

Is Liquidity Related to Prediction Market Efficiency?

Paul C. Tetlock*

December 2007

Abstract

I investigate the relationship between liquidity and market efficiency using data from short-horizon binary outcome securities listed on the TradeSports exchange. I find that liquidity does not reduce—and sometimes increases—deviations of prices from financial and sporting event outcomes. One explanation is that limit order traders are naïve about other traders' knowledge and unwittingly bet against them, which can slow the response of prices to information. Consistent with this explanation, the limit orders that execute during informative events have negative expected returns; and limit orders often execute against traders who exploit the well-known favorite-longshot bias in prices.

* Please send all comments to paul.tetlock@yale.edu. I am currently visiting the finance department at Yale, on leave from Texas. I am extremely grateful to TradeSports for their cooperation in allowing me to collect the liquidity data, and their generosity in providing additional trading data. I thank Yakov Amihud, Malcolm Baker, Ed Glaeser, Larry Glosten, John Griffin, Robert Hahn, Bing Han, Alok Kumar, Terry Murray, Chris Parsons, Matt Spiegel, Laura Starks, Avaniidhar Subrahmanyam, Sheridan Titman, Justin Wolfers, Eric Zitzewitz, and seminar participants at Kansas and Texas for helpful comments.

I. Introduction

There is a growing body of theoretical and empirical research on securities markets for contingent claims on uncertain events, often called prediction or information markets. A key finding in this literature is that the prices in prediction markets can help to produce forecasts of event outcomes with a lower mean squared prediction error than conventional forecasting methods. For example, incorporating the prices in prediction markets increases the accuracy of poll-based forecasts of election outcomes (Berg, Forsythe, Nelson and Rietz (2003)), official corporate experts' forecasts of printer sales (Chen and Plott (2002)), and statistical weather forecasts used by the National Weather Service (Roll (1984)).¹

This study provides new evidence on how the forecasting accuracy of event prediction markets depends on market liquidity. I use three years of intraday data on one-day binary outcome securities based on sports and financial events traded on an online exchange, TradeSports.com. I use three standard measures of securities liquidity, low quoted bid-ask spreads, high market depth, and high trading volume. The idea is that, in a liquid market, traders can cheaply buy and sell large quantities—e.g., O'Hara (1995).

I measure event forecasting accuracy as the difference between securities prices and terminal cash flows. Unlike securities in conventional financial markets, I study TradeSports' securities that pay a single terminal cash flow at the end of their one-day horizons, allowing me to repeatedly observe securities' fundamentals (Thaler and Ziemba (1988)). In addition, many of the securities are exposed to negligible systematic risks.

¹ See Wolfers and Zitzewitz (2004) for a survey of the literature on prediction markets.

Thus, I can directly measure absolute pricing efficiency as the deviation of prices from terminal cash flows.

I find that liquidity does not reduce—and sometimes increases—deviations of securities prices from financial and sporting event outcomes. Prices can deviate from outcomes because average prices differ from average event outcomes (i.e., poor calibration), or because prices do not distinguish between events with different probabilities (i.e., poor resolution). I find that the prices of liquid securities are not better calibrated, and actually exhibit poorer resolution than the prices of illiquid securities.

These results are important for both direct and indirect reasons. Directly, the findings address a key concern in using prediction markets prices to forecast events: their accuracy could depend critically on market liquidity and trading activity. Several researchers emphasize the potential of prediction markets to improve decisions (e.g., Hanson (2002), Hahn and Tetlock (2005), Sunstein (2006), and Cowgill, Wolfers, and Zitzewitz (2007)). In principle, the range of applications is virtually limitless—from helping businesses make better investment decisions to helping governments make better fiscal and monetary policy decisions. For example, decision makers in the Department of Defense, the health care industry, and multi-billion-dollar corporations, such as Eli Lilly, General Electric, Google, France Telecom, Hewlett-Packard, IBM, Intel, Microsoft, Siemens, and Yahoo, conduct internal prediction markets. The prices in these markets reflect employees' expectations about the likelihood of a homeland security threat, the nationwide extent of a flu outbreak, the success of a new drug treatment, the sales revenue from an existing product, the timing of a new product launch, or the quality of a recently introduced software program.

Indirectly, the results shed light on the general relationship between liquidity and absolute pricing efficiency, which could teach us something about financial markets. Formally testing how liquidity and absolute pricing efficiency relate in conventional financial markets would require that researchers observe securities terminal cash flows. Even in cases where this is possible, such as in options or bond markets, other obstacles exist. In options markets, inferences may depend critically on assumptions about the market price of risk. In bond markets, there is often insufficient volatility in cash flows to make the exercise interesting. By necessity, researchers in finance must rely on strong assumptions relating absolute pricing efficiency to observable proxies for market efficiency to estimate its relationship with liquidity. For example, many researchers argue that liquidity increases market efficiency based on evidence from short-horizon return predictability tests (e.g., Chordia, Roll, and Subrahmanyam (2006), and Chordia et al. (2007)). Yet large deviations in prices from fundamentals could be associated with very little or even no return predictability (Summers (1986)). Because nearly all efficiency tests in equity markets do not employ measures of stocks' fundamentals, one cannot infer whether liquidity increases or decreases absolute pricing efficiency.

Despite these obstacles, assessing the relationship between liquidity and absolute pricing efficiency remains important because absolute pricing efficiency—not return predictability—determines whether capital is being efficiently allocated in financial markets. A better understanding of the impact of liquidity on absolute pricing efficiency could inform the decisions of firms and governments whose actions affect liquidity, and thereby affect the efficient allocation of capital.

Beyond offering a particularly clean test of efficiency, there are other reasons to expect that findings from TradeSports data are relevant for other financial markets. The TradeSports exchange uses standard continuous double auctions, comparable to the mechanisms used in the world's major stock, currency, commodity and derivatives exchanges.² Many professional financial traders from New York, Chicago, and London wager thousands of dollars in sports and financial markets on TradeSports.

Prices on TradeSports and conventional financial markets are likely to be correlated. Several traders conduct automatic arbitrage transactions across conventional exchanges and TradeSports, forcing similar pricing to prevail in these markets. Zitzewitz (2006) reports that program trading accounts for over 95% of orders in the TradeSports binary options securities on the daily Dow Jones index. Moreover, he finds surprising evidence that some price *discovery* takes place in these daily Dow options. Specifically, the implied volatilities from TradeSports binary options on the Dow can help forecast the level and volatility of the Dow Jones index, even after controlling for multiple lags of implied and historical volatilities from conventional CBOT, CBOE, and NYSE securities. Arbitrage promotes the transmission of newly discovered price information from the TradeSports exchange to other financial markets, and vice versa, within minutes.

There is also likely to be a common component in liquidity for similar contracts traded on TradeSports and conventional financial exchanges because the popularity of underlying securities indices is positively correlated across markets. For example, the Dow securities are the most popular and actively traded on TradeSports; and the Dow

² In both theory and practice, double auctions are particularly robust mechanisms that promote rapid adjustment towards market equilibrium even in the presence of market frictions and trader irrationality—e.g., Gode and Sunder (1993), Cason and Friedman (1996), Gjerstad and Dickhaut (1998), Noussair *et al.* (1998), and Satterthwaite and Williams (2002).

stocks are also the most heavily traded stocks on the NYSE. Whatever motivates investors to trade certain securities in one market may induce them to trade similar securities in another market. Based on these similarities in liquidity and pricing across markets, an understanding of the relationship between liquidity and efficiency on TradeSports complements the indirect evidence from conventional financial markets, where one generally cannot observe the terminal cash flows of highly volatile securities.

There are compelling theoretical reasons to think that liquidity and efficiency could be either positive or negatively related. In many rational models of liquidity provision, liquidity is beneficial to efficiency because it enhances the incentives for traders to become informed and reveal their information through trading. Usually, the indirect source of liquidity is random noise trading, which does not explicitly counteract informed trading but does allow rational market makers to break even. An alternative outside the scope of rational models is that some agents are naïve about the adverse selection problem. If naïve agents submit near-market limit orders with insufficient regard for the future release of information, they will unwittingly and systematically trade with informed traders, effectively subsidizing liquidity provision. In such a market with excessive liquidity provision, prices could respond more slowly to information.³

On TradeSports, I document two facts that are consistent with naïve liquidity provision. Both facts could help to explain why forecasting accuracy is not higher in the most liquid prediction markets, despite the strong incentives for information acquisition and dissemination in liquid markets. First, limit orders that passively execute during informative events have negative expected returns to expiration. This suggests that limit

³ Linnainmaa (2007) and Kaniel, Liu, Saar, and Titman (2007) report indirect evidence in financial markets, but they cannot observe firms' terminal cash flows. Baker and Stein (2004) and Tetlock and Hahn (2007) present theoretical models in which liquidity subsidies lead to market underreaction.

orders retard the response of prices to new information and thereby inhibit forecasting resolution. Second, limit orders traders passively buy low-priced securities—e.g., \$4 and below—and sell high-priced securities—e.g., \$6 and above. Such limit orders could sustain the overpricing of low-priced securities and the underpricing of high-priced securities, a pattern known as the favorite-longshot bias (e.g., Thaler and Ziemba (1988)). This second finding identifies a mechanism for how prices in liquid markets could remain more poorly calibrated than prices in illiquid markets.

The layout of the paper is as follows. In Section II, I discuss the relationship between this study and the prediction market and finance literatures. In Section III, I describe the structure of the securities data from the TradeSports exchange and the measures of liquidity used throughout the paper. In Section IV, I assess absolute pricing efficiency in liquid and illiquid securities markets, focusing on forecasting calibration and resolution. In Section V, I examine whether liquidity providers are naïve, documenting the two key facts described above. In Section VI, I conclude and suggest directions for further research on liquidity and securities market efficiency.

