Short sales constraints and market quality: An experimental approach

Juan Cabrera, Ph.D., CFA®[a,*]
Ramapo College of New Jersey

Eleni Gousgounis, Ph.D.[b]
Stevens Institute of Technology

Abstract

We use an experimental setting to test the effect of short sale constraints on key aspects of market quality: market efficiency, price discovery, and liquidity. We find that when short sale constraints bind, securities are overpriced, trading at a premium, which is often higher than the difference between the view of the most optimistic investor and the fundamental value. Opinion dispersion contributes to this overpricing; a result that is more pronounced at the time that opinion dispersion is induced. As time progresses, traders update their valuations to incorporate information from trading. As a result, opinion dispersion narrows and its effect on overpricing dissipates. We also observe a slower speed of price adjustment when short sale constraints bind, yet prices exhibit persistence never converging to their fundamental values. Finally, short sale constraints are associated with lower trading volume and higher bid-ask spreads, with the latter effect being significant only when opinion dispersion is high.

First Draft: January 5th, 2015

JEL Classification Codes: G12, G14, G18, D40, D82, D83

Keywords: short sale constraints, opinion dispersion, market efficiency, price discovery, liquidity, experimental

---

* Please address correspondence to Juan F. Cabrera
[a] Juan F. Cabrera is an Associate Professor of Finance at the Anisfield School of Business, Ramapo College of New Jersey, Mahwah, NJ, Phone: (201) 684-7354, E-mail: jcabrera@ramapo.edu
[b] Eleni Gousgounis is an Assistant Professor of Finance at Howe School of Technology Management, Stevens Institute of Technology, Hoboken, NJ, Phone: (201) 216-8567, Fax: (201)-216-8355, E-mail: egousgou@stevens.edu

Short sales constraints and market quality: An experimental approach

1. Introduction

Short sale constraints have been in the spotlight since the financial crisis of 2007-2009, when a series of such restricting rules were applied worldwide. In the case of the U.S., the SEC placed a three-week long short sale ban in September 2008, covering most financial stocks in an effort to prevent a sharp decline in their prices. While the effectiveness of the ban in supporting prices remains unclear, it is suspected that the short sale ban led to a liquidity squeeze. The challenge in studying short sale bans is that they are typically applied at times of rapidly changing market conditions, which makes it difficult to isolate the effects of the ban on distinct measures of market quality. For example, Beber and Pagano (2013), who study the effects of short sale constraints around the world during the 2007-2009 financial crisis, find that the U.S. market is unique in that it is the only market where short sale constraints were associated with overpricing. However, as underlined by the authors, it is difficult to discern whether this effect was due to the constraints or the announcement of bank bailouts that occurred around the same time period.

Our paper utilizes an experimental setting to study the effect of short sale constraints on key aspects of market quality: market efficiency, price discovery, and liquidity. The benefit of the experimental setting is that it allows us to control for suspected drivers of market inefficiencies (i.e. opinion dispersion) ex ante, isolate the effects of short sale constraints, and observe the exact market mechanism leading to these inefficiencies. Moreover, the experimental setting allows us to accurately measure variables that cannot be observed in practice and test theoretical hypotheses, which are difficult to test empirically. For example, the empirical literature on the effect opinion dispersion on market efficiency in a market with short sale constraints, typically, relies on proxies to measure opinion dispersion and mispricing. In contrast, there is no such need in the experimental setting: since we deliberately induce opinion dispersion, we can also effectively control for it. Furthermore, we predetermine the fundamental value of the traded assets, allowing us to measure mispricing accurately.

Our paper relates closely to two experimental papers on short sale constraints (Haruvy and Noussair, 2006; Fellner and Theissen, 2014), yet neither of these papers analyzes all facets of market quality. In more detail, Haruvy and Noussair (2006) study short sale constraints in an experimental setting, focusing on the effect of short sales on price discovery and the creation of bubbles and busts. They find that restricted short selling capacity causes bubbles[1], whose magnitude and duration are reduced as short selling capacity increases. However, even then, prices do not seem to track fundamentals. Fellner and Theissen (2014) also find that stocks are overvalued when short selling constraints bind. They test whether this overpricing increases with opinion dispersion, but they do not find a significant effect. However, their explanation is rooted to their design: they highlight the difficulty in inducing opinion by endowing traders with asymmetric information while also abiding to the established principles of experimental economics. Fellner and Theissen (2014) also find that short selling constraints reduce liquidity, expressed in higher quoted spreads and lower trading volume.

[1] They define bubbles as a sustained episode of high transaction volume at prices that greatly exceed fundamental values, usually followed by a crash.

Our paper is the first experimental paper to study the effect of short sale constraints on key features of market quality, including market efficiency, price discovery, and market liquidity. It also complements the extant literature in that it successfully induces opinion dispersion in an experimental setting by providing asymmetric information, without violating the commonly accepted experimental principles. Therefore, we are able to evaluate the effect of short sale constraints on market quality for various levels of opinion dispersion in both a static (at the time opinion dispersion is induced) and a dynamic setting (as time progresses). We are able to gauge to what extent opinion dispersion is the main driver of market inefficiencies under short sale constraints. Finally, the experimental setting allows us to capture the view of the most optimistic investor and test the theoretical argument of Harrison and Kreps (1978), according to which short sale constraints can cause prices to soar and even surpass the view of the most optimistic investor.

In more detail, we first test whether short sale constraints cause overpricing, as predicted by Miller (1977), and similar to others before us, we find evidence of overpricing. Second, we examine whether the overpricing is extreme (i.e. price bubble) in that the price surpasses not just the asset's fundamental value, but also the view of the most optimistic investor, consistent with the predictions of the theoretical model of Harrison and Kreps (1978). Our results show that indeed price bubbles occur often. This is the first paper to our knowledge test this hypothesis by comparing prices to the view of the most optimistic investor. Third, we explore the role of opinion dispersion on overpricing. We induce opinion dispersion by providing our participants with different pieces of non-conflicting information using overlapping information ranges and the behavioral concepts of framing and anchoring. We do so by also abiding to the generally acceptable experimental principles. Moreover, we are the first, to our knowledge, to experimentally explore the role of opinion dispersion on overpricing in a dynamic setting, where constrained traders update their estimates based on information they extract from actual trading in the market. Our findings support the hypothesis that higher opinion dispersion increases overpricing when short sale constraints bind. The results are stronger at the time that opinion dispersion is induced, which in our experiment coincides with the opening time interval of our market. As time proceeds, the difference in mispricing between markets with different levels of originally induced opinion dispersion narrows. Our explanation is that our original inducement of opinion dispersion becomes less relevant as time proceeds, because traders extract additional information from trading and update their valuations accordingly. Fourth, we are testing whether short sale constraints are associated with a slower speed of adjustment in the price discovery process and whether prices eventually return to their fair value, as predicted by Diamond and Verrechia (1987). While we find that short sale constraints are associated with a slower speed of price adjustment, prices do not converge to their fundamental value. Finally, similar to Fellner and Theissen (2014) we study the effect of short sale constraints on liquidity, measured by percentage spreads and trading volume. We also investigate the interaction between liquidity and opinion dispersion when short selling is not allowed. We find that spreads are higher when short sale constraints bind, but only when opinion dispersion is also high. Trading volume is significantly less when short sale constraints bind, irrespective of the level of opinion dispersion.

The rest of the paper proceeds as follows. Section 2 reviews the relevant literature and develops our testable hypotheses. Section 3 presents our experimental design, while Section 4 discusses our statistical methodology. Section 5 presents our results and Section 6 concludes the paper.

2. Hypotheses

Our objective is to use our experimental setting to test the effect of short sale constraints on all aspects of market quality. More specifically, we investigate the effect of a short sale ban on market efficiency, price discovery, and market liquidity.

The theoretical literature on the effects of short sale constraints on market efficiency starts with the seminal work of Miller (1977), which shows that overpricing is inevitable in a market with heterogeneous expectations and short sale constraints. He argues that short sale constraints can inflate market prices, as bearish investors cannot act on their market views. Figlewski (1981) adds that price inefficiencies will arise, even if investors correct their expectations at the market level: short sale constrained assets with very adverse information will be overpriced, whereas those with little adverse information will be underpriced.

However, the literature has also presented theoretical reasons supporting the counterargument: the application of short sale constraints may not cause overpricing. Jarrow (1980) extends Miller's model to include multiple assets and shows that there are cases that short sale constrained assets would exhibit underpricing. Diamond and Verrecchia (1987) argue that overpricing should quickly resolve as investors revise their expectations to account for overstated prices due to the ban. Hong and Stein (2003) develop a model where prices can decline rapidly as optimistic investors revise their expectations. These revisions account for the views of pessimistic investors who are not entering the market as support buyers.

While many empirical papers (i.e., Desai, Ramesh, Thiagarajan and Balchandran, 2002; Asquith, Pathak and Ritter, 2005) provide evidence in support of the overpricing hypothesis, there has also been some evidence against it. Beber and Pagano (2013) investigate the effect of the recent short sale constraints applied in numerous countries during the period of 2007-2009. They find no evidence of price support in most countries. U.S. appears to be the only exception, a result that could also be attributed to the concurrent announcement of TARP funding to financial institutions.

Therefore, our first hypothesis addresses the question of overpricing in a market with short sale constraints. More specifically,

Hypothesis 1: In a market with short sale constraints there is overpricing.

Building on the overpricing hypothesis, there are theoretical arguments supporting the notion that in a market with short sale constraints, overpricing will be extreme. In fact, Harrison and Kreps (1978) assert that short sale constraints can contribute to the creation of speculative bubbles. They present a model of a short sale constrained market with heterogeneous investors often willing to pay a price even higher than the most optimistic investor's view. This hypothesis has been challenging to test empirically given the practical limitations in capturing the view of the most optimistic investor. The small number of traders and our ability to interact with them allows us to record their valuations, making the experimental setting ideal for testing this hypothesis. Therefore, the Harrison and Kreps (1978) argument comprises our second hypothesis:

Hypothesis 2: In a market with short sale constraints we observe bubbles with price often surpassing the view of the most optimistic investor.

According to Miller (1977), the main driver of overpricing in a market with short sale constraints is the degree of opinion dispersion among market participants: when opinion dispersion is high, the prohibition of short sales leads to overpricing, as pessimistic investors cannot take short positions and remain out of the market. Chen, Hong, and Stein (2002) extend Miller (1977) work by modeling opinions as uniformly distributed. They show that opinion dispersion leads to overpricing if short sale constraints bind at least for some investors and if opinion dispersion is above a minimum threshold.

While there is a vast empirical literature providing support of the association of opinion dispersion and overpricing in a market with short selling constraints, most studies rely on statistical proxies for overpricing and opinion dispersion of uncertain accuracy. Overpricing is frequently measured by negative subsequent returns, based on the assumption that overpricing will be always followed by a price correction. The employed timeframe varies from study to study: Goetzmann and Massa (2005) employ an overnight framework; Boehme, Danielsen, and Sorescu (2006) use a monthly timeframe; whereas Chen, Hong, and Stein (2002) and Nagel (2005) find a negative effect of opinion dispersion on quarterly subsequent returns. Frequently used proxies of opinion dispersion include market turnover (Miller, 1977; Goetzmann and Massa, 2005; Nagel, 2005; Boehme, Danielsen and Sorescu, 2006), dispersion of analysts' forecasts (Diether, Malloy, and Scherbina, 2002; Boehme, Danielsen, and Sorescu, 2006), historical equity return volatility (Goetzmann and Massa, 2005; Nagel, 2005; Boehme, Danielsen, and Sorescu, 2006), and heterogeneity of trade among investor classes (Goetzmann and Massa, 2005).

The experimental setting allows us to test the role of opinion dispersion as a driver of overpricing in a market with short sale constraints without having to rely on proxies to measure overpricing and opinion dispersion. We can measure overpricing accurately as the difference between the market price and the predetermined "true" security value. Also, given that we induce opinion dispersion, we can control for it. The first part of our third hypothesis is presented below:

Hypothesis 3a: In markets with short sale constraints, higher levels of opinion dispersion lead to higher overpricing.

We expect that the effect of opinion dispersion will be more pronounced right after any incident causing information asymmetry and dispersion in investors' beliefs. However, as time passes, information is disseminated in the market through trading, which should lead to a narrower information gap between optimists and pessimists. In other words, the level of opinion dispersion will eventually decrease which should result in overpricing induced by opinion dispersion to resolve.

Hypothesis 3b: The effect of opinion dispersion on overpricing should dissipate over time as investors incorporate market information in their valuation process and revise their views over asset prices accordingly, narrowing the information gap between optimists and pessimists.

