268

Running the regression on each individual year of sportsbook data, as summarized in Table 6, shows that the results found in section 4.2 are highly dependent on a multiyear data sample. One year's games are insufficient to produce standard errors small enough to reject efficiency with any level of confidence.

The higher magnitude of inefficiencies in the sports exchange data, although not statistically significant, may be a product of the differing data sources, as opposed to a result of larger biases in the market. While sportsbooks prices were calculated using an average of the available "buy" and "sell" prices, there are no historical markets available for exchanges—only transaction logs. For exchanges, the price was estimated as a weighted sum of transactions, and may not reflect the available market if one side of the

Table 6

Regressions on sportsbook data from each individual season

Year	1999	2000	2001	2002	2003	2004	2005	2006
p^2	-0.0751	-0.117	0.0172	0.0292	-0.3513	-0.2395	-0.0142	-0.0350
	(0.1755)	(0.1952)	(0.1797)	(0.1509)	(0.1805)	(0.1612)	(0.1753)	(0.1822)
fav	0.0401	-0.0243	-0.0266	-0.0246	0.0172	0.0241	-0.0261	-0.0195
	(0.0341)	(0.0351)	(0.0339)	(0.0334)	(0.0354)	(0.0339)	(0.0334)	0.0320
pctdif	0.0496	-0.0307	0.3034	0.2135	0.2816	0.0724	0.1105	0.0421
	(0.1046)	(0.1225)	(0.1328)	(0.1153)	(0.1131)	(0.0930)	(0.1039)	(0.1271)
forbes	-0.0152	0.0041	-0.0402	-0.0417	-0.027	0.0248	0.0043	-0.0225
	(0.0190)	(0.0192)	(0.0196)	(0.0195)	(0.0194)	(0.0189)	(0.0210)	(0.0210)
pop	-0.0146	-0.0035	0.0160	-0.0060	-0.0107	-0.0232	0.0009	-0.0035
	(0.0125)	(0.0128)	(0.0134)	(0.0134)	(0.0137)	(0.0131)	(0.0149)	(0.0150)
Constant	-0.0183	0.0529	0.0032	0.0122	0.1069	0.0532	0.0199	0.0322
	(0.0405)	(0.0464)	(0.0478)	(0.0367)	(0.0438)	(0.0370)	(0.0425)	(0.0442)
Test on price-based factors (p^2, *fav*)								
F-statistic	0.88	1.35	0.57	0.35	2.51	1.28	0.63	0.51
p-value	0.4132	0.2592	0.5669	0.7032	0.0813	0.2784	0.5326	0.6018
Test on previous results (*pctdif*)								
t-statistic	0.47	-0.25	2.29	1.85	2.49	0.78	1.06	0.33
p-value	0.6350	0.8020	0.0220	0.0640	0.0130	0.4360	0.2870	0.7410
Test on popularity factors (*forbes*, *pop*)								
F-statistic	1.12	0.05	2.19	3.24	2.12	1.86	0.04	1.07
p-value	0.3267	0.9534	0.1122	0.0395	0.1202	0.1557	0.9587	0.3427
Test for efficiency (all factors)								
F-statistic	1.08	0.86	1.42	1.62	2.14	1.05	0.29	0.74
p-value	0.3736	0.5199	0.2023	0.1367	0.0463	0.3932	0.9398	0.6147

market was "hit" more than the other side, i.e., if more transactions were made at the buy price or sell price. If we assume that price-makers (those who post the odds they would like to receive) are more sophisticated than price-takers (those who accept posted odds) and draw upon Levitt's findings that a majority of bets (by price-takers) are placed at the "worse" price, it is possible that more transactions occur on the less efficient price than the more efficient price. These assumptions are reasonable because price-taking is the equivalent to betting at a traditional sportsbook. If more trades were made at the "worse" price, the simple weighted sum of transactions might give an artificially poor price that is not representative of the posted markets. The resulting analysis would deduce greater inefficiencies than a study of true posted markets would imply, and that method of price inference could have produced the given results.