II. The Prediction Market Literature and Its Relationship to Finance

This study on prediction market efficiency is related to studies on the favorite-longshot bias in wagering markets (e.g., Jullien and Selanie (2000), Wolfers and Zitzewitz (2004), and Zitzewitz (2006)) and financial markets (e.g., Rubinstein (1985), Brav and Heaton (1996), and Barberis and Huang (2007)). A key distinction is that these prior studies do not link informational efficiency and market liquidity.

Several theoretical and empirical papers, however, do link these concepts. One view is that illiquidity represents a transaction cost for informed arbitrageurs whose trades make prices more efficient. For example, when liquidity increases in Kyle's (1985) model, informed traders bet more aggressively based on their existing information because their trades have a smaller impact on prices. In addition, informed traders have greater incentives to acquire more precise information in liquid markets. If informed arbitrageurs are less active in illiquid markets where trading is expensive, securities' prices in these markets may deviate by large amounts from their fundamental values. An alternative view is that liquidity is a proxy for non-informational or noise trading, which may harm informational efficiency.⁴ In behavioral finance models, various limits to arbitrage prevent rational agents from making aggressive bets against noise traders—e.g., DeLong, Shleifer, Summers, and Waldmann (1990a) and Shleifer and Vishny (1997). If liquid markets have more noise trading than illiquid markets and rational agents do not fully offset noise traders' systematic biases, then securities prices in liquid markets may be inefficient relative to prices in illiquid markets.

Some recent papers provide indirect empirical support for the view that securities mispricing is greater in illiquid markets—e.g., Wurgler and Zhuravskaya (2002), Chordia, Roll, and Subrahmanyam (2006), and Chordia et al. (2007). Yet other papers provide empirical evidence that suggests mispricing is greater in liquid markets—e.g., Bloomfield, O'Hara, and Saar (2007), Linnainmaa (2007), and Tetlock (2007). Of all these papers, only Bloomfield, O'Hara, and Saar (2007) directly examines the deviation of securities prices from fundamental values because the terminal cash flows of the

⁴ Variations in liquidity correspond to variations in noise trading in Kyle (1985), Glosten and Milgrom (1985), and Baker and Stein (2004), but they may also result from variations in the search costs that buyers and sellers incur in their efforts to transact—e.g., Duffie, Garleanu, and Pedersen (2005).

securities are unobservable in most real-world situations. The efficiency results from TradeSports, where professional traders exchange large sums of money, demonstrate that several laboratory results in Bloomfield, O’Hara, and Saar (2007) could generalize to real-world financial markets.

III. Securities Data and Measures of Liquidity

I construct an automatic data retrieval program to collect comprehensive limit order book and trading history statistics about each security traded on the TradeSports exchange.⁵ The program runs at 30-minute intervals almost continuously from March 17, 2003 to October 23, 2006.⁶ All empirical tests in this paper include only data from the single-day sports and financial securities that the program records. The vast majority of TradeSports’ securities are based on one-day sports or financial events, such whether the Yankees will win a particular baseball game or whether the Dow Jones Index will close 50 or more points above the previous day’s close. I focus on these securities to limit the number of factors needed as controls in the statistical analysis that follows.⁷ Roughly 70% of TradeSports’ securities are based on sports events, 25% are based on financial

⁵ I am grateful to TradeSports Exchange Limited for granting me permission to run this program.

⁶ The program’s 30-minute interval is approximate because it records securities sequentially, implying that the exact time interval depends on whether new securities have been added or subtracted and precise download speeds. In practice, these factors rarely affect the time interval by more than two minutes. The program stops running only for random author-specific events, such as software installations, operating system updates, power failures, and office relocation—and technical TradeSports issues, such as daily server maintenance and occasional changes in the web site’s HTML code.

⁷ For sports events, I consider only securities with an official TradeSports description that includes either the text “game,” “bout,” or “match”; for financial events, I consider only securities based on the daily level of stock indexes, which is the vast majority of all financial securities. Because almost all of the uncertainty for these events is resolved during the scheduled event time, I focus on observations occurring within one event duration of the event starting time—e.g., no more than 3 hours prior to the start of a 3-hour event.

events and fewer than 5% are based on events in all other categories combined—*e.g.*, economic, political, entertainment, legal, weather, and miscellaneous.

The TradeSports exchange solely facilitates the trading of binary outcome securities by its members, and does not conduct transactions for its own account.⁸ Securities owners receive \$10 if a pre-specified, verifiable event occurs and \$0 otherwise—*e.g.*, the owner of the Dow Jones security mentioned above receives \$10 if and only if the index goes up by 50 points or more. For ease of interpretation, the exchange divides its security prices into 100 points, worth \$0.10 per point. The minimum securities price increment, or tick size, ranges between 0.1 and one point in this sample.

TradeSports levies a commission of no more than 0.4% of the maximum securities price (\$10) on a per security basis whenever a security is bought or sold.⁹ At the time of security expiration, when the payoff event is verifiable and the owner receives payment, traders must liquidate all outstanding securities positions and incur commissions. Note that the maximum \$0.08 round-trip transaction fee is smaller than the value of one point (\$0.10) for most securities. This implies arbitrageurs have an incentive to push prices back towards fundamental values if they stray by even one point.

Following conventions in other studies of financial markets, I eliminate observations from TradeSports markets with poor quality price data. Specifically, I exclude observations with a cumulative trading volume below 10 securities (\$100), market depth below 10 securities (\$100), or bid-ask spreads exceeding 10% (\$1.00).¹⁰

⁸ TradeSports limits the risk that the counterparty in a security transaction will default by imposing symmetric margin requirements for each sale or purchase of a security by one of its members. Generally, 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.

⁹ The exchange eliminated commissions for non-marketable limit orders on November 9, 2004.

¹⁰ The observations excluded by the combination of all three restrictions represent only 5% of volume on the exchange. Using more stringent market activity criteria tends to strengthen the statistical results below.

These restrictions are designed to exclude securities without well-established market prices for which tests of efficiency are not meaningful.

I use three measures of securities market liquidity: quoted bid-ask spreads, market depth, and cumulative trading volume. I define the quoted spread as the difference between the inside (lowest) ask and (highest) bid quotations. I compute market depth as the sum of all buy and sell limit orders within the maximum 10% bid-ask spread, divided by two. To enhance comparability across time periods and contracts, I use ad hoc cutoffs of 1 point, 3 points, and 5 points—i.e., \$0.10, \$0.30, and \$0.50—to partition observations into groups sorted according to bid-ask spreads. The results are similar if I define partitions based on the historical distribution of spreads during a rolling time window. I follow an analogous procedure to sort securities by market depth and cumulative trading volume, using ad hoc cutoffs of 100, 300, and 500 securities for both variables.

Table I shows that the three measures designed to capture liquidity exhibit some similarity. The table reports the pair-wise correlations between the logarithms of quoted spreads (*LnSpread*), market depth (*LnDepth*), and trading volume (*LnVolm*) for all securities, sports securities and financial securities. All nine of the correlations in Table I have the expected signs, and are statistically significant at the 1% level. For the group including all securities, the volume correlations with spreads and depth are low because the mean volume-based liquidity measure is higher in financial securities, whereas the spread- and depth-based liquidity measures are higher in sports securities. The volume correlations with spreads and depths are higher in the separate analyses of sports and financial securities. The table also reports the average number of program observations per security, which is between four and six for both sports and financial securities.

[Insert Table I here.]

IV. Tests of Market Efficiency

Now I analyze the absolute pricing efficiency of securities on the TradeSports exchange. I conduct these tests separately for groups of securities sorted by each of the three liquidity measures. I also analyze the efficiency of sports and financial markets separately to investigate the possibility that pricing in these markets differs significantly. Fortunately, the key results in this study apply to both sports securities and financial securities, regardless of their exposure to market risk.¹¹

A. Forecasting Calibration in Liquid and Illiquid Markets

Microstructure theory measures absolute pricing efficiency as the expected mean squared error of prices minus cash flows, which can be decomposed into two components as in Equation (1):

$$(1) \quad E[(Payoff - Price)^2] = E[Payoff - Price]^2 + E[(Payoff - 100 * Event Probability)^2]$$

First, I focus on measuring the squared bias component of absolute pricing efficiency (*i.e.*, the first term), as opposed to the conditional variance component (*i.e.*, the second term). This emphasis is standard practice in the literature on market efficiency in binary prediction markets because expected cash flows are not directly observable for each

¹¹ In unreported tests, I allow for the possibility that financial securities with positive exposure to market risk have different expected returns from those with negative risk. I find a positive, but insignificant, risk premium of less than 1% for the typical financial security with positive exposure to the market. This is not surprising because three years of data is usually insufficient for estimating market risk premiums.

security. Thus, researchers cannot measure securities' conditional variances without making assumptions about biases, but they can estimate average biases by comparing securities' average cash flows to their prices. In this subsection, I estimate biases in prices following conventions from related work—*e.g.*, Wolfers and Zitzewitz (2004), Tetlock (2004), and Zitzewitz (2006). In the next subsection, I make additional assumptions in an effort to indirectly estimate the conditional variance component of mean squared error.