The academic literature on short selling constraints highlights their effect on the speed of adjustment in the price discovery process. Diamond and Verrecchia (1987) present a model in which a short sale ban prevents pessimists from trading slowing the price discovery process. However, since active traders are aware of the restrictions on pessimists, they update their estimates accordingly. This eventually leads to security prices being traded close to their fair values. Bris, Goetzmann, and Zhu (2007) find that the price discovery process indeed slows down in countries with more short selling restrictions, but price inefficiencies persist. Beber and Pagano (2013), who examine the effect of short sale restrictions worldwide during the period 2007-2009, also find that the price discovery process is slower for stocks with short selling restrictions. Our fourth hypothesis is summarized below:

Hypothesis 4: In a market with short sale constraints the price discovery process slows down and reaches the true efficient price as investors revise their expectations to account for the fact that short sale constraints prohibit pessimists from participating in the market.

The effect of short sale constraints on liquidity has also been debated. Diamond and Verrechia (1987) argue that the slower price discovery process is associated with uncertainty about the value for a prolonged timeframe, which is translated into increased bid-ask spreads. Boehmer, Jones, and Zhang (2009) find that liquidity worsened during the short sale ban on US financial stocks in September 2008. This finding is consistent with the experience of market participants as highlighted in the comments of Christopher Cox, SEC Chairman at the time, who admitted that the ban may have been a mistake, as it led to a liquidity squeeze[2] without succeeding in preventing the prices from tumbling (Paley and Hilzenrath, 2008). Moreover, Sadka and Scherbina (2009) provide evidence of the link of overpricing and low liquidity, which is more likely to appear when opinion dispersion is high and short sale constraints bind. Our last hypothesis is summarized below.

Hypothesis 5: Short sale constraints lead to lower market liquidity, especially when opinion dispersion is high.

3. Experimental Study

This section provides a description of our experimental design, the unique features of our market, and the controls we employ to ensure the validity of our experimental findings. First, we cover some general terminology. A *cohort* is defined as a group of eight traders who always trade together. A *security* is a claim on a terminal liquidating dividend, denominated in laboratory dollars, which are converted into cash (US$) at the end of the trading session. A *trading session* refers to the 75-minute period in which a cohort participates in twenty consecutive trading trials. A *trading trial* is a period during which traders receive a trading target or information on the value of a specific security and trade accordingly. During each trial, traders can only trade one specific security.

[2] SEC approved Rule 201 which reintroduces a revised uptick rule along with other short sale restrictions: Securities Exchange Commission, "SEC Approves Short Selling Restrictions", Press Release, 2010-26.

3.1. Overview

We aim to study the effect of short sale constraints on market efficiency under various levels of opinion dispersion. We use a fully factorial repeated measures (balanced) design with the following factors:[3] short sale constraints (constrained, unconstrained), opinion dispersion (low, high), time (five intervals of 24 seconds each), trader type (informed, liquidity) and cohort (eight cohorts of 8 traders each, for a total of 64 participants). Trader type and cohort are manipulated across traders (i.e. between-subject factors) and all other factors are manipulated within traders (i.e. within-subject factors). Each cohort is repeatedly observed under all treatments.

More specifically, each cohort consists of four informed and four liquidity traders, with trader roles assigned randomly at the beginning of each trial. The security value, drawn from a uniform distribution, is determined before the trading trial starts. Informed traders receive some information about the liquidating "true" security value at the beginning of each trading trial, while liquidity traders receive a target. Opinion dispersion, which is induced by the distribution of information to informed traders at the beginning of every trading trial has two levels: "High" and "Low" and is described in detail in a dedicated section of this paper.

Cohort participants trade twenty securities sequentially. All cohorts trade the same securities in the same order. However, half of the cohorts face short sale restrictions during the first ten trading trials, whereas the remaining cohorts face short sale constraints during last ten trials. The short sale constraints we impose represent a complete ban on short selling. In every trading session there are ten trials reflecting markets with high opinion dispersion, five of which are also markets with short sale constraints. The other ten trials reflect markets with low opinion dispersion, five of which are markets with short sale constraints.

Table 1 describes how our two primary factors (short sale constraints and opinion dispersion) are manipulated in the 2x2 within-subjects balanced factorial design. Panel A in Table 1 presents the order of treatments assigned to each cohort and Panel B presents the details of each treatment order. As shown in Table 1, all cohorts trade the securities in the same order and experience all four settings.

We also introduce time as a third within-subjects factor aiming to test the dynamic effects of short sale constraints on market efficiency. With this goal in mind, we split the trading day in five time intervals (24 seconds each). We focus on the differences between the opening (first) interval, the midday (third) interval and the closing (fifth) interval. Therefore, when examining the effects of short sale constraints in a dynamic setting, we add a third factor (time) with three levels (opening, midday, and closing), and manipulate our factors (short sale constraints, opinion dispersion, and time) in a 2x2x3 within-subjects balanced factorial design.

Insert Table 1

3.2. Laboratory Market

Our laboratory market, which is comparable to the market developed in Bloomfield et al. (2005), resembles a double auction electronic market with a book-building period and a main trading

[3] A repeated measures design has the advantage being economical because each member is measured under all treatments or conditions. This advantage is particularly important when the number of treatments is large.

period. The market features an electronic limit order book, with pre- and post-trade transparency, which allows continuous trading and automated order execution following price-time priority rules. The trading functionality provides traders with the flexibility to submit limit and market orders. The maximum number of shares per order allowed is one. However, traders can enter multiple orders at the same time. Also, traders are able to cancel individual orders or all bids (or asks) at once. Traders can see their own standing orders as well as the orders submitted by other traders.

Each trading trial includes three periods: (1) information period, (2) pre-trading period, and (3) main trading period.

Information period: During this period participants receive information regarding their trader role for the particular trial. Informed traders receive information about the value of the security and liquidity traders receive a trading target. Then, informed traders are asked to provide an estimate of the security value based on the information received.

Pre-trading period: The pre-trading period lasts 20 seconds during which traders can enter and cancel as many orders as they wish. During this period no trades are executed. Any marketable orders do not result in a trade. These crossing orders are simply kept on the order book and at the end of the pre-trading period the order book is purged of these orders in the following way: if the highest bid crosses with the lowest ask, the more recent of the two orders is deleted from the book. This process is repeated until the highest bid price is below the lowest ask price.

Main trading period: The main trading period lasts 120 seconds during which traders are permitted to enter and cancel as many limit and market orders as they wish. During this period, orders can be executed and may result in trades. At this stage, traders can trade continuously and are free to pursue dynamic order placement and cancellation strategies.

3.3. Trader types

The market contains two types of traders: (1) four informed traders who possess information about the fundamental value of the security, and (2) four liquidity traders[4] who face liquidity needs (i.e. they trade on the basis of exogenous non-informational reasons). All traders whose roles are assigned randomly in each trial have unlimited cash in all trials. They also have unlimited short selling capacity when short selling is allowed.

Informed traders possess superior information regarding the value of the security. Information provided to informed traders is designed to induce opinion dispersion: two informed traders receive positive information about the value of the asset, while the other two receive negative information about the value of the asset. A detailed description of the opinion dispersion generation process is presented in the section that follows. All informed traders receive an endowment of five shares each, which was determined based on practice trials outside of the experiment. This endowment level is low enough to allow short sale constraints to be binding while at the same time it is high enough to warrant trading engagement from pessimist traders.

In each trial, informed traders earn a profit/loss equal to the sum of their realized and unrealized profit/loss in laboratory dollars. The realized profit/loss is the based on the trade prices

[4] These traders are also known as uninformed traders. In fact, Harris (1998) refers to this type of trader as an uninformed liquidity trader.

from any roundtrip trade and can be tracked during the trial. The unrealized profit/loss, which is not displayed during the trial, is calculated at the end of each trial by valuing any outstanding long or short position at the "true" security value. The total profit for the trial is disclosed at the end of each trial. Therefore, informed traders earn a profit every time they buy (sell) shares at a price below (above) the true security value, which provides them the incentive to actively participate in the price discovery process.

Liquidity traders do not possess information regarding the value of the security. Instead, they have a trading target which is communicated to them during the information period. The target may have a different direction for different liquidity traders. In every trial, two liquidity traders are not given any endowment and are asked to reach a trading target of 20 shares (liquidity net buyers), whereas the other two liquidity traders are given an endowment of 20 shares and a trading target of zero shares (liquidity net sellers). The goal of liquidity traders is to meet their target at some point during the trial at the most favorable prices possible. However, they are not obliged to end the trial with an inventory level equal to their trading target.

Similar to informed traders, liquidity traders earn a profit/loss equal to the sum of their realized profit/loss, which can be tracked during the trial, and their unrealized profit/loss, which is not displayed during the trial and is calculated based on the "true" value. Liquidity traders are more likely to lose money given that they have to complete directional trades and are at an informational disadvantage. Liquidity traders take a loss every time they buy (sell) shares at a price above (below) the security value. Additionally, liquidity traders who fail to reach their target during the trial are penalized 1000 laboratory dollars for every unfulfilled share at the end of the trading trial. For example, a liquidity trader with a target of 20 shares who never reaches a position of 20 shares and ends the trading trial with an inventory of 18 shares will receive a penalty of $2000. The penalty is large enough so that liquidity traders are better off hitting the exact target even if the prices at which they transact are unfavorable. The use of targets captures the notion that liquidity traders are trading for exogenous reasons related to liquidity needs (i.e. they are uninformed). Similar to informed traders, liquidity traders can see their total profit for the trial once it is completed.

In summary, the two types of traders have different levels of information, leading to substantially different motivations for trading. Informed traders trade with the aim of maximizing their profit based on their informational advantage. Liquidity traders are at a natural disadvantage and therefore trade with the aim of minimizing their loss while meeting their exogenous liquidity needs. However, once liquidity traders reach their targets they behave more like informed traders with the difference that their estimate of the true value depends solely on information gained through trading.

3.4. Opinion Dispersion

Inducing opinion dispersion is critical for testing many of our hypotheses. We achieve this with (1) our choice of security values, (2) the information distributed to traders at the beginning of each trading trial, and (3) our request to informed traders to provide a security value estimate after they receive information. Our goal is to influence informed traders in such a way that there are two optimists and two pessimists in the market, and to induce a gap between the beliefs of the two groups which alternates between two states: low and high.

Choice of security values: Security values are drawn from a uniform distribution within the range of [0, 1000]. The uniform distribution has the attractive feature that any value is equally likely. Therefore, participants assigned the role of informed traders do not have a reason to converge to the middle of the proposed range, a concern we would have if security values were drawn from a bell shaped distribution.

Information: We induce two levels of opinion dispersion by varying the information we distribute to informed traders. There are four randomly chosen informed traders in every trial. At the beginning of each trial, they receive a combination of hard (numerical) and soft (descriptive) information regarding the value of the security.

The hard information refers to a numerical range containing the "true value" of the security. In more detail, similar to Bloomfield et al. (2014), in the high dispersion setting two informed traders (optimists if "x" is positive and pessimists if "x" is negative) observe the true value of the security plus a predetermined random variable "x", which is drawn from the interval (-100,100) and is different from every security, while the other two informed traders (pessimists if "x" is positive and optimists if "x" is negative) observe the true value of each security minus the same predetermined variable "x" for that security. This way, while each trader observes imperfect information, there is perfect information released to the market, which should warrant market efficiency (the price should be equal to the true value of each security) when short sale constraints do not bind. In the low dispersion setting, two informed traders (optimists) observe the true value of the security plus a predetermined random variable "y", which is drawn from the interval (-20, 20), while the other two informed traders (pessimists) observe the true value of each security minus the same predetermined variable "y". Variables "x" and "y" are not independent from each other; they are both linked to a predetermined variable "z", which is drawn from a uniform distribution within the range (-1, 1). Then, "x" ("y") is equal to "z" multiplied by 100 (20), resulting in a predetermined random variable between [-100,100] ([-20, 20]).

For example, if true value = 363 and z = 0.55, then x = 55 (= 0.55*100) and y = 11 (= 0.55*20). This would mean that in the corresponding trial with high opinion dispersion, two informed traders (pessimists) would receive a true value signal equal to 308 (= 363 – 55) and would infer that the true value is between [208, 408], while two informed traders (optimists) would receive a true value signal equal to 418 (= 363 + 55) and would infer that the true value is between [318, 518]. In the low dispersion setting for the same security, two informed traders (pessimists) would receive a signal for the true value equal to 352 (= 363 - 11) and would infer that the value is in the range of [332, 372] while two informed traders would receive a true signal equal to 374 and would infer that the value is between [354, 374].

During the information period of each trial, we provide informed traders with the inferred ranges. In the directions emailed to participants prior to the experiment (Appendix A), we inform them how this range is computed. However, we do not disclose the link between "x" and "y". We also do not explicitly alert them about the opinion dispersion in each trial, although this could be easily inferred from the width of the provided information range.