VI. Conclusion

This study attempts to construct a model that isolates inefficiencies in Major League Baseball betting markets and can determine the relative efficiency of markets that differ in structure. The OLS regression model is able to effectively isolate market inefficiencies over an eight year sample of baseball games using publicly-available statistics, and the type and direction of inefficiencies found in the market are in general agreement with those identified in previous studies. While efficiency can be rejected with high confidence using sportsbook data over an eight year period, there is no preliminary evidence that removing the discriminating market-makers from traditional sportsbooks and creating a market-clearing exchange environment yields any greater level of efficiency. This is a result of two weaknesses in the best available data sources used in this study—insufficient observations and inconsistent pricing mechanisms.

Although results were statistically significant over the eight year sample, one year of data proved to be insufficient in rejecting efficiency in either market. The standard errors were too high for anything less than a three- to four-year sample for the factors to be statistically significant. Unfortunately, as sports exchanges are in their relative infancy in the US sports markets, exchanges likely lack the history and liquidity to have sufficient data to test for market efficiency at this time. Additionally, the magnitudes of the coefficients in the 2006 sample suggest that sports exchanges are *less* efficient than traditional sportsbooks, but this is likely the product of the different ways in which market prices were calculated, not an effect of market efficiency itself.

As a result, I conclude that although the model developed in this study effectively isolates market inefficiencies and can be used to compare the efficiency of different markets, the data to test this relationship properly does not currently exist. A multi-year sample of posted exchange prices is necessary for proper analysis, and a program would have to be developed and run over time to collect this information, as the major exchanges do not keep (or are not willing to distribute) data in this format.

If and when this data is available, I believe the model will find that the sports exchanges arrive at a price somewhere between sportsbooks and the fair price. Removing discriminating linesmakers will increase efficiency, but the high interconnectedness of the markets and inherent bettor biases will keep the exchange from ever becoming truly efficient.

Despite any current or future conclusions that may be drawn from studies like these, one must be careful when attempting to generalize the results to real financial markets. A study by Tetlock (2004) uses both sports and financial markets from TradeSports.com to find that inefficiencies in the sports betting markets do not necessarily have counterparts in real financial markets. While results from future study may reveal the true effect of market structure on sports betting pricing efficiency, the same conclusion might not hold for financial markets.

Most studies in this field view the efficiency of these markets from a purely financial sense, i.e., if prices do not accurately approximate the value of a bet, the market is working suboptimally. This analysis assumes all bettors are investors, and utility is derived only from economic gain. Market participants must be rational, however, and Vernón (2003) and others incorporate utility by assuming it is the residual in the efficiency calculations and solve for it by knowing the other parameters. This study has used team popularity-based factors to incorporate non-financial utility, but assumes priced-based and historical results-based factors are public misperceptions and truly inefficient outcomes. Perhaps a future, more psychologically-based study can analyze the preferences of the betting public in more color and complexity, instead of just inferring it from the sum of its bets. Doing so will allow us to separate both financial utility and emotional utility from the market's true underlying inefficiencies.

As Gabriel and Marsden (1990) pose, "Are we observing an inefficient market or simply one in which the tastes and preferences of the market participations lead to the observed results?" The model successfully developed in this piece, when applied to more homogenous and robust data, will hopefully be able to determine whether market inefficiencies are a product of the market structure or whether they lie entirely in the bias of bettors.

Appendix: Deriving Probabilistic Form of Prices

For the expectation value of the bet to be zero, the bet must win *p* percent of the time.

For underdogs, one is betting 100 to win X.

$$0 = Xp_u - 100(1 - p_u)$$
$$0 = Xp_u - 100 + 100p_u$$
$$100 = p_u(X + 100)$$
$$p_u = \frac{100}{X + 100}$$

For favorites, one is betting X to win 100:

$$0 = 100p_f - X(1 - p_u)$$
$$0 = 100p_f - X + Xp_f$$
$$X = p_f(X + 100)$$
$$p_f = \frac{X}{X + 100}$$

References

Borghesi, Richard (2004). Weather Biases in the NFL Betting Market: Explaining the Home Underdog Effect. Retrieved December 8, 2006, from http://www.business.txstate.edu/users/rb38/Weather_Biases.pdf

Camerer, Colin F. (1989). *The American Economic Review*, *79*(5), 1257-1261.