I employ a straightforward regression methodology to test the null hypothesis that securities prices are unbiased predictors of securities' cash flows. The null hypothesis is that securities' expected returns to expiration are zero, regardless of the current securities price. The alternative hypothesis is that Kahneman and Tversky's (1979) theory of probability perception describes the pattern of expected returns across securities with different current prices. A single observation consists of a security's current price and its returns until expiration. I measure current prices using the midpoints of the inside bid and ask quotations to avoid the problem of bid-ask bounce that could affect transaction prices. The results are robust to using the most recent transaction price instead.

I calculate a security's percentage returns to expiration by subtracting its current price from its payoff at expiration, which is either 0 or 100 points, then dividing by 100 points.¹² This is the standard measure of returns in the prediction markets literature—*e.g.*, Wolfers and Zitzewitz (2004), Tetlock (2004), and Zitzewitz (2006). Relative to alternative measures that divide by price or the duration of the holding period, the measure of returns to expiration described above possesses the advantages of being much closer to homoskedastic and normally distributed, and it is symmetric for buyers and

¹² I exclude the very small fraction of TradeSports contracts that do not expire at 0 or 100 points. I divide by 100 points to represent the combined amount of capital that buyers and sellers invest in the security.

sellers. For the securities that I examine, the natural unit of information release is an event, rather than a given amount of time. In addition, the opportunity cost of invested funds is trivial over the daily time horizon of these securities.

The S-shaped form of the probability weighting function hypothesized in Kahneman and Tversky (1979) and formalized in Prelec (1998) informs my choice of pricing quantiles and statistical tests. The S-shape refers to a graph of subjective versus objective probabilities, as shown first in Kahneman and Tversky (1979). Prelec (1998) derives the S-shape in probability misperception from axiomatic foundations. His theory predicts that agents overestimate the likelihood of events with objective probabilities less than $1/e = 37\%$ and underestimate the likelihood of events with objective probabilities greater than $1/e$. There is also an ample body of empirical evidence that is consistent with a probability weighting function having a fixed point in the neighborhood of $1/e$ (Tversky and Kahneman (1992), Camerer and Ho (1994), and Wu and Gonzalez (1996)).

Based on this evidence, I construct dummy variables, Price1 through Price5, for five equally-spaced pricing intervals: (0,20), [20,40), [40,60), [60,80), and [80,100) points. I then measure the returns until expiration for securities in each pricing quantile. I test the null hypothesis that all returns to expiration are equal to zero against the alternative that securities in the first two quantiles—Price1 and Price2—based on small probability events ($p < 40\%$), are overpriced and securities in the last three categories—Price3, Price4, and Price5—based on large probability events ($p \geq 40\%$), are underpriced.

I report the results from three Wald (1943) tests based on this simple idea. The first Wald test measures whether small probability events are overpriced on average:¹³

¹³ Despite the directional nature of over- and underpricing predicted by Kahneman and Tversky (1979) and the favorite-longshot bias, I use two-tailed Wald tests to be conservative.

$$(2) \quad (\text{Price1} + \text{Price2}) / 2 = 0$$

The second Wald test assesses whether large probability events are underpriced:

$$(3) \quad (\text{Price3} + \text{Price4} + \text{Price5}) / 3 = 0$$

The third Wald test measures whether large probability events are more underpriced than small probability events—*i.e.*, whether the mispricing function is S-shaped:

$$(4) \quad (\text{Price3} + \text{Price4} + \text{Price5}) / 3 - (\text{Price1} + \text{Price2}) / 2 = 0$$

Of the three, this is the most powerful test of the null hypothesis against the Kahneman and Tversky (1979) alternative because it accounts for other factors that could influence average returns in both small and large probability events.¹⁴ I also interpret equation (4) as a test of the favorite-longshot bias: the expected returns on longshots, with prices less than 40, are lower than the expected returns on favorites. In the original Ali (1977) racetrack data, expected returns are roughly zero for horses with a 30% probability of winning, and increase nearly monotonically with a horse's probability of winning.

I use standard ordinary least squares to estimate the coefficients of the five pricing categories. For all regression coefficients, I compute clustered standard errors, as in Froot (1989), to allow for correlations in the error terms of all securities expiring on the same calendar day, which simultaneously corrects for the repeated sampling of the same security and the sampling of related events.¹⁵ This clustering procedure exploits the fact that all event uncertainty is resolved on the day of expiration (see footnote 7).

To illustrate the efficiency tests and give an overview of the data, I first examine the returns to expiration for all sports securities, all financial securities, and both groups

¹⁴ The choice of how to partition the pricing categories has little effect on the Wald tests because, regardless of the partitioning, these tests assess whether the returns to expiration of securities priced below 40 points differ from the returns of securities priced above 40 points.

¹⁵ Using a finer clustering unit based on the expiration day and type of security does not affect the results.

together. Table II displays the regression coefficient estimates for Price1 through Price5 along with the three Wald tests described above. One interesting result is that neither sports nor financial securities exhibit substantial mispricing, which is consistent with Wolfers and Zitzewitz (2004) and Tetlock (2004).

[Insert Table II around here.]

The qualitative patterns in the pricing of both sets of securities and in their aggregate suggest, however, that the favorite-longshot bias could play a role in any mispricing that does exist. To aid the reader in identifying the S-shaped favorite-longshot bias, Figure 1 provides a visual representation of mispricing in each pricing category for sports, financial, and both types of securities.

[Insert Figure 1 around here.]

The securities based on small probability sports events in pricing quantiles 1 and 2 are slightly overpriced by 1.15 points (p -value = 0.127); and financial securities based on large probability events in quantiles 3, 4, and 5 are underpriced by 2.21 points (p -value = 0.030). The Wald test for the favorite-longshot bias rejects the null hypothesis of no difference in expected returns at the 5% level for both sports and financial securities. Interestingly, the magnitude of this bias decreases from an average of 2.13 points across the sports and financial groups to just 1.24 points in the aggregate group, which is barely statistically significant at the 5% level. This reduction occurs because of differences in the pricing patterns of sports and financial securities and the changing relative composition of sports and financial securities within pricing quantiles.¹⁶

¹⁶ The disparity between the average of the individual estimates and the aggregate estimate illustrates the importance of estimating the effects on sports and financial securities separately.

Having established that both sports and financial securities show a limited degree of inefficiency on average, I now turn to the key test of whether the favorite-longshot bias is more pronounced in illiquid or liquid securities. Panels A and B of Table III analyze the favorite-longshot bias in sports and financial securities. The four regressions in each panel in Table III sort securities according to their liquidity quantiles as measured by bid-ask spreads. In both panels, the main result is that the favorite-longshot bias is no smaller in liquid securities than the bias in illiquid securities. If anything, the bias appears larger in liquid securities—e.g., the bias in liquidity quantiles 3 and 4 minus the bias in quantiles 1 and 2 is positive (2.23%), but not statistically significant. Securities in the two intermediate bid-ask spread quantiles exhibit a somewhat smaller favorite-longshot bias, but this conclusion does not generalize for other liquidity measures. For the bid-ask spread measure, the same patterns in the favorite-longshot bias apply to both sports and financial securities.

[Insert Table III here.]

Based on the bid-ask spread measure, the financial securities are less liquid than the sports securities. Panel B shows that relatively few financial securities have bid-ask spreads less than \$0.10, or 1% of the \$10 maximum contract value. The financial securities, however, are not less liquid based on the other liquidity measures, such as limit order book depth and trading volume.

Next, I repeat the same regression analysis for sports and financial securities sorted based on the other two liquidity measures: market depth and trading volume. For brevity, Figure 2 summarizes only the favorite-longshot bias measures from these

analyses, together with the bid-ask spread results above. Figure 2 shows that the favorite-longshot bias is virtually always positive, regardless of securities type or liquidity.

[Insert Figure 2 here.]

Overall, liquid and illiquid securities exhibit favorite-longshot biases of a similar magnitude. In sports markets, liquid securities have slightly larger biases, but often not significantly larger. In financial markets, the relationship between liquidity and the favorite-longshot bias is quite weak, and depends on the liquidity measure. For example, in financial securities with market depth in the top quantile, there is no favorite-longshot bias; yet, in financial securities with bid-ask spreads in the bottom quantile, the favorite-longshot bias exceeds 6%. This 6% bias is large relative to the bid-ask spreads of less than 1% and the trading commissions of 0.8% or lower. More generally, the surprising finding is that prices in liquid prediction markets are not better calibrated than prices in illiquid markets.

A lack of statistical power does not seem to explain the lack of evidence for liquidity improving efficiency. The favorite-longshot bias is consistently positive. Moreover, the bias is actually larger in the top two liquidity quantiles than the bottom two quantiles in five out of six cases in Figure 2. In addition, the confidence intervals are sometimes narrow enough to reject the hypothesis that liquid securities are better calibrated than illiquid securities at the 5% level—e.g., in the case of the trading volume measure in sports securities. I offer a partial resolution to this puzzle in Section V.

In unreported tests, I evaluate the possibility that other security characteristics, such as return volatility and time until expiration could explain the relationship between mispricing and liquidity. The inclusion of the controls based on volatility and time

horizon does not significantly increase the explanatory power of the regressions in Table III. Proponents of market efficiency can take comfort in the fact that few variables, other than liquidity and price, are useful predictors of securities' returns to expiration.

B. Forecasting Resolution in Liquid and Illiquid Markets

In this subsection, I measure how quickly securities prices converge toward event outcomes and whether these rates depend on liquidity. In models of rational adverse selection, periods of illiquidity tend to precede the release of information about securities' fundamental values and its incorporation into prices. Intuitively, traders are reluctant to submit limit orders because they are wary that traders who submit market orders possess information that they do not have. Thus, a lack of limit orders, or illiquidity, could predict the release of information even if illiquidity does not cause prices to incorporate information more rapidly.