Additionally, we utilize framing (Tversky & Kahneman, 1981) to further increase the difference between the two opinion dispersion levels. We do this by providing soft (descriptive) information about the earnings performance of the company. In the trials with high opinion dispersion, we provide optimists with positive information about the company's earnings (e.g. the company's earnings were 10% higher than last quarter), while at the same time we provide pessimists with negative information about the company's earnings (e.g. the company's earnings were 14% lower than the earnings of the closest competitor's), aiming to push the valuation estimates of the two groups apart. On the contrary, in the case of low opinion dispersion we provide optimists with negative information and pessimists with positive information about the earnings with the objective to bring their opinions closer. The information we provide to the two groups relies on the principle of framing and is not contradictory, in order to comply with the generally acceptable experimental principles.

Requesting a security value estimate: After receiving the relevant information, and before they start trading, informed traders are asked to provide us with their estimate of the "true value." Our objective for this request is twofold: (1) we force our subjects to think about the information provided to them and anchor them to one price before the trading starts, (2) we collect information on their views, which allows us to gauge whether we were successful in generating two distinctively different levels of opinion dispersion.

3.5. Experimental Design issues: Controls

A primary benefit of an experimental study is the ability to control for features of the experimental design that might influence behavior but are not the focus of the study (Bloomfied & O'Hara, 1999). We control for extremity, order, and carry-over effects (O'Rourke, Hatcher, & Stepanski, 2005), as well as differences in personal characteristics of individual subjects and cohorts.

Extremity

One concern, which may arise from our design, is the issue of extremity. First, for securities with extreme values (either very small or very large), the resulting information ranges for opinion dispersion surpass the natural zero bound or the upper bound set by the experiment (Lab$ 1000), thus, forcing those information ranges to be narrower. Second, according to Bloomfield (1996) securities with extreme values tend to be priced less efficiently. Third, given that our security values are declared to be between (0, 1000), we expect extreme low values to have much more room for overpricing than extreme high values, which are bounded by the upper limit of Lab$ 1000. Therefore, the extremity effect may be so large that securities with extremely low values could exhibit higher mispricing compared to high value securities, irrespective of the opinion dispersion in the particular trial.

To control for such differences across securities, our participants trade the security values included in the experimental design of Panel A in Table 1, which are uniformly distributed between (0, 1000). However, we exclude from our analysis securities with extreme values, or in more detail the securities with the upper and lower four fundamental values. This results in discarding 22% of our sample on each side. Table 2 presents the securities values of each cohort (in the order traded) with the values included in our analysis in bold.

Moreover, our security values are uniformly distributed between (0, 1000) instead of between (0, 100), which is often encountered in many experimental studies (Bloomfield et. al. 2009, 2014). This design could further limit the extremity effect. The rationale is that the perception of extremity in security values is stronger if the range of feasible security values is smaller, i.e. within (0-100) compared to (0-1000). Finally, an added benefit of a wider range of security values offers a bigger pool of potential values making it harder to draw inferences about those values from trial to trial.

Insert Table 2

Carry-over and order effects

Carry-over effects (also known as learning or practice effects) can be significant in a repeated measures design. These effects occur when an effect from one treatment changes (carries over to) participants' responses in the following treatment condition (O'Rourke, Hatcher, & Stepanski, 2005).[5] In this experimental design, the traders' actions may change simply because they become more familiar with the trading features of the market.

In order to control for these effects, an in-depth training session is held prior to running the experimental market. Participants receive extensive training in the mechanics of the trading platform (i.e. trading functionality of the software) and of the market (i.e. trader types, nature of the limit order book, short sale constraints).[6] They also take part in a series of practice trials that are identical to the actual experimental trials.[7]

Furthermore, altering the order of the treatments may also help mitigate these carry-over effects. As mentioned in an earlier section, we alternate the order of trials with short sale constraints and opinion dispersion, while keeping the order of traded securities fixed.

Order effects result when the ordinal position of the treatments biases participant responses (O'Rourke, Hatcher, & Stepanski, 2005). In order to eliminate possible order effects, we vary the treatment order across cohorts. Four cohorts are presented first with the no short sale constraints setting followed by the short sale constraints setting, while the other four cohorts trade in the opposite order.[8]

Different levels of intelligence, motivation, and familiarity

Another concern relates to the different levels of intelligence, motivation, and familiarity with the experimental environment across different participants and cohorts (Bloomfied & O'Hara, 1999). Besides having all eight cohorts trade in all experimental settings, each trader in a given cohort is randomly assigned to a trader type (informed trader or liquidity trader) at the start of trading for each security. The random assignment of participants helps minimize the possibility that differences across trader types and market dynamics are driven by individual characteristics. In order to control for differences across cohorts and avoid learning effects, we vary all settings across cohorts (as shown in Table 1) yet we keep the order of securities fixed. Through this, we attempt to control for differences between the aptitude, experience, and motivation of our subjects.

3.6. Subjects, Training, and Incentives

[5] Learning effects may violate the independence assumption for the error terms in standard ANOVA. Fortunately, a special feature of repeated measures analysis is a series of corrections to the standard statistics tests if violations are detected (Freund, Wilson, & Mohr, 2010). Also, by using one observation from each cohort in the statistical analysis, the assumptions underlying the ANOVA tests are not violated by these carry-over effects (Bloomfield, O'Hara, & Saar, 2005).

[6] Also, all participants had been exposed to this trading simulation before they took part of this experiment. This further mitigates any learning effects that may arise during the trading sessions.

[7] The data from these practice trials is not included in the statistical analysis.

[8] The technique where conditions are presented to different participants in a different order is known as counterbalancing. This technique is commonly used in experimental designs to control for order effects (O'Rourke, Hatcher, & Stepanski, 2005).

Our experiments were conducted in the Hanlon Financial Systems Lab at Stevens Institute of Technology. The trading simulation software was provided by Financial Trading System (FTS) via its FTS Interactive Markets trading simulator. This trading simulator was largely adapted to match this paper's experimental market design.

Our participants, primarily a mixture of undergraduate and graduate students, received extensive training before participating at the actual experiment. They were emailed detailed written instructions and were told to carefully review these instructions prior to the trading session. A copy of these instructions is provided in Appendix A. Also, prior to their participation in the trading simulation, all participants attended a 90-minute training session, which consisted of three parts: (1) a 30-minute discussion of the written instructions including an overview of the experimental market and the trading software (e.g. trading functionality, mechanics of the limit order book, trading rules such as price-time priority, the role of the market maker, etc.), (2) a 15-minute trading simulation where participants learn the basics of the trading software by trading and openly discuss any challenges they may have in using the software, and (3) a 45-minute practice session where participants trade in trials replicating the exact dynamics of the experiment (e.g. random assignment of trader types, trials with a pre-trading and a main trading phase, etc.). Participants were notified when short sale constraints were binding. They were not notified explicitly of the level of opinion dispersion in every trading trial. However, this could be inferred by the width of the provided information ranges. Some participants completed the training on a different day from the day they traded (typically a few days before). In this case they participated in the 30-minute practice session, which served as a "refresher", prior to participating in the actual experiment.

Notwithstanding the importance of adequate training, financial experiments must offer participants monetary incentives. The fundamental method of experimental economics is to create a setting that implements some institutional features of interest and then provide participants with incentives to maximize utility within that setting (Bloomfield & Anderson, 2010). In order to create these incentives, we adopt the reward structure in Bloomfield et al. (2005). Participants receive a final profit/loss in laboratory dollars, which is equal to the sum of all their profits/losses across all trials during the trading session. Actual U.S. dollars winnings for each session are calculated by subtracting a floor from each trader's winnings in laboratory dollars and then multiplying by a U.S. dollar conversion rate. Neither the floor value nor the conversion rate is known to the participants.[9] These two parameters are set equal to values that yield average cash winnings of US$ 15.00 per participant per trading session with a minimum payment is US$ 5.00.[10]

4. Statistical Methodology

In this section we provide a broad overview of the statistical methodology used to analyze the experimental data. The exact nature of the methodology depends on the hypothesis being tested and, thus, may be different for each hypothesis. Generally, the statistical methodology used to test

[9] These two values remain unknown to participants during the entire session to mitigate any potential risk-seeking behavior among participants who have low of negative balances (i.e. house money effects) and any other gaming behavior.
[10] A detailed description of the reward structure for this experiment is given to participants in the written instructions (see Appendix A).

this paper's hypotheses is repeated-measures analysis of variance (ANOVA) with between-subject factors. As with any ANOVA, repeated-measures ANOVA tests the equality of means. However, repeated-measures ANOVA is used when all members of a sample are measured under a number of different treatments. As the sample is exposed to each treatment in turn, the measurement of the dependent variable is repeated. A repeated measures analysis effectively treats each cohort as providing a single independent observation of the dependent variable. This design reduces a common problem in experimental economics of overstating statistical significance by assuming that repetitions of the same actions by the same group of subjects are independent events (Bloomfield, O'Hara, & Saar, 2009). Repeated measures designs are classified by the number of between-subject and within-subject factors. In order to understand the statistical analysis, it is necessary to first specify the applicable between- and within-subject combination. This paper's experimental design has a total of two between-subject factors (cohort and trader type) and three within-subject factors (short sale constraints, opinion dispersion, and time). However, statistical tests for each of the hypotheses may require the use of only a subset of factors.[11] Thus, to test for statistical significance, we compute the average of the dependent variable for each treatment (or cell) as defined by the appropriate subset of factors relevant to a given hypothesis.

A factorial ANOVA provides a methodology to test a variety of effects: (1) a significant *main effect* means that there is a difference between at least two levels of a factor with respect to mean scores on the response (or dependent) variable, (2) an interaction is a condition in which the effect of one factor on the response variable is different at different levels of another factor, and (3) a significant *simple effect* means that there is a significant relationship between a factor and the response variable at *a given level* of another factor (O'Rourke, Hatcher, & Stepanski, 2005).

To illustrate these three effects, lets take a look at hypothesis 3a: in markets with short sale constraints, a higher level of opinion dispersion across traders leads to higher overpricing. In this case the response variable is our measure of mispricing. A significant short sale constraint main effect (without a constraint/dispersion interaction effect) means that short sale constraints exert a similar influence on market efficiency across all levels of opinion dispersion. A significant constraint/dispersion interaction effect means that the impact of short sale constraints on market efficiency is different in markets with different levels of opinion dispersion. Finally, a significant constraint/dispersion interaction effect suggests the presence of significant simple effects. The opinion dispersion factor has two levels: low and high. Therefore, we could examine two separate simple effects: (1) a simple effect of short sale constraints for the low-dispersion market and (2) a simple effect of short sale constraints for the high-dispersion market. In other words, testing for simple effects of short sale constraints is similar to testing for main effects, but it is done for one opinion dispersion level at a time.

The relevant type of effect (i.e. interaction, main, or simple effect) depends on the nature of the hypothesis but it also dependent on the significance of the interaction effect. If the interaction effect is significant, there is no need to test for main effects, as these results would be

[11] In experimental sciences, it is not unusual to find a design described as "one between-subject and two within-subject repeated measures" (Freund, Wilson, & Mohr, 2010). The type of design is, thus, hypothesis-specific. For example, to test the hypothesis 3A, only two factors may be needed: short sale constraints and opinion dispersion. Hypothesis 3B, however, requires an additional between-subject factor: time.

misleading.[12] Instead, it is appropriate to test for simple effects. If the interaction effect is not significant, main effects should be tested instead (simple effects also remain meaningful). For the sake of consistency, and more importantly appropriateness, the statistical tests will always begin with a test of the interaction effect and then they will be sequenced as follows:

- If the interaction effect is not significant, we will proceed to test for main effects for each of the factors being considered. If any of the main effects are significant, we will provide contrast tests and/or multiple comparison tests for any significant main effects.[13]
- If the interaction effect is significant, we will proceed to interpret a profile plot to gain some insight on the nature of the interaction.[14] Then, we will test for any relevant simple effects. Similar to main effects, if a given simple effect is significant, we will provide contrast tests and or multiple comparison tests.

5. Results

The focus of our analysis is on the effect of short sale constraints on the quality of markets under different degrees of opinion dispersion. More specifically, we analyze market efficiency, price discovery, and liquidity. We begin with preliminary summary statistics to provide a sense of how typical is the aggregate behavior resulting from these experiments. We then examine the market quality effects mentioned above.

Before we analyze the experimental data, we proceed to clean the order book data from any instances when the bid-ask spread is either negative or zero. The bid-ask spread could be negative (zero) if a single trader enters an ask order priced lower (equal to) than her own highest bid order. These two orders do not cross since they are entered by the same trader and, therefore, remain standing on the book. The bid-ask spread could also be negative if the ask order book is empty while the bid order book is not.

The total number of observations collected for our experimental analysis is 53,425. Out of this total, the proportion of actions that resulted in a trade was 21.5% or 11,508 total trades. Our cleaning rule eliminated 347 observations (0.65% of the total number of observations) and 99 trades (0.89% of the total number of trades).