Dixon, Mark J. & Stuart G. Coles (1997). Modeling Association Football Scores and Inefficiencies in the Football Betting Market. *Applied Statistics*, *46*(2), 265-280.

Doc's Sports Service (2007). Retrieved February 25, 2007, from http://www.docsports.com/mlb-schedule.html

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

Gabriel, Paul & James Marsden (1990). An Examination of Market Efficiency in British Racetrack Betting. *Journal of Political Economy*, *98*(4), 874-885.

Gandar, John M., William H. Dare, Craig R. Brown, & Richard A. Zuber (1998). Informed traders and price variations in the betting market for professional basketball games. *Journal of Finance*, *53*, 385-401.

Gandar, John, Richard Zuber, Thomas O'Brien, & Ben Russo (1988). Testing Rationality in the Point Spread Betting Market. *Journal of Finance*, *43*(4), 995-1008.

Gilovich, Thomas, Robert Vallone & Amos Tversky (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, *17*, 295-314.

Golec, Joseph & Maurry Tamarkin (1991). The degree of inefficiency in the football betting markets. *Journal of Financial Economics*, *30*, 311-323.

Gray, Philip K. & Stephen F. Gray (1997). Testing Market Efficiency: Evidence from the NFL Sports Betting Market. *Journal of Finance*, *52*(4), 1725-1737.

Hetherington, Alex (2006). Betting Against Efficiency: Behavioral Finance in an NFL Gambling Exchange. Retrieved December 4, 2007, from the Social Science Research Network, http://www.ssrn.com

Levitt, Steven D. (2004). Why Are Gambling Markets Organised So Differently from Financial Markets?. *The Economic Journal*, *114*, 223-246.

Ozgit, Alper E. (2006). Three Essays on Market Structure and Pricing. University of California – Los Angeles. Retrieved November 6, 2006, from http://www.bol.ucla.edu/~ozgit/AlperOzgit-Dissertation.pdf

Sauer, Ray (1998). The Economics of Wagering Markets. Journal of Economic Literature, *36*(4), 2021-2064.

Snowberg, Erik & Justin Wolfers (2005). The Favorite-Longshot Bias: Understanding a Market Anomaly. *Efficiency of Sports and Lottery Markets*, Elsevier: Handbooks in Finance Series.

Snowberg, Erik & Justin Wolfers (2006). Explaining the Favorite-Longshot Bias: Is it Risk-Love, or Misperceptions? Working draft. Retrieved March 19, 2007, from http://cbdr.cmu.edu/seminar/Wolfers.pdf

Strumpf, Koleman (2003). Illegal sports bookmakers. University of North Carolina, Department of Economics. Retrieved November 26, 2006, from http://www.unc.edu/~cigar/papers/Bookie4b.pdf

Tetlock, Paul (2003). How Efficient are Information Markets? Evidence from an Online Exchange. Retrieved December 5, 2006, from http://www.mccombs.utexas.edu/faculty/Paul.Tetlock/papers/Tetlock-Efficient_Info_Markets-01_02.pdf

Thaler, Richard & William Ziemba (1988). Anomalies: Parimutual Betting Markets: Racetracks and Lotteries. *Journal of Economic Perspectives*, *2*(2), 161-174.

Vernón, Antonio M. R. (2003). Market Efficiency and March Madness: Empirical Tests of Point Spread Betting. Working paper. Retrieved November 6, 2006, from Social Science Resource Network, http://www.ssrn.com

Wolfers, Justin (2006). Point Shaving: Corruption in NCAA Basketball. *American Economic Review*, *96*(2), 279-283.

Woodland, Linda M. & Bill M. Woodland (1994). Market Efficiency and the Favorite-Longshot Bias: The Baseball Betting Market. *The Journal of Finance*, *49*(1), 269-279.

Woodland, Linda M. & Bill M. Woodland (2001). Market Efficiency and Profitable Wagering in the National Hockey League: Can Bettors Score on Longshots?. *Southern Economic Journal*, *67*(4), 983-995.