In an effort to avoid this reverse causality problem, I estimate the impact of liquidity on information release using instrumental variables (IV) for changes in liquidity. Interestingly, the simple OLS estimates are qualitatively similar to the IV estimates, suggesting that accounting for traders' endogenous responses to time-varying adverse selection does not affect the main findings. In Section V, I directly assess how limit order traders respond to time-varying information release.

I use two instruments for market design that could have a significant impact on prediction market liquidity, but do not have a direct impact on the release of event information in sports and financial markets. On September 1, 2004, the TradeSports

exchange began a series of three changes in its trading fee structure. First, the exchange reduced the trading fees on securities priced above \$9.50 or below \$0.50 by half. Within two weeks, the exchange reduced the minimum tick size on its securities from \$0.10 to \$0.01. Finally, on November 9, 2004, the exchange eliminated all trading fees on limit orders that did not immediately execute. I construct two dummy variables to instrument for changes in liquidity that occurred during the transition phase, September 1, 2004 to November 9, 2004, and for changes occurring after all new rules were implemented. I use observations on securities from March 1, 2004 through February 9, 2005 to evaluate the impact of these changes.¹⁷

I measure the extent of information release during an event as the convergence of prices toward securities' terminal payoffs at expiration—*i.e.*, the reduction in the squared difference between prices and terminal payoffs. To analyze the change in this squared difference, I use the identity in Equation (1) that decomposes mean squared error into squared bias (calibration) and conditional variance (resolution) components.

The measurement of information release as the reduction in mean squared pricing error (MSPE) requires an important assumption: reductions in MSPE occur only from reduction in the conditional variance component—*i.e.*, the squared bias component does not change. This may be a reasonable approximation for short time horizons, but the previous results suggest this assumption could fail for long time horizons because squared bias approaches zero as an event ends. Specifically, if liquid markets exhibit a greater favorite-longshot bias that is corrected over time, periods of liquidity will appear

¹⁷ Because an idiosyncratic author-specific event limits the number of observations in the pre-event period, I use a six-month pre-event window to obtain a similar number of observations before and after the two rule changes. The key regression coefficients are quantitatively sensitive to the specific event window chosen, but the qualitative conclusions are robust.

to precede greater information release even if there is no causal relationship between the two. This reasoning suggests my assumption biases the results toward finding that liquidity causes prices to incorporate information more rapidly—but the bias is small because there is not a large difference in mispricing between liquid and illiquid markets.

I use a two-stage least squares approach to examine how liquidity and the rate of price convergence change as TradeSports changes its market design. I control for several factors other than liquidity that influence the release of event information. For each sports and financial security, I create a continuous scaled measure of the event time: $\text{event time} = \text{current time} - \text{event start time} / (\text{event duration})$. I then create a series of dummy variables to indicate the discrete event time position of a security in 20% increments, where $EvTime_X$ corresponds to an event time interval of $((X-1)*0.2, X*0.2]$ and $X = 0, 1, \dots, 6$. Although some sports events last longer than expected, financial events almost never do because they are based on index levels at pre-specified times. I also control for the logarithm of the event time, in case the dummies fail to capture within-event trends in liquidity or information release. In addition, I control for the logarithm of the date on which the event is scheduled to expire to ensure that the excluded market design instruments do not capture long-term trends in liquidity and information release.

Finally, I include a control for expected information release (*ExpInfo*) based on the current securities price level. *ExpInfo* is defined as $(100 - Price) * Price$, which is the expected value of squared pricing error under the assumption of market efficiency. To see this, note that market efficiency implies prices equal true event probabilities, so that Equation (5) holds:

$$(5) \quad E[(\text{Payoff} - \text{Price})^2] = E[(\text{Payoff} - 100 * \text{Event Probability})^2]$$

$$= (100 - \text{Event Probability}) * \text{Event Probability} = (100 - \text{Price}) * \text{Price}$$

In robustness checks, I find that the liquidity coefficient estimates are not particularly sensitive to the use of alternative sets of control variables.

Table IV reports the results from six IV regressions each of which measures the causal effect of liquidity on the reduction in MSPE over the next data recording period (roughly 30 minutes). The six (2 x 3) IV regressions correspond to the results for sports and financial securities for each of the three measures of liquidity: the logarithms of bid-ask spreads (*LnSpread*), market depth (*LnDepth*), and trading volume (*LnVolm*). All standard errors in Table IV are robust to heteroskedasticity and arbitrary intra-group correlation within securities that expire on the same day (Froot (1989)).

[Insert Table IV here.]

Table IV consistently shows that an increase in liquidity slows the rate at which prices incorporate event information. In four out of six IV regressions, I reject the null hypothesis of no effect with a 1% Type I error; and, in the other two regressions, I reject the null at the 10% level. The coefficient estimates are reasonably similar across the six specifications. As bid-ask spreads decline, as market depth goes up, and as trading volume increases, MSPE decreases more slowly. In all six specifications, the magnitude of the effect is between 200 and 600 squared percentage points of MSPE per one standard deviation increase in liquidity. To think about this magnitude, consider a security priced at 55 that will eventually expire at 100, implying that the initial MSPE is 2025 points. A drop in bid-ask spreads from 4 points to 2 points is roughly one standard deviation. On average, this decrease in the spread would prevent the security priced at 55 from

converging to a price of 60—a new MSPE of 1600 points, which is 425 points lower—in the next 30 minutes. In other words, the decrease in information release can be larger than the increase in spreads.

Regression misspecification does not seem to explain this puzzling finding. Table IV shows that I cannot reject any of the six overidentifying restrictions for the excluded instrument at even the 10% level—see the Hansen (1982) *J*-statistics, which have a chi squared distribution with one degree of freedom. In addition, the regression equations are not underidentified: the Andersen (1984) likelihood-ratio statistic strongly rejects underidentification in all six cases. Finally, a comparison of the Cragg-Donaldson *F*-statistic to either the Cragg-Donaldson (1993) or the Stock and Yogo (2005) critical values shows that the two instruments are not weak.

The intuition for the IV estimates is that the change in TradeSports' market design causes opposing changes in the speed of price responses and market liquidity. These opposing changes occur in both sports markets and financial markets. Interestingly, although TradeSports' new rules increase all three measures of liquidity in financial markets, they actually decrease all three measures of liquidity in sports markets (after controlling for other factors). One possible explanation is that the optimal tick size differs in sports and financial securities (e.g., Harris (1991) and Angel (1997)). As liquidity increases in financial markets, prices respond more slowly to information; and, as liquidity decreases in sports markets, prices respond more quickly to information.

V. Exploring the Liquidity Provision Mechanism

The evidence in Section IV suggests that forecasting calibration is not better in liquid prediction markets, and that resolution is actually worse. This is surprising because liquid markets generally provide greater incentives for traders to acquire and act on information, whereas illiquid markets inhibit informed trading. In an effort to understand the two main findings in Section IV, I study the mechanism whereby prices incorporate information in prediction markets. I gather comprehensive trade-by-trade data from the TradeSports Exchange Limited archive to monitor how prices respond to information. This individual trade data indicates whether the buyer or seller initiated the trade. Because the TradeSports exchange operates as a pure limit order book, I deduce that the party initiating the trade submitted a market order (or a marketable limit order), and the counterparty initially submitted a (non-marketable) limit order.

Two key properties of the limit orders that execute shed light on the liquidity provision mechanism. First, during periods of frequent trading activity, limit orders that execute have negative expected returns. This suggests that market orders during these time periods convey valuable information that prices do not fully incorporate, perhaps because limit orders naively stand in the way of price discovery. Second, limit order traders passively buy low-priced securities and sell high-priced securities. This suggests that traders submitting market orders exploit and counteract the favorite-longshot bias, but that standing limit orders delay the correction in prices.

A. Limit and Market Order Behavior during Informative Events

First, I analyze the performance of active (market order) and passive (limit order) trades when traders act on event-relevant information. Although I do not directly observe event-relevant information, I develop a simple empirical proxy based on the clustering of trades during time periods when the event is scheduled to be in-progress. The TradeSports exchange supplies the official starting and ending event times—e.g., a typical baseball game lasts 3 hours; and a typical financial index event lasts from 9:30am Eastern time to 4:00pm Eastern time.

Table V summarizes the daily performance of market orders that trigger trades on the TradeSports exchange. I define a time period to be informative if the preceding three trades occurred within 60 seconds of each other and the event is currently in-progress. I define a time period to be uninformative if the preceding three trades did not occur within 10 minutes of each other and the event is not yet in-progress. On each day, I aggregate the properties of the trades in both groups from the perspective of the market order trader. The value-weighted statistics weight trades by the number of securities traded, whereas equal-weighted statistics weight trades equally. Table V reports the daily averages of several trading statistics for the trades during informative periods, uninformative periods, and all periods. I form these three groups separately within sports and financial securities. The goal of this analysis is to assess whether traders submitting market orders or limit orders receive and act on superior information during periods of high adverse selection.

[Insert Table V here.]

The main finding in Table V is that the trader submitting the market order consistently outperforms the limit order trader following trade clusters during the event. The magnitude of the market order trader's value-weighted daily return is greater than 1% and highly statistically significant for both sports and financial securities. These returns to expiration narrowly exceed the maximum round-trip commission costs of 0.8% for market orders.

Interestingly, when the prior three trades are not clustered within a 10-minute period and the event is not taking place, limit order traders perform significantly better than market order traders. In these situations, the returns to expiration for limit order traders are 0.77% and 1.12% for sports and financial securities, and are statistically significant at the 5% level. In addition, limit order traders incur commissions of no more than 0.4% upon expiration, and incur no commission on the initial trade after November 9, 2004. Thus, their returns to expiration are positive even after including fees.