5.1. Summary Statistics

Is the market well behaved?

Figure 1 presents the evolution over time of three market-wide variables: trading volume, bid–ask spread, and pricing error. These statistics were computed only for markets with no short sale constraints to evaluate how well our market behaves. Each panel divides trading into five 24-

[12] If there is a significant constraint/dispersion interaction, a significant short sale constraints main effect becomes misleading as it suggests that short sale constraints exerts a significant influence on market efficiency across both levels of opinion dispersion contradicting the interaction effects. For this reason, a significant interaction effect precludes the existence of a meaningful main effect.

[13] Contrasts tests examine statistical significance of the difference between a given factor level (e.g., high opinion dispersion) and the selected benchmark factor level (e.g., low opinion dispersion). Multiple comparison tests determine which pairs of factor levels are statistically different. Of course if the factor for which the main effect is significant has only two levels, the contrast test and the multiple comparison test would be equivalent. In this case, there would only be a need to use one of these tests.

[14] A profile plot graphs the means of the response variable for each of the factor/level combinations or treatments. For example, hypothesis 3A consists of two factors (i.e., short sale constraints and opinion dispersion) with two levels each. Thus, the profile plot would consist of four points, each point representing the mean of the response variable (i.e., mispricing measure) for a given constraint/dispersion combination.

second intervals. While approximately 150 shares were traded in a typical market, Panel A shows that volume exhibits the usual ''U'' shape observed in equity markets. The steep increase in trading volume during the last trading interval reflects the trading behavior of liquidity traders rushing to hit their trading targets. Panel B shows the time series behavior of the market's bid-ask spread. Each data point is computed as the average bid-ask spread for the 24-second time interval. The spread declines from an average of $200 in the first interval to an average of near $100 in interval four. During the last interval the spread increases to approximately $140. This increase in the spread near the closing can be explained by the liquidity-taking behavior of liquidity traders.[15] Panel C shows the average pricing error, calculated as the absolute value of the deviations of the mid-quote from the true security value. The pricing error decreases by an average of roughly 30% from the opening time to its lowest point near closing time. These patterns suggest that markets behave reasonably well, in light of theoretical, archival, and experimental studies. Of particular importance is that our experimental markets appear to gradually incorporate information, a feature consistent with market efficiency.

Insert Figure 1 – Panels A, B, & C

As a robustness check, we also compute the average percentage mispricing in markets where no short sale constraints were imposed. Table 3 provides a summary of the results. The average percentage mispricing across all unconstrained markets is equal to 1.49%. This average is not statistically different from zero (p = 0.7804) implying that our unconstrained experimental markets are efficient. In order to isolate any potential effect of opinion dispersion on market efficiency, we compute the average percentage mispricing across unconstrained markets with different degrees of induced opinion dispersion. The mispricing measure is not statistically different from zero under neither the low opinion dispersion nor the high opinion dispersion markets (the p-values are 0.7006 and 0.4595, respectively).

Insert Table 3

The ability of markets to incorporate information into security prices depends on informed traders exploiting their informational advantage, which should result in positive profits for the informed group as a whole. In a zero-sum game, uninformed or liquidity traders should then lose money. Our experimental markets indeed produce the expected results. Figure 2 shows that informed traders generated an average of $847 while liquidity traders lose on average the same amount. These results imply that informed traders not only have an informational advantage but they successfully trade on it. This provides further evidence that our markets are well behaved.

Insert Figure 2

Opinion dispersion

A particular important aspect of our study is the inducement of different levels of opinion dispersion across informed traders. In order to evaluate the success of our efforts, we construct two ex-post opinion dispersion measures: the absolute forecast error and the absolute percentage

[15] This pattern is consistent with observed behavior in equity markets and with the theoretical model of Brock and Kleidon (1992).

forecast error. The absolute forecast error is calculated as the dollar difference between the informed traders' estimate of value (provided during the information session of each trial) and the predetermined true value. The absolute percentage forecast error represents the same measure in percentage terms. Table 4 provides the summary statistics of these two measures across the two opinions dispersion levels: low and high. We expect that averages of both measures are significantly lower when ex-ante opinion dispersion is low as opposed to when induced opinion dispersion is high.

Insert Table 4

Panel A shows the average absolute forecast error for each opinion dispersion level. In markets with low dispersion levels, traders misestimate the true value by an average of $25 while in high dispersion markets traders exhibit a much larger forecast error ($92). The difference in means is statistically different from zero with a p-value lower than 0.01. Panel B shows similar results when we use the absolute percentage forecast error as the measure of opinion dispersion. The high-dispersion average error is 21.74% compared to an average error of only 5.78% for the low dispersion case. These two means are statistically different ($p < 0.01$).

A quick look at the other summary statistics further supports the finding that opinion dispersion was induced properly across markets. The standard deviation of the forecast errors is roughly three times larger for the high dispersion markets relative to the low dispersion markets. This result is consistent across the two different measures of opinion dispersion. Finally, the highest available absolute forecast error in low dispersion markets is roughly half the size of the smallest absolute forecast error provided by traders in high dispersion markets. Overall, these results provide strong support for the ability of our experimental design to induce the desired level of opinion dispersion across informed traders.

5.2. Hypotheses Testing

Overpricing

Short Sale Constraints and Overpricing

Our first hypothesis predicts that, when short sale constraints are binding, the market price will be significantly higher than the true value of the security. However, in the absence of short sale constraints the market should be efficient, which we have already shown to be true. In order to test this hypothesis, we look at the main effect of short sale constraints on mispricing. We compute two measures of mispricing: (1) dollar mispricing computed as the difference between the market price and the true security value and (2) percentage mispricing which is obtained by dividing the dollar mispricing by the true security value. Table 5 summarizes our main effect findings.

Insert Table 5

As we can see, under unconstrained markets, both mispricing variables are, on average, very close to zero and are not statistically different from zero. On the other hand, the mispricing measures for constrained markets are both positive and significantly higher than zero. The average percentage mispricing in markets with short sales constraints is 33.55%. The average dollar mispricing in these markets is roughly $110. More importantly, the difference in average mispricing across markets with and without short sale constraints is statistically significant with p-values lower than 0.01. This test provides strong support for our hypothesis. Short sale constraints have a significant main effect on market efficiency (or mispricing) with markets where short sale constraints are binding exhibiting an economically meaningful overpricing.

Short Sale Constraints and Bubbles

Under rational expectations, the average trade price should not be higher than the view of the most optimistic trader. Harrison and Kreps (1978) predict that in markets with short sales constraints the overpricing will often be such that the price could surpass the view of the most optimistic investor (i.e. a price bubble). In other words, in constrained markets, the price could be higher than the highest valuation among all market participants, as investors take long speculative positions aiming to later sell their shares at a profit to optimists, which dominate short sale constrained markets. We test this prediction by first comparing the market's average trade price to the view of the most optimistic trader. In our experimental market, informed traders are asked to estimate the true value of the security prior to the market opening and after they receive some information about the value of the asset. Here, we define a price bubble as a market where the average trade price is higher than the highest value estimate.[16]

In order to test our second hypothesis, we measure price bubbles by computing the speculative premium as the average difference between the trade price and the highest value estimate, which is the view of the most optimistic investor. We examine the speculative dollar and percentage premium across trials with and without short sale constraints. Panel A in Table 6 shows that across 48 constrained trials the average speculative premium is $45. This average is statistically different from zero ($p = 0.0655$). Furthermore, the speculative premium is, on average, $129 higher in constrained markets relative to unconstrained markets, since, as expected, the speculative premium is negative in unconstrained markets. This difference is statistically significant with a p-value of 0.0048. In order to control for the effect of the price level on the speculative premium, we also calculate the speculative percentage premium as a proxy of a price bubble using the highest value estimate as the base. The average speculative percentage premium in constrained markets is 11% (this average is statistically different from zero with a p-value of 0.0175). This average is 28% higher than the average premium under unconstrained markets. Also, note that both proxies in unconstrained markets are significantly lower than zero. In other words, there is no price bubble in when short sales are allowed.

Insert Table 6

Our experimental markets provide informed traders with an information range as well as earnings information. Informed traders were then asked to provide a single estimate of value. This estimate may not be necessarily equal to either bound of the originally information range. For example, the most optimistic trader may have estimated a value lower than the upper bound of her information

[16] The average trade price in our experimental markets is equivalent to the volume-weighted average price (VWAP) since we impose a restriction on the order quantity of 1 share.

range. As a robustness check, we also calculate the difference between the market's average trade price and the upper bound of the most optimistic information range. The rationale is that investors could be adapting their valuation during the trial to account for information they gain through trading. However, we would still expect them to value the asset below the upper bound of the original information range, provided they do not engage in any speculative behavior. Panel B in Table 6 presents the results. Similar to Panel A, we find that the speculative premium using the highest upper bound of the most optimistic information range is significantly different from zero. For example, the percentage difference between the average trade price and the upper bound of the most optimistic range is, on average, 9%. This average is statistically higher than zero ($p = 0.0243$).[17]

Our experimental study provides strong evidence supporting the notion that short selling constraints can lead to price bubbles. In fact, in approximately 52% of our trials, the average trade price is higher than the valuation estimate of the most optimistic trader, which rarely happens in unconstrained markets. We then conclude that under short sale constraints markets behave irrationally, leading to price bubbles.

<u>Short Sale Constraints, Opinion Dispersion, and Overpricing</u>

We now consider how opinion dispersion influences the ability of the market to incorporate information into prices efficiently. To conduct this analysis, we define mispricing as the percentage difference between the trade price and the true security value (using the true value as the base). Figure 3 plots the interaction effect of short sale constraints and opinion dispersion on mispricing.

Insert Figure 3

Previously, we have shown that the there is a main effect of short sale constraints on mispricing. Here we examine whether the effect of short sale constraints is different across various levels of opinion dispersion: markets with low opinion dispersion across informed traders and markets where the dispersion of opinion across traders is high. Miller (1977) predicts that in markets with short sale constraints, a higher level of opinion dispersion leads to higher overpricing. Our experimental markets incorporate both trials with low opinion dispersion and trials with high dispersion.[18] To test this prediction, we compare the average mispricing of trials across these two dispersion levels and we interact this effect with the effect of short sale constraints.

The interaction effect shown in Figure 3 is not significant ($p = 0.1849$). This is because of the dominant main effect of short sale constraints on mispricing. In order to find support for our third hypothesis and Miller (1977) prediction is only necessary to find a significant simple effect of opinion dispersion on mispricing under short sale constraints.

In markets with low opinion dispersion, the average mispricing is significantly higher (p =

[17] In order to check that our results are not mainly driven by a highly right skewed distribution of the overpricing measure (which would inflate the average), we counted the number of constrained markets where the average trade price was higher than our bubble benchmark. We find that in 52% of our constrained markets, the average trade price was higher than the valuation estimate of the most optimistic trader. Also, in 50% of the constrained markets, the average trade price was higher than the upper bound of the most optimistic range.

[18] Note that all of our experimental markets induce some level of opinion dispersion. Therefore, we expect to find overpricing across all constrained markets. However, markets with high opinion dispersion should exhibit a higher overpricing than markets with low opinion dispersion. This is implies that the simple effect of opinion dispersion on mispricing under constraints markets is our most important test.

0.0166) when short sale constraints are binding (24.29%) than when no short sale constraints are imposed (-2.29%). In markets with high opinion dispersion across traders, the average mispricing is also significantly higher (p = 0.0006) when short sale constraints are binding (42.34%) than when no short sale constraints are imposed (4.42%). More importantly, however, is the difference in average mispricing under short sale constraints in markets with low opinion dispersion relative to markets with high opinion dispersion. Our results show that under short sale constraints, markets with high dispersion exhibit a significantly higher overpricing (42.34%) than markets with low dispersion (24.29%). The p-value of the difference in means is 0.0304. The difference is economically significant. The size of the overpricing in constrained markets with high dispersion is roughly twice the size of the overpricing in low dispersion constrained markets. These results are consistent with the theoretical predictions of Miller (1977).[19]

Short Sale Constraints, Opinion Dispersion, Overpricing, and Time

Miller (1977) provides a static prediction of the effects of short sale constraints on mispricing across different levels of opinion dispersion. In this section, we provide a dynamic analysis of Miller's prediction by incorporating a new factor: time. Specifically, we examine how the mispricing effect of opinion dispersion evolves over time. We focus our analysis on markets where short sale constraints are binding given that when short sale constraints are absent there is no overpricing.

We expect that the effect of opinion dispersion will be most pronounced at the opening of each trial, right after opinion dispersion is induced. As trading progresses, informed traders update their valuation. Therefore, we expect the information gap between optimists and pessimists to narrow regardless of the level of the originally induced opinion dispersion. Towards the closing, both the level of opinion dispersion and its effect on overpricing decreases.

We study the interaction effect of short sales constraints and opinion dispersion on mispricing at three different times of the day: opening, midday, and closing.[20] Figure 4 plots the time evolution of mispricing across four relevant market conditions: (1) no short sale constraints with low dispersion, (2) no short sale constraints with high dispersion, (3) short sale constraints with low dispersion, and (4) short sale constraints with high dispersion. Our analysis focuses mainly on the last two market conditions.[21]

Insert Figure 4

First, we test the interaction effect of time and condition (where condition refers to all four market conditions listed above) on mispricing. The interaction effect is significant (p = 0.0097). This implies that the average mispricing varies across time for one or more of the four market

[19] For robustness, we examine the data using the average dollar mispricing as the dependent or response variable and results remain unchanged. We also quantified mispricing as the deviation of the true value from the bid-ask midquote. In this case, we did not find a significant interaction effect. This result can be explained by the asymmetric overpricing of the ask orders relative to the bid orders in constrained environments. In other words, under short sale constraints, the average ask price was significantly higher than the average bid price.

[20] Our experimental trials consist of 120 seconds of continuous trading (excluding the 20-second pre-trading session). In order to examine the effect of time on mispricing, we divide the total time into five time intervals of 24 seconds each and we define the opening period as the first interval, the midday period as the third interval, and the closing period as the fifth interval.

[21] In markets with no short sale constraints there is no significant over/under pricing across all three levels of the time factor (opening, midday and closing) and across both levels of opinion dispersion (low and high). Theoretically, this implies that the plotted lines for both unconstrained cases are flat and horizontal at zero.

conditions described. To better understand this interaction effect (and given the lack of mispricing in unconstrained markets), we now proceed to analyze the behavior of mispricing across time in markets with short sale constraints.

At the opening, there is a significant difference in the average overpricing of markets with different levels of opinion dispersion. As predicted by Miller (1977), markets with high opinion dispersion have a higher average overpricing (33.53%) than markets with low opinion dispersion (8.33%). This difference is statistically significant ($p = 0.0177$). However, if we examine the evolution of overpricing across the trading day, we can see that the difference in overpricing decreases. By midday, the high dispersion markets exhibit an overpricing of 50.20% while low dispersion markets exhibit an average overpricing of 34.74%. Although still different, the significance of the difference in means during the midday period is weaker ($p = 0.0751$). Finally, at the market closing, the statistical significance of the difference in average overpricing disappears ($p = 0.3719$).

These results suggest that the effect of opinion dispersion on mispricing disappears through the trading day. Although mispricing is always present in constrained markets, the difference in this mispricing due to opinion dispersion is only strongly significant at the opening, but it disappears toward the end of the trading day. This result can be intuitively explained within a model of adaptive valuations. As they trade, informed traders incorporate market information (e.g. market prices, volume, and order flow) and revise their valuations. These revisions will bring their valuations closer to each other reducing the degree of opinion dispersion. By the end of the trading day, the market should exhibit low opinion dispersion and a corresponding level of overpricing.

Our findings are consistent with the predictions of Miller (1977) but they extend the analysis by incorporating the time factor. Overall, the time behavior of overpricing shows that trading and valuation revisions dissipate the level of opinion dispersion and, therefore, its effect on overpricing. To test whether Miller (1977) prediction holds at the opening, we test the interaction effect using data from the opening period only. Our expectation is that results will be much more pronounced at the opening, right after opinion dispersion is induced, than the overall sample.

Insert Figure 5

Figure 5 shows the interaction effect of short sale constraints and opinion dispersion on mispricing during the opening period. The interaction effect is significant ($p = 0.0402$). A quick look at the profile plot shows that the interaction effect is caused by the significant difference in average mispricing across levels of opinion dispersion only for constrained markets. We test this simple effect and the difference is significant ($p = 0.0177$). We also perform the same test of simple effects for the midday and closing periods separately. We do not find significant interaction effects in either case. This evidence provides additional support of an interaction between opinion dispersion and time.

Price Discovery

In this section, we analyze our fourth hypothesis. In a market with short sale constraints the price discovery process slows down, but still reaches the true efficient price as investors revise their expectations to account for the fact that short sale constraints prohibit pessimists from participating in the market. This hypothesis is consistent with the prediction of Diamond and Verecchia (1987). In order to test this hypothesis, we examine how short sale constraints influence the ability of the market to incorporate information into prices over the trading day. Figure 6 presents the average mispricing for constrained and unconstrained markets over time.

Insert Figure 6

As we have shown above, markets with no short sales constrained are efficient. Here we show the time evolution of mispricing. At the opening, unconstrained markets have an average overpricing of 7%. This overpricing decreases to 2% by the middle of the trading day and it continues to decrease toward the closing period. None of these averages are statistically different from zero. These results provide supporting evidence for the market efficiency of our unconstrained markets.

When short sale constraints are in place, the time evolution of mispricing is in stark contrast with unconstrained markets. Under short sale constraints, our markets exhibit significant overpricing across the entire trading day. At the opening, the average overpricing is 21% (statistically different from zero with a p-value equal to 0.0309). The overpricing increases over the day to an average of 43% by the middle of the day. At the closing, although the average overpricing drops to 30%, the significance remains strong at 1% ($p = 0.0005$). To further examine the time evolution of overpricing under short sale constraints, we contrast mispricing in constrained markets at different time periods. Comparing the closing period relative to the opening period, we find that there is no statistical difference in the average mispricing ($p = 0.2288$). On the other hand, when we contrast the average mispricing during the middle of the day to the opening period, we find a significant difference ($p = 0.0071$).

The humped shape of the mispricing measure over time provides strong partial support for the prediction of Diamond and Verecchia (1987) and for our fourth hypothesis. Another interesting finding is the overpricing persistence toward the end of the trading day (e.g. closing period). Although the overpricing is significantly lower than the midday average, it remains near the level of overpricing exhibited during the opening period. This result suggests that short sale constraints slow down the price discovery process and prices do not fully adjust to true value. In our experimental markets, the average closing price is approximately 30% higher than the true security value.

Liquidity

Our final hypothesis analyzes the effect of short sale constraints and opinion dispersion on the degree of liquidity. As we have shown above, short sale constraints reduce the speed of price discovery. This delayed resolution of uncertainty about the true value may tend to reduce liquidity (Beber & Pagano, 2013). Sherbina and Sadka (2009) present evidence that mispricing is often accompanied by low liquidity, especially when opinion dispersion is high. In this section, we test this notion.

To test our hypothesis, we measure liquidity using the average percentage spread and the average trading volume. The experimental nature of our study allows us to calculate the percentage spread relative to the true value of the security. This measure of liquidity is superior to most empirical spread measures. Percentage spreads, using a market-based price (e.g. midquote or ask price) as a scaling factor to the dollar spread, incorporates too much noise in the form of transaction costs.

Panel A in Figure 7 shows the interaction effect of short sales constraints and opinion dispersion on our spread measure of liquidity (this effect is significant with a p-value equal to 0.0273). As expected the percentage spread is higher in markets with short sale constraints. This difference, however, is only significant for markets with high opinion dispersion across informed traders (p = 0.0788). In markets with low opinion dispersion, the difference in liquidity across short sale constraint levels is statistically insignificant (p = 0.2425).

Also note that in unconstrained markets, the average degree of liquidity is the same regardless of the level of opinion dispersion.[22] On the other hand, opinion dispersion seems to have an effect on liquidity within trials with short sale constraints. The average percentage spread in constrained trials with high opinion dispersion is 57%. This average spread is significantly higher relative to trials with low opinion dispersion (p = 0.0398).

Insert Figure 7 – Panel A & Panel B

Panel B in Figure 7 performs the same statistical analysis provided in Panel A but in this case we use trading volume as our proxy for liquidity. Our results are generally consistent with the results obtained with the percentage spread. The key difference is that trading volume exhibits a significant decrease across short sale constraint levels regardless of the degree of opinion dispersion. In other words, there is a strong and significant main effect (p = 0.0131) of short sale constraints on trading volume that dominates the interaction effect (p = 0.1298). Similar to our spread analysis, the simple effect of opinion dispersion on trading volume under short sale constraints is significant (p = 0.0252), suggesting that the effect of short sale constraints on liquidity is greater when the degree of opinion dispersion in the market is high. Overall, these results provide supporting evidence for our hypothesis. We find a significant interaction effect where short sale constraints influence market liquidity only within markets that exhibit high opinion dispersion across participants. These findings are consistent with most available evidence showing that short sale constraints damage liquidity.

6. Conclusion

In this study, we investigate the effect of short sale constraints on the main characteristics of market quality. In particular, we examine how short sale constraints influence market efficiency, price discovery, and market liquidity. We find that markets with short sale constraints exhibit significant overpricing. In fact, we find evidence of price bubbles where the market price surpasses the view of the most optimistic investor in the market.

Miller (1977) predicts that the degree of opinion dispersion across investors causes overpricing in markets with short sale constraints. We test this prediction and find that under short sale constraints, markets with higher opinion dispersion exhibit higher overpricing. We then extend this analysis to examine the behavior of overpricing under short sale constraints and different levels of opinion dispersion across time. At the opening, markets with high opinion

[22] Markets with low opinion dispersion exhibit an average percentage spread of 39.68% while markets with high opinion dispersion exhibit an average percentage spread of 35.02%. These two averages, however, are not statistically different from each other (p = 0.3235).

dispersion exhibit high overpricing, but as trading takes place, market participants revise their valuation and opinion dispersion dissipates. This results in the size of the overpricing converging across opinion dispersion levels toward the closing of the day.

Short sale constraints do not only create overpricing, but this overpricing may be persistent over the trading day. We find that the price discovery process slows down when short sale constraints bind, with prices never converging to their true value. In fact, we find that the average closing price is roughly 30% higher than the true security value.

This slower and partial adjustment of prices is also coupled with lower market liquidity, measured by the percentage spread and trading volume. We find that liquidity decreases substantially when short sale constraints are in place. Furthermore, this decline in liquidity is more significant in markets with a high level of opinion dispersion.

An important contribution of our research is to make clearer how short sale constraints impact market quality. To this end, we provide robust experimental evidence of the adverse effects of short sale restrictions on market efficiency, liquidity, and price discovery, while accounting for varying levels of opinion dispersion.

References

Asquith, P., P.A. Pathak, and J.R. Ritter, 2005. Short interest, institutional ownership, and stock returns, *Journal of Financial Economics* 78(2), 243–276.

Beber, A., & Pagano, M. (2013). Short selling bans around the world: Evidence from the 2007-2009 crisis. *Journal of Finance 68*, 343-381

Bloomfied, R., & O'Hara, M. (1999). Market transparency: Who wins and who loses? *The Review of Financial Studies 12* (1), 5-35.

Bloomfield, R. (1996). Quotes, prices, and estimates in a laboratory market. *Journal of Finance 51* (5), 1791-1808.

Bloomfield, R., & Anderson, A. (2010). Experimental finance. In H. K. Baker, & J. R. Nofsinger (Eds.), *Behavioral Finance* (pp. 113-130). Hoboken, NJ: John Wiley & Sons, Inc.

Bloomfield, R., O'Hara, M., & Saar, G. (2005). The "make or take decision" in an electronic market: Evidence on the evolution of liquidity. *Journal of Financial Economics 75* (1), 165-199.

Bloomfield, R., O'Hara, M., & Saar, G. (2009). How noise trading affects markets: An experimental analysis. *The Review of Financial Studies 22* (6), 2275-2302.

Bloomfield, R., O'Hara, M., & Saar, G. (2014). Hidden Liquidity: Some new light on dark trading. *Journal of Finance, forthcoming*.

Boehmer, R.D., B.R. Danielsen, and S. M. Sorescu (2006). Short sale constraints, differences of opinion and overvaluation. *Journal of Financial and Quantitative Analysis* 41(2), 455–487.

Boehmer, E., Jones, C., & Zhang, X. (2014). Shackling short sellers: The 2008 shorting ban. *Review of Financial Studies 26*, 1363-1400.

Bris, A., Goetzmann, W., & Zhu, N. (2007). Efficiency and the bear: Short sales and markets around the world. *Journal of Finance 62*, 1029-1079.

Brock, W. A., & Kleidon, A. (1992). Perdiodic market closure and trading volume: A model of intraday bids and asks. *Journal of Eonomic Dynamics and Control 16*, 451-489.

Chen, J., H. Hong, and J.C. Stein (2002). Breadth of ownership and stock returns, *Journal of Financial Economics* 66, 171–205.

Desai, H., K. Ramesh, S.R. Thiagarajan, and B.V. Balachandran (2002). An investigation of the informational role of short interest in the Nasdaq market, *Journal of Finance* 57, 2263–2287.