The main result in Table V is robust to several alternative methodological approaches. The second row of Table V shows that using equal-weighted returns, rather than weighting each trade by the quantity traded, generates qualitatively similar results. The third row (EW Return w/ Lag) shows that a market order trader can use public information about past trades to exploit the market friction during informative time periods. This row computes the hypothetical equal-weighted return for a trader taking the same position in the current trade as the previous market order trader—e.g., taking the buy-side of the current trade if the previous trade was buyer-initiated. The strongly positive and significant return for this strategy shows that lags in price discovery prevent complete adjustment to the direction of trade initiation. In unreported tests, I find that

prices only partially adjust to the initiation direction of the previous five trades. In addition, sorting trades according to alternative proxies for information—e.g., using recent price volatility instead of recent trading clustering—produces results qualitatively similar to those in Table V.¹⁸

An implication of Table V is that limit order traders are systematically losing money in time periods when prices are incorporating new information. It is as though these traders pay insufficient attention to time-varying adverse selection. During informative periods, limit order traders provide liquidity in excessive amounts relative to a competitive market maker who tries to break even on each transaction. This naïve liquidity provision, or failure to withdraw liquidity, delays the response of prices to informative market order flows during the event. Still, during uninformative pre-event periods, liquidity can be beneficial for market efficiency, as shown in columns Two and Five in Table V: limit order traders prevent the overreaction of prices to uninformative market orders during these periods. Figure 3 visually depicts these patterns in returns to expiration around informative and uninformative events. The key point is that limit order trades are associated with less efficient pricing during informative event periods, but more efficient pricing during quiescent pre-event periods.

[Insert Figure 3 here.]

As an aside, the buy-sell imbalance row in Table V shows that there are more buyer-initiated trades in both sports and financial securities, mainly during the pre-event periods without trading clusters. I define buy-sell imbalance as the difference in the quantities of buyer- and seller-initiated trades divided by the total quantity of trades.

¹⁸ Row four in Table V reveals that there are far more market buy orders than market sell orders in pre-event periods without clustered trades. This imbalance could indicate that aggressive pre-event traders suffer from a framing bias, but I have no additional evidence to support this theory.

Because margin requirements and limits on trading are symmetric on the exchange, they cannot explain the asymmetry between buyers' and sellers' aggressive market orders. Thus, the positive buy-sell imbalance suggests that the preferences or prior beliefs of aggressive pre-event traders are aligned with the way the TradeSports exchange frames the events underlying securities—e.g., pre-event traders prefer to buy or speculatively buy the security based on the Dow going up when the underlying event is framed as the Dow going up, rather than down. However, the framing effect is much larger for securities with negative risk exposure, suggesting that pre-event hedging explains much of the effect in financial securities. Because TradeSports' choices of event frames are not random, it is unclear whether this framing effect is causal—but other researchers have found a causal framing effect in similar contexts (Fox, Rogers, and Tversky (1996)).

B. Limit and Market Order Behavior across Pricing Quantiles

Next, I investigate whether limit orders could impede the response of prices to aggressive market order traders who exploit market inefficiencies, such as the favorite-longshot bias. The purpose of these tests is to reconcile the lack of evidence for better forecasting calibration in liquid prediction markets in Section IV with the strong incentives for informed trading that liquidity provides. Specifically, I explore the possibility that limit order traders exhibit a systematic behavioral tendency that could offset the beneficial impact of improvements in liquidity.

To test this hypothesis, I examine how the quantities of limit buy and sell orders vary according to the prices of these binary-outcome (0 or 100 points) securities. I sort both sports and financial securities into the five 20-point pricing quantiles described earlier. For each of these ten (2x5) groups of securities, I aggregate the buy-sell imbalance and returns over all trades taking place during scheduled event times. Again, I aggregate trades on a daily basis from the perspective of the trader submitting the market order. Panels A and B in Table VI report the value-weighted buy-sell imbalance and returns for each of the pricing quantiles. I use the securities price from the previous transaction as the basis for these sorts to avoid any mechanical correlation induced by price pressure or bid-ask bounce. The results are slightly stronger if I use the contemporaneous transaction price instead.

Table VI reveals that the patterns in the buy-sell imbalances of market orders are similar to the favorite-longshot patterns in expected returns across pricing quantiles. Market order traders in both sports and financial securities are more likely to sell (longshot) securities with prices below 40 points, which have lower expected returns than other securities. Market order traders tend to buy (favorite) securities priced above 60 points (or above 40 points in sports securities), which have higher expected returns than other securities. The buy-sell imbalance results in Table VI demonstrate that limit order traders may systematically, albeit passively, exacerbate the favorite-longshot bias in prices. Executed limit orders tend to buy overpriced longshots and sell underpriced favorites in both sports and financial securities. This tendency is strong in economic terms: in sports securities purchases account for 55.6% and 43.7% of trades in pricing

quantiles 2 and 4, respectively; in financial securities, purchases account for 54.1% and 46.3% of trades in pricing quantiles 2 and 4, respectively.

Table VI also shows that market orders have significantly positive expected returns in most pricing quantiles, suggesting market orders are based on superior information. Expected returns in the middle pricing quantiles generally exceed expected returns in the extreme quantiles. One interpretation is that the securities in the middle quantiles have higher conditional variances—e.g., in an efficient market, conditional variance is $Price*(100 - Price)$, which is maximized at $Price = 50$ points—and limit order traders do not set appropriately wide bid-ask spreads to account for this fact.

Lastly, I try to determine whether the favorite-longshot pattern in limit order execution is active or passive by looking at order imbalances in the limit order book. Two facts could explain the relatively large quantities of executed limit buy orders for longshots: either limit traders submit more buy orders for longshots than for favorites, or market order traders selectively sell to the limit order traders who submit buy orders for longshots. In the first explanation, the buy-sell imbalance of unexecuted limit orders should be higher for longshots than for favorites, whereas the second explanation predicts a lower buy-sell imbalance for longshots. To test these explanations, on each day, I compute the limit order book statistics using the same procedure described earlier for trades. The only difference is that I use market depth rather than trading quantities to generate value-weighted statistics.

[Insert Table VII here.]

Most of the evidence in Table VII supports the passive explanation of limit order traders' buy-sell imbalances, particularly in the financial securities shown in Panel B.

Sell-side limit orders are conspicuously missing in increasing amounts at higher pricing quantiles, mimicking the patterns of buyer-initiated market order trades in Table VI. In financial securities in Panel B of Table VII, the difference in buy-sell imbalances across the pricing quantiles is highly statistically and economically significant—e.g., the percentage of unexecuted limit sell orders is 12.4% lower in pricing quantile 1 than in pricing quantile 5. The decreases in depth and increases in volume at higher pricing quantiles reinforce the interpretation that buyer-initiated market orders actively take sell-side liquidity away from high-priced securities. These trading patterns are less obvious in the sports securities in Panel A of Table VII, though they appear within the top three pricing quantiles. In the bottom two pricing quantiles, there is some evidence to suggest that sports traders actively submit limit orders to purchase longshots.

[Insert Figure 4 here.]

Figure 4 summarizes the results on market and limit order buy-sell imbalances in Table VI and Table VII. The solid and dashed black lines, describing financial market order and limit order imbalances across the five pricing quantiles, correlate strongly with each other (0.56). By contrast, the solid and dashed gray lines, describing sports market order and limit order imbalances, are weakly positively correlated (0.13). In both cases, the evidence is broadly consistent with the naïve and passive explanation of limit order trader behavior.¹⁹

¹⁹ In unreported tests, I show that buy-sell imbalances in unexecuted limit orders do not predict event outcomes, providing further confirmation that unexecuted limit orders are not informed.

VI. Concluding Discussion

This study's main contribution is to illustrate the complex efficiency consequences of liquidity provision, exploiting the simple structure of securities in prediction markets with observable terminal cash flows and limited exposure to systematic risks. In both sports and financial prediction markets, the calibration of prices to event probabilities does not improve with increases in liquidity; and the forecasting resolution of market prices actually decreases with increases in liquidity.

A careful examination of the behavior of limit order traders, who provide liquidity, shows that these agents do not conform to the assumptions in standard rational models. In times when market order flow is informative, limit order traders supply liquidity in excessive amounts, systematically losing money to market order traders. This finding could help to explain why prices in liquid markets are slow to incorporate event information. Furthermore, limit order traders passively trade against the favorite-longshot bias, possibly explaining why prices in liquid prediction markets are not better calibrated than prices in illiquid markets.

More broadly, these results show that liquidity providers are not always rational, and do not always attain zero profits. The findings also suggest that traders' passivity and naïveté partially explain underreaction to information in prediction markets. Given that Linnainmaa (2007) observes similar behavioral and pricing phenomena in financial markets, the same limit order mechanism could hinder efficiency in financial markets with unusually high depth.

Anecdotally, cases such as technology stocks during the Internet boom raise doubts about whether pricing in liquid markets is always more efficient than pricing in illiquid markets. Further exploration of the limit order mechanism identified here and in Linnainmaa (2007) could yield insights into these puzzling cases. The effect of excessive liquidity provision on equilibrium prices depends critically on how arbitrageurs respond to liquidity. Because liquidity gives arbitrageurs the option to enter and exit their trading positions easily, it could induce them to take short-run outlooks and conduct trades that actually exacerbate mispricing (e.g., DeLong et al. (1990b) and Brunnermeier and Nagel (2004)). In other cases in which arbitrageurs have incentives to correct mispricing, they may be unable to do so because they have less capital than naïve liquidity providers. Theoretical models of the limit order mechanism could make new predictions about the situations in which liquid markets are more or less efficient than illiquid markets.

References

- Ali, Mukhtar M., 1977, Probability and utility estimates for racetrack bettors, *Journal of Political Economy* 85, 803-815.
- Anderson, T.W., 1984. *Introduction to Multivariate Statistical Analysis*. 2nd ed. New York: John Wiley & Sons.
- Angel, James J., 1997, Tick size, share price, and stock splits, *Journal of Finance* 52, 655-681.
- Baker, Malcolm, and Jeremy Stein, 2004, Market liquidity as a sentiment indicator, *Journal of Financial Markets* 7, 271-299.
- Barberis, Nicholas, and Ming Huang, 2007, Stocks as lotteries: the implications of probability weighting for security prices, NBER Working Paper #12936.
- Berg, Joyce, Robert Forsythe, Forrest Nelson and Thomas Rietz, 2003, Results from a Dozen Years of Election Futures Markets Research, University of Iowa working paper.
- Bloomfield, Robert, Maureen O'Hara, and Gideon Saar, 2007, How noise trading affects markets: an experimental analysis, Cornell University Working Paper.
- Brav, Alon, and John B. Heaton, 1996, Explaining the underperformance of initial public offerings: a cumulative prospect utility approach, Duke University Working Paper.
- Brunnermeier, Markus K., and Stefan Nagel, 2004, Hedge funds and the technology bubble, *Journal of Finance* 59, 2013-2040.
- Camerer, Colin and Teck H. Ho, 1994, Violations of the betweenness axiom and nonlinearity in probability, *Journal of Risk and Uncertainty* 8, 167-196.
- Cason, Timothy and Daniel Friedman, 1996, Price formation in double auction markets, *Journal of Economic Dynamics and Control* 20, 1307-1337.
- Chen, Kay-Yut and Charles Plott, 2002, Information Aggregation Mechanisms: Concepts, Design, and Implementation for a Sales Forecasting Problem, CalTech Working Paper No. 1131.
- Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam, 2006, Liquidity and market efficiency, Emory University Working Paper.

- Chordia, Tarun, Amit Goyal, Gil Sadka, Ronnie Sadka, and Lakshmanan Shivakumar, 2007, Liquidity and the post-earnings-announcement drift, Emory University Working Paper.
- Cowgill, Bo, Justin Wolfers, and Eric Zitzewitz, 2007, Prediction Markets Inside the Firm: Evidence from Google, University of Pennsylvania Working Paper.
- Cragg, J.G., and S.G. Donaldson, 1993, Testing identifiability and specification in instrumental variables models, *Econometric Theory* 9, 222-240.
- DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990a, Noise trader risk in financial markets, *Journal of Political Economy* 98, 703-738.
- DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990b, Positive feedback investment strategies and destabilizing rational speculation, *Journal of Finance* 45, 379-395.
- Duffie, Darrell, Garleanu, Nicolae, and Lasse Heje Pedersen, 2005, Over-the-counter markets, *Econometrica* 73, 1815-1847.
- Fox, Craig R., Brett A. Rogers, and Amos Tversky, 1996, Options traders exhibit subadditive decision weights, *Journal of Risk and Uncertainty* 13, 5-17.
- Froot, Kenneth A., 1989, Consistent covariance matrix estimation with cross-sectional dependence and heteroskedasticity in financial data, *Journal of Financial and Quantitative Analysis* 24, 333-355.
- Gjerstad, Steven, and John Dickhaut, 1998, Price formation in double auctions, *Games and Economic Behavior* 22, 1-29.
- Glosten, Lawrence, and Paul Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71-100.
- 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, 119-137.
- Hahn, Robert W., and Paul C. Tetlock, 2005, Using information markets to improve public decision making, *Harvard Journal of Law and Public Policy* 29, 214-289.
- Hansen, Lars, 1982, Large sample properties of the generalized method of moments estimators, *Econometrica* 50, 1029-1054.

- Hanson, Robin, 2002, Decision markets, in *Entrepreneurial Economics: Bright Ideas from the Dismal Science*, Oxford University Press.
- Harris, Lawrence, 1991, Stock price clustering and discreteness, *Review of Financial Studies* 4, 389-415.
- Jullien, Bruno, and Bernard Selanie, 2000, Estimating preferences under risk: the case of racetrack bettors, *Journal of Political Economy* 108, 503-530.
- Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: an analysis of decision under risk, *Econometrica* 47, 263-291.
- Kaniel, Ron, Shuming Liu, Gideon Saar, and Sheridan Titman, 2007, Individual investor trading around earnings announcements, Duke University working paper.
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315-1336.
- Linnainmaa, Juhani, 2007, The limit order effect, University of Chicago working paper.
- Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703-708.
- Noussair, Charles, Stephane Robin, and Bernard Ruffieux, 1998, The effect of transactions costs on double auction markets, *Journal of Economic Behavior and Organization* 36, 221-233.
- O'Hara, Maureen, 1998, *Market Microstructure Theory*, John Wiley & Sons, Inc.
- Prelec, Drazen, 1998, The probability weighting function, *Econometrica* 66, 497-527.
- Roll, Richard, 1984, Orange juice and weather, *American Economic Review* 74, 861-880.
- Rubinstein, Mark, 1985, Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active CBOE option classes from August 23, 1976 through August 31, 1978, *Journal of Finance* 40, 455-480.
- Sadka, Ronnie, and Anna Scherbina, 2007, Analyst disagreement, mispricing, and liquidity, *Journal of Finance* 62, 2367-2403.
- Satterthwaite, Mark A., and Steven R. Williams, 2002, The optimality of a simple market mechanism, *Econometrica* 70, 1841-1863.
- Shleifer, Andrei, and Robert W. Vishy, 1997, The limits of arbitrage, *Journal of Finance* 52, 35-55.

- Stock, James H., and Motohiro Yogo, 2005, Testing for weak instruments in linear IV regression. In Donald W.K. Andrews and James H. Stock, eds. *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*. Cambridge: Cambridge University Press.
- Summers, Lawrence H., 1986, Does the stock market rationally reflect fundamental values? *Journal of Finance* 41, 591-601.
- Sunstein, Cass R., 2006, *Infotopia: How Many Minds Produce Knowledge?* Oxford University Press.
- Tetlock, Paul C., 2004, How efficient are information markets? Evidence from an online exchange, Yale University Working Paper.
- Tetlock, Paul C., and Robert W. Hahn, 2007, Optimal liquidity provision for decision makers, Yale University Working Paper.
- Tetlock, Paul C., 2007, All the news that's fit to reprint: Do investors react to stale information? Yale University Working Paper.
- Thaler, Richard, and William Ziemba, 1988, Anomalies: parimutuel betting markets: racetracks and lotteries, *Journal of Economic Perspectives* 2, 161-174.
- Tversky, Amos, and Daniel Kahneman, 1992, Advances in prospect theory: cumulative representation of uncertainty, *Journal of Risk and Uncertainty* 5, 297-323.
- Wald, Abraham, 1943, Tests of statistical hypotheses concerning several parameters when the number of observations is large, *Transactions of the American Mathematical Society*, 54, 426-482.
- White, Halbert, 1980, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* 48, 817-838.
- Wolfers, Justin, and Eric Zitzewitz, 2004, Prediction markets, *Journal of Economic Perspectives* 18, 107-126.
- Wu, George, and Richard Gonzalez, 1996, Curvature of the probability weighting function, *Management Science* 42, 1676-1690.
- Wurgler, Jeffrey, and Ekaterina Zhuravskaya, 2002, Does arbitrage flatten demand curves for stocks? *Journal of Business* 75, 583-608.
- Zitzewitz, Eric, 2006, Price discovery among the punters: using new financial betting markets to predict intraday volatility, Dartmouth University Working Paper.

Table I: Summary of Liquidity Statistics for Prediction Market Securities

All securities are based on one-day binary outcome events: either sports games or levels of financial indexes. At 30-minute intervals, a web crawler records statistics for all securities in which active trading occurs from March 17, 2003 to October 23, 2006. Table I summarizes three measures of securities market liquidity: quoted bid-ask spreads, market depth, and cumulative trading volume—see text for definitions. The table reports the pair-wise correlations between the logarithms of quoted spreads (*LnSpread*), market depth (*LnDepth*), and trading volume (*LnVolm*) for all securities, sports securities and financial securities. The means for each liquidity measure are equal-weighted (EW) across all order book observations.

	Securities Included	All	Sports	Financial
Correlations	<i>LnSpread-LnDepth</i>	-0.434	-0.334	-0.208
	<i>LnSpread-LnVolm</i>	-0.012	-0.101	-0.205
	<i>LnDepth-LnVolm</i>	0.124	0.177	0.414
Means	EW Spread	2.51	2.08	4.37
	EW Depth	537	605	238
	EW Volume	455	405	674
	# of Observations	247613	201492	46121
	# of Securities	46670	35089	11581

Table II: Returns to Expiration for Sports and Financial Securities

This table reports the results from three ordinary least squares (OLS) regressions of securities' returns to expiration on five dummy variables, representing five equally-spaced pricing intervals: (0,20), [20,40), [40,60), [60,80), and [80,100) points.. The three regressions include observations of the returns and prices of securities based on one-day sports, financial, and all events recorded at 30-minute intervals in which there is active trading. I compute returns to expiration as the payoff at expiration, 0 or 100 points, minus the bid-ask midpoint divided by 100 points. The small probabilities row displays the magnitude and significance of the average coefficient on Price1 and Price2. The large probabilities row displays the magnitude and significance of the average coefficient on Price3, Price4, and Price5. The large minus small row reports the magnitude and significance of the difference in these two averages. Standard errors clustered for securities that expire on the same calendar day appear in parentheses (Froot (1989)). The symbols *, **, and *** denote significance at the 10%, 5%, and 1%, levels, respectively.