Diamond, D.W., Verrechia, R. (1987). Constraints on Short-Selling and Asset Price Adjustment to Private Information. *Journal of Financial Economics* 18, 277-311.

Diether, K.B., C.J. Malloy, and A. Scherbina (2002). Differences of opinion and the cross section of stock returns, *Journal of Finance* 57, 2113–2141.

Miller, E.M. (1997). Risk, uncertainty, and divergence of opinion, *Journal of Finance* 52, 1151–1168.

Fellner, G. and E. Theissen (2014). Short sale constraints, divergence of opinion and asset prices: Evidence form the laboratory, *Journal of Economic Behavior and Organization* 101, 113-127.

Figlewski, S. (1981). The informational effects of restrictions on short sales: Some empirical evidence, *Journal of Financial and Quantitative Analysis* 16, 463–476.

Freund, R. J., Wilson, W. J., & Mohr, D. L. (2010). *Statistical Methods* (3rd ed.). Academic Press.

Goetzmann, N.W. and M. Massa (2005). Dispersion of opinion and stock returns, *Journal of Financial Markets* 8, 325–350.

Harris, L. (1998). Optimal dynamic order submission strategies in some stylized trading problems. *Financial Markets, Institutions & Instruments 7* (2), 27-74.

Harrison, M.J., Kreps, D.M. (1978). Speculative investor behavior in a stock market with heterogeneous expectations. *Quarterly Journal of Economics* 92, 323-336.

Haruvy, E. and C.N. Noussair (2006). The Effect of Short Selling on Bubbles and Crashes in Experimental Spot Asset Markets. *Journal of Finance* 61, 1119-1157.

Hong, H. and J.C. Stein (2003). Difference of opinion, short-sales constraints, and the market crashes, *Review of Financial Studies* 16, 487–525.

Nagel, S. (2005). Short sales, institutional investors and the cross-section of stock returns. *Journal of Financial Economics* 78, 277–309.

O'Rourke, N., Hatcher, L., & Stepanski, E. J. (2005). *A Step-By-Step Approach to Using SAS for Univariate and Multivariate Statistics* (2nd ed.). Cary, NC: SAS Institute Inc.

Paley, A.R. and D.S. Hilzenrath (2008). SEC chief defends his restraint: Cox rebuffs criticism of leadership during crisis, *The Washington Post*, 24 December. p. A01.

Sadka, R. and A. Scherbina (2009). Mispricing and Costly Arbitrage. *Journal of Investment Management* 7, 11-13.

Securities Exchange Commission. SEC approves short selling restrictions. Retrieved on 2010-02-26 from http://www.sec.gov/news/press/2010/2010-26.htm.

Tversky, A. and D. Kahneman (1981). The framing of decisions and the psychology of choice. *Science*, New series 211:4481, 453-458.

Tables and Figures

Table 1
Our experimental design includes eight cohorts and each cohort consists of four informed and four liquidity traders. It is a balanced 2x2 within-subject factorial design, with factors for short sale constraints and opinion dispersion. Panel A shows that cohorts vary depending on whether they trade in a market with a short sale ban in the first ten trials (SSC - Order 2) or the last ten trials (SSC - Order 1). They also vary depending on the order trials they face with low and high of opinion (OD – Block 1 and OD – Block 2). As shown in Panel B, the short sale constraints factor has two levels: ban and no ban and the opinion dispersion factor has two levels: high and low. Also, all cohorts trade 20 securities in the same order.

Panel A

Cohorts	Short Sale Constraints Order	Opinion Dispersion Block
1	SSC - Order 1	OD - Block 1
2	SSC - Order 1	OD - Block 2
3	SSC - Order 2	OD - Block 1
4	SSC - Order 2	OD - Block 2
5	SSC - Order 1	OD - Block 1
6	SSC - Order 1	OD - Block 2
7	SSC - Order 2	OD - Block 1
8	SSC - Order 2	OD - Block 2

Panel B

	Short sale constraints		Opinion Dispersion	
Security No	SSC - Order 1	SSC - Order 2	OD- Block 1	OD - Block 2
1	No ban	Ban	Low	High
2	No ban	Ban	Low	High
6	No ban	Ban	Low	High
7	No ban	Ban	Low	High
8	No ban	Ban	Low	High
3	No ban	Ban	High	Low
4	No ban	Ban	High	Low
5	No ban	Ban	High	Low
9	No ban	Ban	High	Low
10	No ban	Ban	High	Low
12	Ban	No Ban	Low	High
15	Ban	No Ban	Low	High
16	Ban	No Ban	Low	High
19	Ban	No Ban	Low	High
20	Ban	No Ban	Low	High
11	Ban	No Ban	High	Low
13	Ban	No Ban	High	Low
14	Ban	No Ban	High	Low
17	Ban	No Ban	High	Low
18	Ban	No Ban	High	Low

Table 2
This table presents the security values, which are drawn from the uniform distribution in the range of [0, 1000] and are common to all cohorts. Moreover, all cohorts trade the same securities in the same order which is presented in the first column of the table. Security values that appear in bold are included in the statistical analysis. Data associated with security values that do not appear in bold are discarded to avoid any extremity influence in our results. Effectively, we are discarding 22% of each side of our uniform distribution.

Security No.	Security value
1	92.0
2	**363.0**
3	**426.0**
4	**578.0**
5	**354.0**
6	177.0
7	**298.0**
8	**520.0**
9	955.0
10	806.0
11	**510.0**
12	**749.0**
13	783.0
14	**342.0**
15	**682.0**
16	**241.0**
17	215.0
18	915.0
19	68.0
20	**370.0**

Table 3
The percentage mispricing is computed as the difference between the market price and the true security value divided by the true security value. In this table we present the overall average percentage mispricing, which is obtained by averaging the mispricing variable within trials and across trials and cohorts. We also show the average percentage mispricing for only trials with a low degree of opinion dispersion. Finally, we present the average percentage mispricing across markets with a high degree of opinion dispersion. For each of these mispricing series, we perform a t-test to test for the significance of the average percentage mispricing. This table presents both the t-statistic and the p-value for each of the three tests.

	Average % Mispricing	t-statistic	p-value
Overall	1.49%	0.29	0.7804
Low Opinion Dispersion	-2.29%	-0.40	0.7006
High Opinion Dispersion	4.42%	0.78	0.4595

Table 4
This table presents the summary statistics for two measures of opinion dispersion: the absolute forecast error (Panel A) and the absolute percentage forecast error (Panel B). The absolute forecast error is calculated as the difference between the value estimates provided by informed traders and the true security value. The absolute percentage forecast error is obtained by dividing the absolute forecast error by the true security value. The summary statistics are presented separately for each level of opinions dispersion: markets with low dispersion and markets with high dispersion. The difference in the average across opinion dispersion levels is presented and t-test of the difference in means is performed.

Panel A: Absolute Forecast Error

	Average	Std. Deviation	Minimum	Maximum
Low Dispersion	\$25	\$10	\$13	\$40
High Dispersion	\$92	\$27	\$69	\$151
Difference in means	\$67			
P-value	< 0.01			

Panel B: Absolute Percentage Forecast Error

	Average	Std. Deviation	Minimum	Maximum
Low Dispersion	5.78%	2.04%	3.06%	8.88%
High Dispersion	21.74%	6.29%	15.90%	35.00%
Difference in means	15.96%			
P-value	< 0.01			

Table 5
This table shows two measures of mispricing: dollar mispricing and percentage mispricing. The dollar mispricing is the difference between the market price and the true security value. The percentage mispricing is obtained by dividing the dollar mispricing by the true security value. The table presents the average dollar and percentage mispricing for both, constrained and unconstrained markets. The difference in the mispricing average across short sale constraint levels is presented and t-test of the difference in means is performed.

	% Mispricing	$ Mispricing
Unconstrained	1.49%	-14.42
Constrained	33.55%	109.73
Difference in means	32.06%	124.15
P-value	0.0026	0.0008

Table 6
Panel A shows the dollar and percentage difference between the average trade price in a trial and the value estimate provided by the most optimistic informed trader (i.e., the highest value estimate in the trial). The percentage difference is computed as the ratio between the dollar difference and the average trade price. There are 48 trials with no short sale constraints and 48 trials with short sale constraints. Panel A shows the average dollar and percentage difference for each of the two short sale constraints levels. The difference in averages across short sale constraint levels is also presented and t-test of the difference in means is performed. Panel B shows a similar analysis but here the focus is on the difference between the average trade price per trial and the upper bound of the optimistic information range.

Panel A: Average Trade Price Relative to Most Optimistic Forecast

	N	Percentage Difference	Dollar Difference
Unconstrained	48	-17%	-$84
		< 0.01	< 0.01
Constrained	48	11%	$45
		0.0175	0.0665
Difference in means		28%	$129
P-value		<.0001	0.0048

Panel B: Average Trade Price Relative to Upper Bound of Most Optimistic Range

	N	Percentage Difference	Dollar Difference
Unconstrained	48	-21%	-$103
		< 0.01	< 0.01
Constrained	48	9%	$40
		0.0243	0.0672
Difference in means		30%	$142
P-value		< 0.01	< 0.01

Figure 1
We divide each trading trial into five time intervals of 24 seconds each. For each trial we sum the number of shares traded in each interval to obtain the trading volume per interval. Panel A shows the trading volume for each interval averaged across trials and cohorts. The bid-ask spread is the difference between the highest-prices market bid and the lowest-priced market ask. We calculate the average bid-ask spread in each time interval. Panel B shows the bid-ask spread for each interval averaged across trials and cohorts. The pricing error is obtained as the absolute value of the difference between the price midquote and the true security value. For each interval, we average the pricing error. Panel C shows the pricing error for each interval averaged across trials and cohorts.

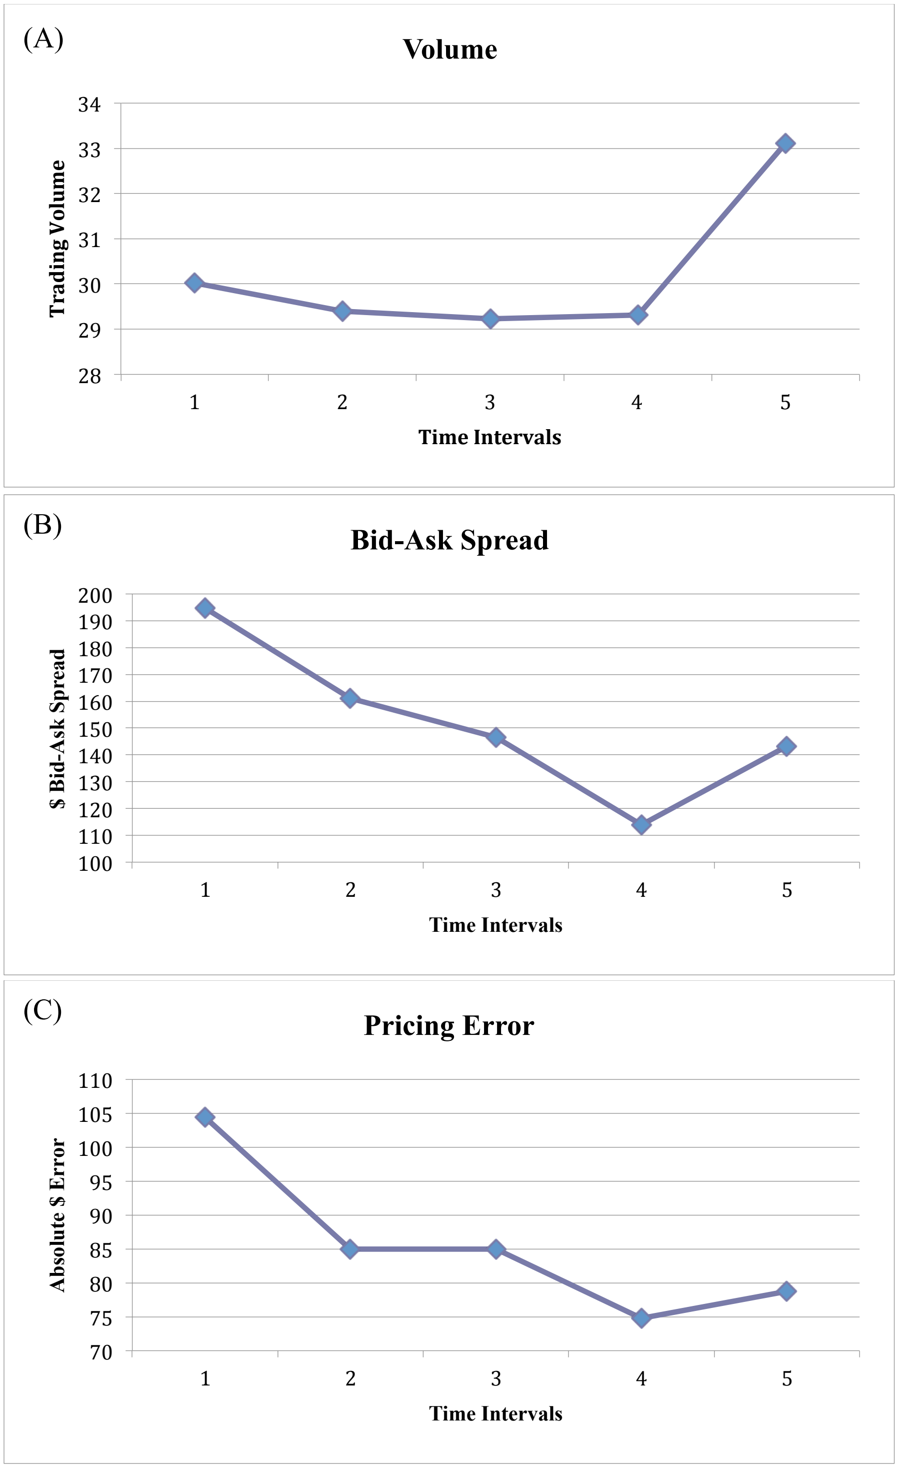

Figure 2
Profit is measured as the dollar gain/loss for each trader in a given trial. This figure shows the average profit across trials and cohorts for each of two different types of traders. On average, informed traders generated a profit of \$847.65 in a trial while liquidity traders took a loss of \$847.65 in a typical trial. Given the zero-sum nature of our experimental markets, the average dollar gain of informed traders is equal to the magnitude of the average loss for liquidity traders.