	Sports	Financial	All
0 < Price < 20	-1.65** (0.67)	-0.30 (0.63)	-0.55 (0.53)
20 ≤ Price < 40	-0.64 (1.25)	-0.23 (1.12)	-0.46 (0.87)
40 ≤ Price < 60	-0.65 (0.53)	2.47* (1.41)	-0.48 (0.50)
60 ≤ Price < 80	0.82 (0.88)	2.81** (1.42)	1.09 (0.78)
80 ≤ Price < 100	1.73** (0.81)	1.34 (0.82)	1.60*** (0.59)
Small Probabilities	-1.15* (0.75)	-0.27 (0.79)	-0.50 (0.59)
Large Probabilities	0.63 (0.52)	2.21** (1.02)	0.74* (0.44)
Large – Small	1.78** (0.83)	2.47** (1.03)	1.24* (0.68)
R^2	0.0003	0.0016	0.0003
Expiration Days	1110	707	1122
Observations	201492	46121	247613

Table III: Returns to Expiration for Sports and Financial Securities Sorted by Bid-Ask Spread

This table reports the results from eight (2x4) ordinary least squares (OLS) regressions of securities' returns to expiration on five dummy variables (Price1 to Price5), representing five equally-spaced pricing intervals: (0,20), [20,40), [40,60), [60,80), and [80,100). The eight regressions include two sets of regressions for securities based on one-day binary outcome sports and financial events. Each set includes regressions for groups of securities sorted into four quantiles by their bid-ask spreads (S). The small, large, and small minus large probabilities rows display the magnitude and significance of the average coefficients on Price1 and Price2, Price3, Price4, and Price5, and their difference, respectively. Standard errors clustered for securities that expire on the same calendar day appear in parentheses (Froot (1989)). The symbols *, **, and *** denote significance at the 10%, 5%, and 1%, levels, respectively.

Bid-Ask Spread (S)	Sports Securities				Financial Securities			
	$S \leq 1$	$1 < S \leq 3$	$3 < S \leq 5$	$S > 5$	$S \leq 1$	$1 < S \leq 3$	$3 < S \leq 5$	$S > 5$
0 < Price < 20	-3.42*** (1.26)	-1.15 (1.00)	-1.35 (1.01)	-0.91 (1.40)	-0.09 (1.34)	-0.61 (0.76)	-0.19 (0.73)	-0.31 (0.77)
20 ≤ Price < 40	-2.22 (3.67)	-0.58 (1.42)	1.06 (1.52)	-2.64 (1.62)	-5.05 (4.64)	2.05 (1.92)	0.71 (1.50)	-1.15 (1.14)
40 ≤ Price < 60	-1.06 (1.16)	-0.45 (0.58)	-0.84 (0.55)	-0.89 (0.88)	5.87 (5.43)	2.26 (2.35)	-0.47 (1.81)	3.98*** (1.50)
60 ≤ Price < 80	0.92 (1.77)	0.88 (0.95)	0.19 (1.05)	1.94 (1.38)	2.41 (4.93)	5.15** (2.17)	0.97 (1.87)	3.33** (1.53)
80 ≤ Price < 100	2.79 (1.87)	1.39 (1.07)	1.45 (0.96)	2.28** (0.98)	2.33** (0.92)	1.23 (1.21)	1.41 (0.88)	1.24 (1.05)
Small Probabilities	-2.82 (1.92)	-0.87 (0.93)	-0.15 (0.96)	-1.77 (1.12)	-2.57 (2.47)	0.72 (1.16)	0.26 (1.00)	-0.73 (0.83)
Large Probabilities	0.88 (1.06)	0.61 (0.59)	0.27 (0.57)	1.11 (0.68)	3.54 (2.42)	2.88** (1.41)	0.64 (1.23)	2.85*** (1.10)
Large – Small	3.70* (2.20)	1.47 (1.03)	0.41 (1.03)	2.89** (1.24)	6.11* (3.45)	2.16 (1.51)	0.37 (1.22)	3.58*** (1.18)
R^2	0.0008	0.0002	0.0003	0.0011	0.0096	0.0038	0.0004	0.0033
Expiration Days	634	1099	1100	998	351	616	669	696
Observations	30818	112839	43380	14455	839	6218	15811	23253

Table IV: The Effect of Liquidity on the Convergence of Prices to Outcomes

This table reports the results from six two-stage least squares regressions of the reduction in squared percentage returns to expiration on liquidity and several control variables. I compute the reduction in squared returns over the 30-minute period following liquidity and control variable measurement. The key independent variables are the logarithms of bid-ask spreads (*LnSpread*), market depth (*LnDepth*), and trading volume (*LnVolm*). I use two dummy variables to instrument for changes in liquidity using TradeSports market design changes on September 1, 2004 and November 9, 2004. Each regression uses observations from March 1, 2004 to February 9, 2004. I include a control variable for an event's expected information release through expiration, *ExpInfo*, which is $(100 - Price) * Price$. All regressions also include controls for six event time dummies, an event time trend, and a date trend. The Anderson, Cragg-Donaldson, and Hansen test statistics for underidentification, weak identification, and overidentification appear below, along with appropriate *p*-values and critical values. Standard errors clustered for securities that expire on the same calendar day appear in parentheses (Froot (1989)). The symbols *, **, and *** denote significance at the 10%, 5%, and 1%, levels, respectively.

	MSPE Reduction in Sports			MSPE Reduction in Financial		
<i>LnSpread</i>	310*** (90)			978** (445)		
<i>LnDepth</i>	-284*** (85)			-465* (242)		
<i>LnVolm</i>	-189*** (47)			-188* (108)		
<i>ExpInfo</i>	0.23*** (0.04)	0.25*** (0.04)	0.13*** (0.04)	0.13* (0.07)	0.23*** (0.04)	0.24*** (0.04)
Anderson <i>LR</i>	711.8***	541.3***	405.4***	42.8***	125.5***	151.3***
Anderson <i>p</i> -value	0.000	0.000	0.000	0.000	0.000	0.000
Cragg-Donaldson <i>F</i>	364.4***	275.5***	205.3***	21.5***	63.7***	77.1***
<i>F</i> for 10% IV Size	19.9	19.9	19.9	19.9	19.9	19.9
Hansen <i>J</i>	0.15	3.07*	0.02	0.65	1.16	2.52
Hansen <i>p</i> -value	0.694	0.080	0.884	0.420	0.282	0.112
<i>R</i> ²	0.231	0.226	0.225	0.010	0.075	0.090
Expiration Days	232	232	232	163	163	163
Observations	14644	14644	14644	3294	3294	3294

Table V: Market Order Performance after Informative and Uninformative Periods

Table V summarizes the average daily returns to expiration of market orders that trigger trades on the TradeSports exchange. The columns report these statistics for the trades during “Info” periods, during “No Info” periods, and all trades. I form these three groups separately within sports and financial securities. The “Info” trades are those occurring during the event and after three trades occurring within 60 seconds of each other. The “No Info” trades are those occurring before the event and after three trades that did not occur within 10 minutes of each other. On each day, I aggregate the properties of the trades in all groups from the perspective of the market order trader. The value-weighted (VW) statistics weight trades by the number of securities traded, whereas equal-weighted (EW) statistics weight trades equally. The “EW Return w/ Lag” row computes the hypothetical equal-weighted return to expiration for a trader taking the same position in the current trade as the previous market order trader. I define the VW buy-sell imbalance as the quantity of buyer-initiated trades minus the quantity of seller-initiated trades divided by the total quantity of trades. Standard errors in parentheses are robust to heteroskedasticity and autocorrelation up to five lags (Newey and West (1987)).

	Sports			Financial		
	Info	No Info	All	Info	No Info	All
VW Return	1.20*** (0.28)	-0.77** (0.34)	0.38*** (0.12)	1.41*** (0.30)	-1.12** (0.56)	0.35*** (0.10)
EW Return	0.48** (0.22)	-0.54*** (0.18)	0.07 (0.08)	0.94*** (0.17)	-0.98*** (0.34)	0.29*** (0.06)
EW Return w/ Lag	0.91*** (0.21)	-0.16 (0.18)	0.26*** (0.09)	0.53*** (0.18)	-0.12 (0.35)	0.34*** (0.06)
VW Buy-Sell Imbal	0.005 (0.008)	0.217*** (0.008)	0.067*** (0.004)	0.013 (0.009)	0.062*** (0.015)	-0.001 (0.003)
# Trades	308	208	1485	303	48	1429
Days	1518	1540	1546	1080	1051	1089

Table VI: Market Order Imbalances and Returns Sorted by Pricing Quantile

Table VI summarizes the average daily returns to expiration and buy-sell imbalances of market orders that trigger trades on the TradeSports exchange. I sort both sports (Panel A) and financial (Panel B) securities into five equally-spaced pricing quantiles—(0,20), [20,40), [40,60), [60,80), and [80,100)—using the securities price from the previous transaction. For each of these ten (2x5) groups of securities, I aggregate each day’s buy-sell imbalance and returns for all trades taking place during the scheduled event time. I compute the properties of the trades in the ten groups from the perspective of the market order trader. The value-weighted (VW) statistics weight trades by the number of securities traded. Table VI reports these daily averages of trading statistics for the ten groups of securities. I define the VW buy-sell imbalance as the quantity of buyer-initiated trades minus the quantity of seller-initiated trades divided by the total quantity of trades. Standard errors in parentheses are robust to heteroskedasticity (White (1980)).