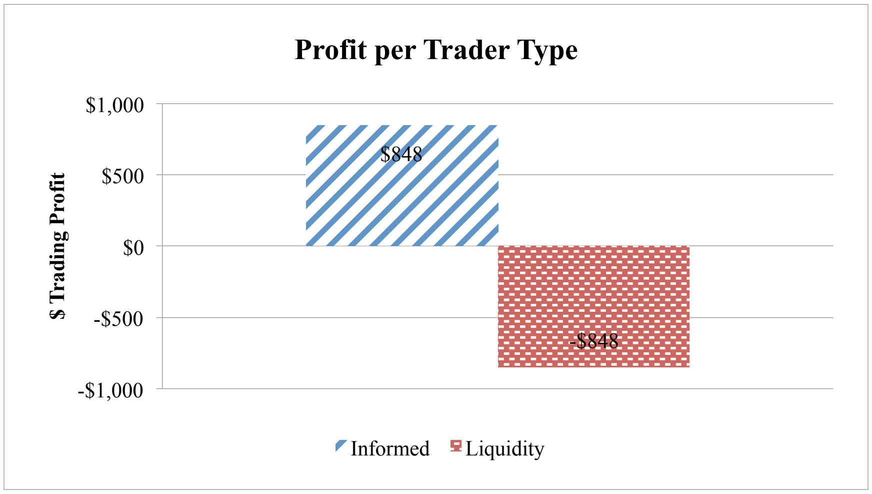

Figure 3

This figure plots the average percentage mispricing (measured as the dollar difference between the market price and the true security value divided by the true security value) for each of four market conditions: (1) unconstrained markets (no short sale constraints) with a low degree of opinion dispersion, (2) unconstrained markets with a high degree of opinion dispersion, (3) constrained markets (short sale constraints are present) with a low degree of opinion dispersion, and (4) constrained markets with a high degree of opinion dispersion. The plotted values represent the averages across trials and all cohorts.

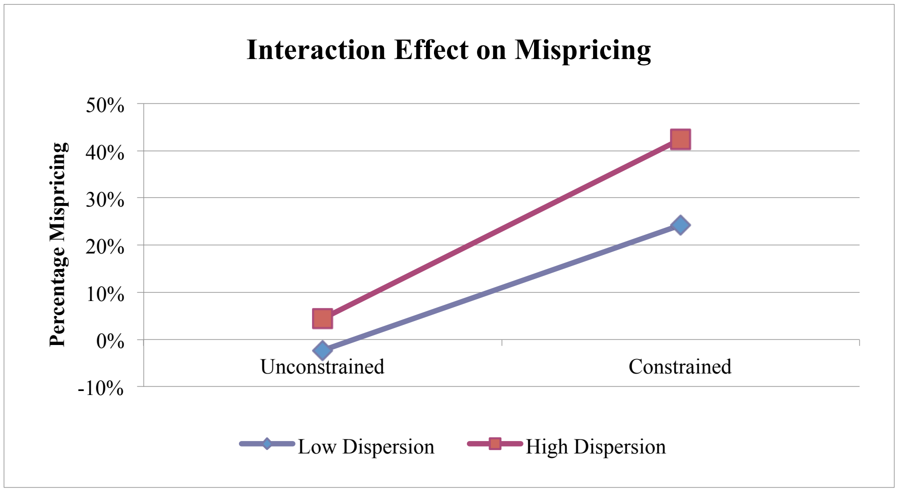

Figure 4
This figure plots the average percentage mispricing (measured as the dollar difference between the market price and the true security value divided by the true security value) for each of four market conditions: (1) unconstrained markets (no short sale constraints) with a low degree of opinion dispersion, (2) unconstrained markets with a high degree of opinion dispersion, (3) constrained markets (short sale constraints are present) with a low degree of opinion dispersion, and (4) constrained markets with a high degree of opinion dispersion. The variable is computed separately for each of three 24-second time interval: the first 24 seconds of trading (opening period), the third 24 seconds of trading (midday period) and the last 24 seconds of trading (closing period). The plotted values represent the averages across trials and all cohorts.

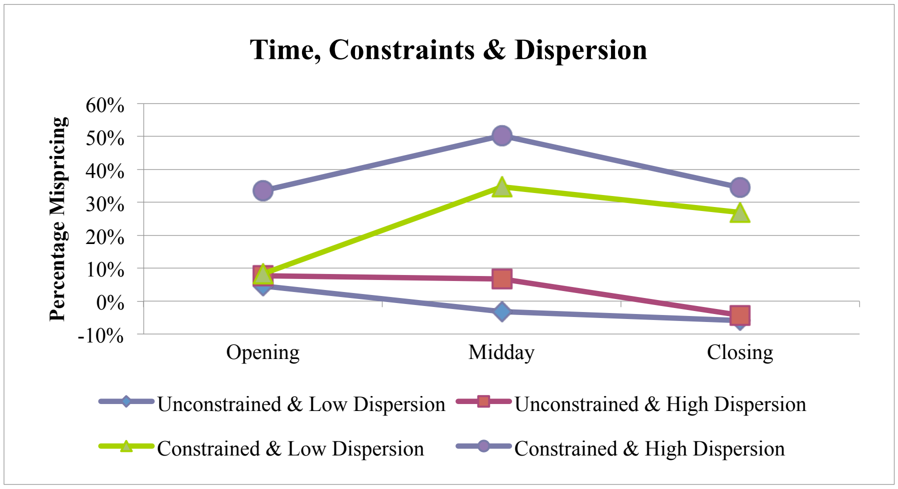

Figure 5
This figure plots the average percentage mispricing (measured as the dollar difference between the market price and the true security value divided by the true security value) for each of four market conditions: (1) unconstrained markets (no short sale constraints) with a low degree of opinion dispersion, (2) unconstrained markets with a high degree of opinion dispersion, (3) constrained markets (short sale constraints are present) with a low degree of opinion dispersion, and (4) constrained markets with a high degree of opinion dispersion. The variable is computed only for the first 24 seconds of trading or opening period and then averaged across trials and all cohorts.

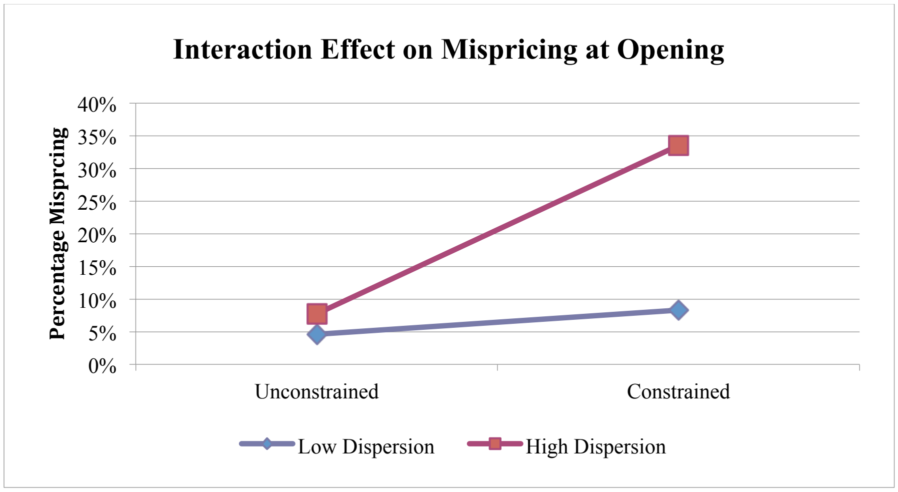

Figure 6
This figure plots the average percentage mispricing (measured as the dollar difference between the market price and the true security value divided by the true security value) for two distinct types of markets: markets with and without short sale constraints. The plotted values are computed separately for each of three 24-second time interval: the first 24 seconds of trading (opening period), the third 24 seconds of trading (midday period) and the last 24 seconds of trading (closing period). The plotted values represent the averages across trials and all cohorts.

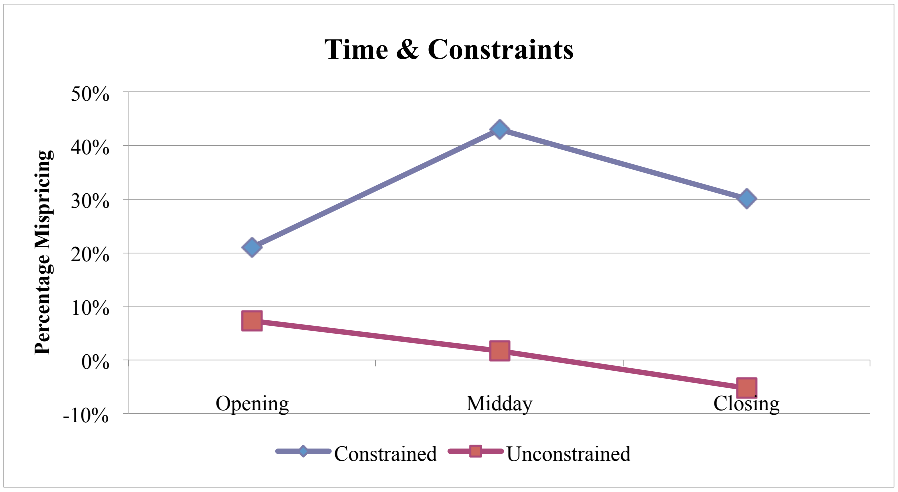

Figure 7
Panel A in this figure plots the average percentage bid-ask spread (measured as the difference between the best bid price and the best ask price divided by the true security value) for each of four market conditions: (1) unconstrained markets (no short sale constraints) with a low degree of opinion dispersion, (2) unconstrained markets with a high degree of opinion dispersion, (3) constrained markets (short sale constraints are present) with a low degree of opinion dispersion, and (4) constrained markets with a high degree of opinion dispersion. The plotted values represent the averages across trials and all cohorts. Panel B presents a similar graph but it here the focus is on the trading volume computed as the number of shares traded in a trial.