Panel A: Sports Securities					
	$0 < P < 20$	$20 \leq P < 40$	$40 \leq P < 60$	$60 \leq P < 80$	$80 \leq P < 100$
VW Buy-Sell Imbal	-0.167*** (0.009)	-0.112*** (0.008)	0.059*** (0.005)	0.127*** (0.007)	0.166*** (0.008)
VW Return	-0.04 (0.31)	1.44*** (0.35)	0.64*** (0.23)	0.89*** (0.29)	0.48* (0.26)
# Trades	153	178	412	242	160
# Days	1511	1531	1543	1539	1515
Panel B: Financial Securities					
	$0 < P < 20$	$20 \leq P < 40$	$40 \leq P < 60$	$60 \leq P < 80$	$80 \leq P < 100$
VW Buy-Sell Imbal	-0.016* (0.009)	-0.081*** (0.006)	0.000 (0.006)	0.074*** (0.008)	0.054*** (0.009)
VW Return	0.03 (0.25)	0.21 (0.28)	0.74*** (0.25)	0.76*** (0.27)	0.40 (0.31)
# of Trades	298	302	298	228	198
# of Days	1082	1081	1080	1075	1059

Table VII: Unexecuted Limit Orders Sorted by Pricing Quantile

Table VII reports average daily statistics from the TradeSports order book recorded at 30-minute intervals. I sort both sports (Panel A) and financial (Panel B) securities into five equally-spaced pricing quantiles—(0,20), [20,40), [40,60), [60,80), and [80,100)—using the securities price from the previous transaction. For each of these ten (2x5) groups of securities, I aggregate each day's buy-sell imbalance and order book liquidity across all data recording periods. The value-weighted (VW) statistics weight observations by order book depth, whereas the equal-weighted (EW) statistics weight observations equally. See text for details on buy-sell imbalance, depth, volume, and spread definitions. Standard errors in parentheses are robust to heteroskedasticity (White (1980)).

Panel A: Sports Securities					
	0 < P < 20	20 ≤ P < 40	40 ≤ P < 60	60 ≤ P < 80	80 ≤ P < 100
VW Buy-Sell Imbal	0.045* (0.026)	0.085*** (0.012)	0.012*** (0.004)	0.034*** (0.007)	0.117*** (0.014)
EW Buy-Sell Imbal	0.051** (0.026)	0.087*** (0.011)	0.019*** (0.004)	0.047*** (0.006)	0.138*** (0.014)
EW Depth	394*** (13)	476*** (16)	641*** (9)	616*** (10)	475*** (10)
EW Volume	3393*** (168)	1683*** (95)	422*** (17)	640*** (36)	1629*** (82)
VW Bid-Ask Spreads	2.67*** (0.07)	2.56*** (0.05)	1.93*** (0.02)	2.00*** (0.03)	2.48*** (0.04)
EW Bid-Ask Spreads	2.75*** (0.07)	2.74*** (0.05)	2.13*** (0.02)	2.20*** (0.03)	2.62*** (0.04)
# of Days	490	784	1089	1001	745
Panel B: Financial Securities					
	0 < P < 20	20 ≤ P < 40	40 ≤ P < 60	60 ≤ P < 80	80 ≤ P < 100
VW Buy-Sell Imbal	0.030** (0.016)	0.117*** (0.010)	0.134*** (0.008)	0.158*** (0.011)	0.278*** (0.018)
EW Buy-Sell Imbal	0.000 (0.017)	0.130*** (0.010)	0.161*** (0.009)	0.185*** (0.011)	0.303*** (0.017)
EW Past Depth	279*** (7)	278*** (7)	263*** (6)	250*** (6)	238*** (6)
EW Past Volume	1009*** (55)	990*** (37)	1103*** (50)	1106*** (48)	1335*** (65)
VW Bid-Ask Spreads	3.95*** (0.05)	4.32*** (0.05)	4.29*** (0.05)	4.29*** (0.06)	3.99*** (0.06)
EW Bid-Ask Spreads	4.08*** (0.05)	4.49*** (0.05)	4.53*** (0.05)	4.51*** (0.05)	4.11*** (0.06)
# of Days	516	535	581	585	522

Figure 1: The Favorite-Longshot Bias in Sports and Financial Securities

This figure depicts the estimated returns to expiration of securities based on one-day events for five equally-spaced pricing quantiles: (0,20), [20,40), [40,60), [60,80), and [80,100). The three series in the figure depict the returns of securities based on sports events, financial events, and all events combined. Thus, the figure plots the three sets of coefficient estimates on the five pricing category coefficients shown in the three columns in Table II (see table for construction).




Figure 2: The Effect of Liquidity on Returns to Expiration

This figure depicts the estimated favorite-longshot biases in returns to expiration of sports and financial securities with differing degrees of liquidity as measured by three proxies—bid-ask spreads, market depth, and trading volume. I measure favorite-longshot biases for securities in each of the four bid-ask spread quantiles, where the bias is the expected returns on favorites minus the expected returns on longshots—see Table III and the text for details. I perform analogous tests for securities within each market depth and trading volume quantile.

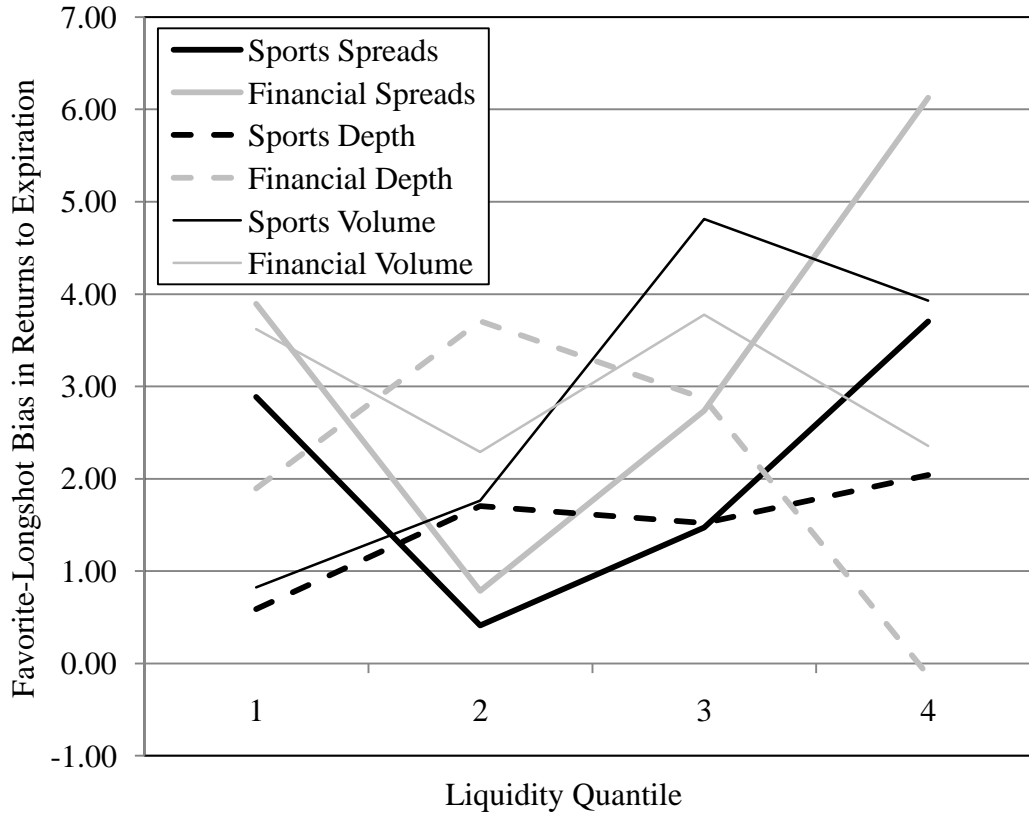


Figure 3: Market Order Performance after Informative and Uninformative Periods

This figure summarizes the average daily returns to expiration for market orders that trigger trades on the TradeSports exchange. Figure 3 reports these statistics for the trades during “Info” periods, during “No Info” periods, and all trades. I form these three groups separately within sports and financial securities. The “Info” trades are those occurring during the event and after three trades occurring within 60 seconds of each other. The “No Info” trades are those occurring before the event and after three trades that did not occur within 10 minutes of each other. On each day, I aggregate the properties of the trades in all groups from the perspective of the market order trader. The value-weighted (VW) return series weights trades by the number of securities traded, whereas equal-weighted (EW) return series weights trades equally. The “EW Return w/ Lag” data series computes the hypothetical equal-weighted return to expiration for a trader taking the same position in the current trade as the previous market order trader.




Figure 4: Buy-Sell Order Imbalances Sorted by Pricing Quantiles

This figure depicts value-weighted buy-sell imbalances of four different order types sorted by the prices of the securities. For limit orders, observations come from the TradeSports order book recorded at 30-minute intervals, and value weights are based on market depth. For market orders, observations come from the TradeSports transaction archive, and value weights are based on the quantity of securities traded. I sort both sports and financial securities into five equally-spaced pricing quantiles—(0,20), [20,40), [40,60), [60,80), and [80,100)—using the securities price from the previous order book recording or previous transaction. For each of these ten (2x5) groups of securities, I aggregate each day's market or limit order buy-sell imbalance. See Table VI and Table VII for further details.