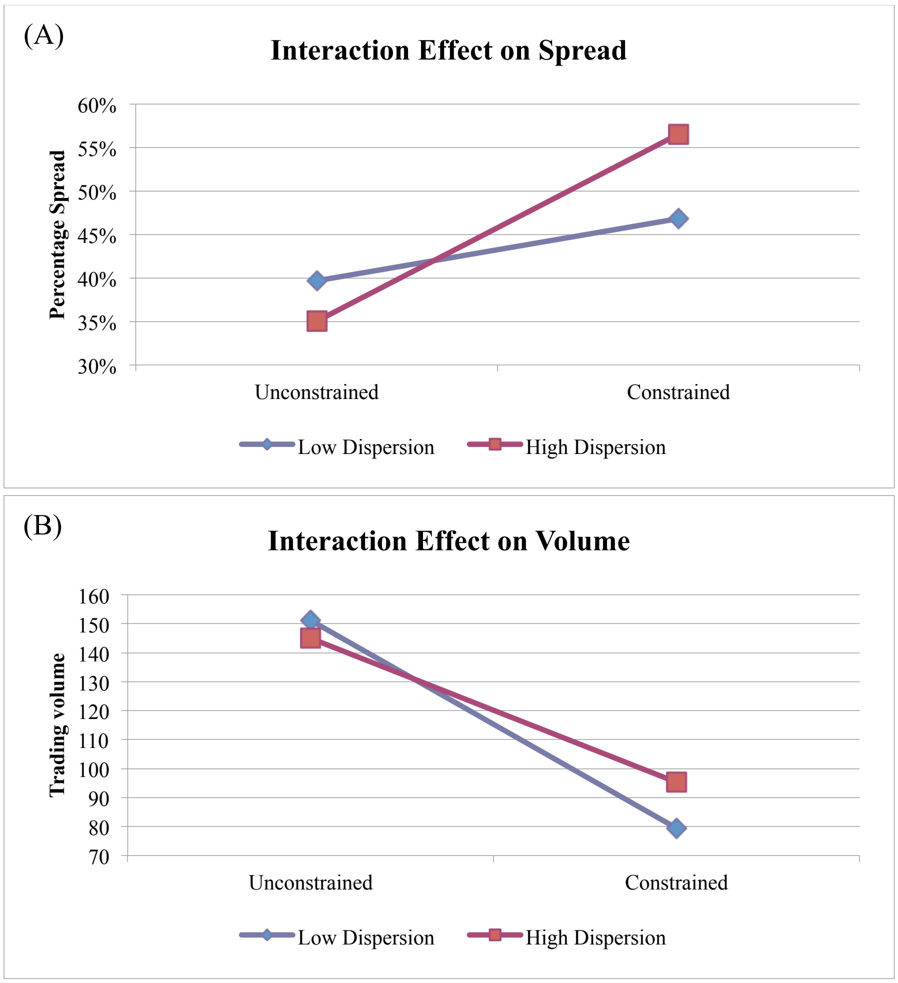

Appendix A

Experimental Trading Simulation
Instructions

In this trading session you will participate in a total of 20 trials. In each trial you will trade a different security that is valued in "laboratory dollars" (LAB$). At the end of the session, we will convert your trading gains into U.S. dollars (US$) to determine your payment. To trade these securities, you will be using the Financial Trading Services (FTS) trading software. Please refer to Appendix I and II for a detailed overview of the FTS trading screen and instructions on how to login.

Ten Basic Terms

1. *Bid* – it is an order to *buy* shares at a stated price (the bid price). The bid price is the highest price the buyer is willing to pay for one share.
2. *Ask* – it is an order to *sell* shares at a stated price (the ask price). The ask price is the lowest price the seller is willing to receive for one share.
3. *Bid book* – it is a list ordered by price (highest price first) of all the bids traders have entered.
4. *Ask book* – it is a list ordered by price (lowest price first) of all the asks traders have entered.
5. *Best bid* – it is the bid with the highest price on the bid book.
6. *Best ask* – it is the ask with the lowest price on the book.
7. *Entering a bid* (limit order to buy) – a trader willing to *buy* shares at a stated price can submit a bid to the book. The bid will be held on the bid book until another trader chooses to "take" it.
8. *Entering an ask* (limit order to sell) – a trader willing to *sell* shares at a stated price can submit an ask to the book. The ask will be held on the ask book until another trader chooses to "take" it.
9. *Taking a bid* (market order to sell) – a trader willing to *sell* shares can *take the highest bid* from the bid book in two ways: (i) she can directly sell to the highest bid or (ii) she can enter an ask with a price lower than the highest bid price.
10. *Taking an ask* (market order to buy) – a trader willing to *buy* shares can *take the lowest ask* from the ask book in two ways: (i) she can directly buy from the lowest ask or (ii) she can enter a bid with a price higher than the lowest ask price.

The Trading Session

A trading session consists of trading 20 securities successively (i.e., 20 trials). Short selling will be allowed in some of these trials and prohibited in others. You will be notified when short selling is not allowed. Information about each security and each trader is generated right before the beginning of each trial. Traders will trade based on this information until the trial ends.[23] Prior to the start of a new trial, new information is generated.

The Trading Trial

Each trial will be split into four parts: receive information, estimate value, pre-trading and main trading.

Receive information:

[23] A trial can be compared to one trading day in real markets. News about the value of the company is released prior to the market opening (i.e., overnight). Once the market opens, traders trade on this information. The difference in this simulated environment is that no additional information is released during the trial.

First, you will receive some information about the liquidating "true" value of the stock (informed traders) or your trading target (liquidity traders).

Estimate the value:
Then, we will ask you to provide us with your estimate of the liquidating value. Naturally, if you are an informed trader you will be able to estimate the liquidating value better.

Pre-trading (20 seconds)
During the pre-trading period, traders can enter orders (bids and asks), but no one can take them. This means that no actual trading will occur during this period. At the end of the pre-trading period, the highest bid and the lowest ask will be paired up. If they "cross" (the bid is priced higher than the ask) the more recent order will be deleted and a new pair will be matched. This will be repeated until there are no crossing orders remaining. The purpose of this pre-trading period is to allow traders to enter orders into the book before trading takes place. In other words, traders will "build" the book before trading can take place.

Main Trading (120 seconds)
During the main trading period, all traders can enter bids and asks, and can also take the bids and asks posted by the other traders (i.e., orders can cross). In other words, during this period trading takes place. Executions follow price/time priority, meaning that orders at most competitive prices will be executed first. Orders at the same price level will be executed in following time priority; orders submitted first are also executed first.

The Security Value
The range of permissible prices in this trading session is between 0 and 1000. The liquidating "true" value of the security is generated using a uniform probability distribution. In other words, all values within the 0-1000 range are equally likely to be selected. The security value will be determined and shown only to *some* traders (i.e., the "informed" traders) prior to the beginning of each trial.
Remember that the *security value* and the *market price* are not necessarily the same thing**.** The market price is determined by the amount traders are willing to pay or accept (i.e., transactions), and may change as trading progresses during each trial. The security value is determined prior to the start of each trial and does not change during the trial.

Types of Traders
There are two broad types of traders: (i) informed traders and (ii) liquidity traders. Each trader type will have different trading objectives and levels of information. Your trading screen will tell you what type of trader you are prior to the start of each trial. You will be randomly assigned a trader type before the beginning of each trial.

Informed traders
They *know* a range for the "value" of the security as well as information on the company's earnings, which they learn right before trading starts. They trade because they have valuable information regarding the security value. Therefore, they will earn a profit every time they buy (sell) a share at a price below (above) the security value.

The range for each security will be (-100,100) or (-20, 20). We draw a random number "x" that can be any integer between (-100,100) or (-20, 20) respectively. Before trading starts two informed traders will be notified of the liquidating "true" value plus x and two different informed traders will be notified of the liquidating "true" value minus x.

For example, if the liquidating "true" value is 300, x=5 and the range is (-20, 20):

Two informed traders will be informed of the following minimum and maximum liquidating "true" value:

Min liquidating "true" value: 300 – 5 – 20 = 275
Max liquidating "true" value: 300 – 5 + 20 = 315

Also, two informed traders will be informed of the following minimum and maximum liquidating "true" value:

Min liquidating "true" value: 300 + 5 – 20 = 285
Max liquidating "true" value: 300 + 5 + 20 = 325

You will be provided with one of the two ranges. You will also be provided with some soft information. The format of the information is described below.

Information received by informed traders on their screen would look like this:

"Your value estimate is within 275 to 315; Earnings are 19% higher than the competitor's earnings"

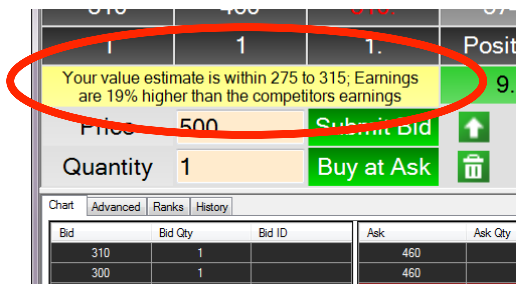

They will start each trading trial with zero cash and a small endowment of shares of the security. They can take unlimited long positions (unlimited cash borrowing is allowed). They can also take unlimited short positions whenever short selling is allowed (short selling will be allowed in some trials and prohibited in other trials).[24]

- They can *enter a bid* (limit order to buy) by entering a price and left clicking on "Submit Bid."
- They can *enter an ask* (limit order to sell) by entering a price and left-clicking on "Submit Ask."
- They can *take bids* (sell at highest bid) by left clicking on "Sell to Bid."
- They can *take asks* (buy at lowest ask) by left-clicking on "Buy at Ask."
- They can cancel individual or multiple orders.

Informed traders can see the market's best bid and ask as well as *all orders* on the book (refer to Appendix I for a detailed view of the trading screen)

[24] Short selling refers to the process of borrowing shares in order to sell them hoping that their value will go down. This strategy allows the trader to sell shares without owning them. You will be notified by the experiment administrator if short-selling is allowed before the beginning of a new trial.

Liquidity traders
They *do not know* the “value” of the security (i.e., they are “uninformed” traders). Therefore, they trade because of reasons other than information related to the security value. Liquidity traders are given a “target” number of shares they are required to trade before the end of the trial.
There are two types of liquidity traders:

- Net buyers – they have a target of +20 shares and receive neither a cash endowment nor a share endowment.
- Net sellers – they have a target of 0 shares and receive no cash endowment. They do, however, receive an endowment of 20 shares.

Typical information received by liquidity traders on their screen would look like this:

“Your target is 20 shares”

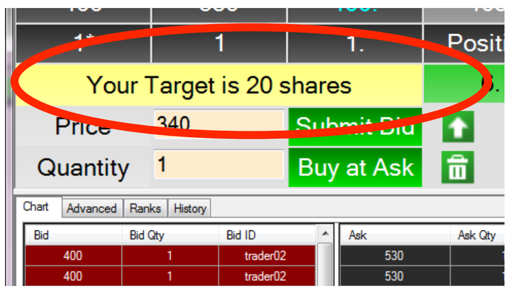

- They can *enter a bid* (limit order to buy) by entering a price and left clicking on “Submit Bid.”
- They can *enter an ask* (limit order to sell) by entering a price and left-clicking on “Submit Ask.”
- They can *take bids* (sell at highest bid) by left clicking on “Sell to Bid.”
- They can *take asks* (buy at lowest ask) by left-clicking on “Buy at Ask.”
- They can cancel individual or multiple orders.

Similar to informed traders, liquidity traders can see the market’s best bid and ask as well as *all orders* on the book (refer to Appendix I for a detailed view of the trading screen).

Liquidity traders incur a *penalty* of LAB$ 1000 per share for failing to achieve their targets (once they reach their targets they are allowed to trade freely without penalty). These penalties are large enough that liquidity traders are always better off trading enough to hit their target, even if they must buy at very high prices or sell at very low prices to do so.

Basic Trading Rules

- Once you enter an order on the book, you can only trade if someone takes the other side of your order. For example, if you have entered a bid to the book at LAB$ 200, it will remain on the book until another trader takes it (or until you cancel the order).
- You **cannot** trade with yourself. If you enter an ask at LAB$ 700 and then you enter a bid at LAB$ 750, these two orders will not cross and they will both remain standing on the book.
- The maximum number of shares (i.e., quantity) you can enter per order is 1. However, you can enter multiple orders at any price point. For example, if you enter 3 *separate* bids (one share each) at LAB$ 520, you will have 3 shares listed at the LAB$ 520 price point.

Getting Paid

You will start each trial with a zero cash balance and a zero share balance. At the end of each trial, the shares you own ("n") pay an amount ("V") equal to the security "true value."

- If you have a *positive* share balance ($n > 0$), then "n x V" LAB$ will be *added* to your cash balance.
- If you have a *negative* share balance ($n < 0$), then "n x V" LAB$ will be *subtracted* from your cash balance.

The resulting cash balance is your trading gain/loss in LAB$. Any penalties assessed for failing to hit your exact target are deducted from your resulting cash balance decreasing your gain or increasing your loss.

Remember that you do not "get paid" in laboratory dollars (LAB$). These LAB$ need to be converted into US$ before you can get paid. The conversion will be done using the following formula:

US$ Winnings = Baseline + (Your LAB$ - Type-Average LAB$) x Conversion Rate

You will not know the exact *baseline* or *conversion rate*. However, we will tell you three key facts:

1. The type-average LAB$ refers to the average LAB$ for only those traders that are of your same type (i.e., informed traders or liquidity traders).
2. The *baseline* is a positive US$ amount. If your LAB$ at the end of the trial is equal to the average LAB$ for traders of your type, you will earn an amount of LAB$ equal to the baseline.
3. The *conversion rate* is positive, meaning that the more LAB$ you win, or the fewer you lose, the more US$ you take home.

The parameters (baseline and conversion rate) are set so that the *average* US$ winnings will be between US$10 and US$12 per person per hour (not including the training session). Finally, US$ are determined separately for each trading session. Losses in one session do not offset gains in another session.

Other Rules

Please do not talk with other traders or look at their computer screens without explicit permission from the experiment administrator. Please ask the administrator before leaving the room for any reason.

Appendix I: The Trading Screen

The FTS Trading Screen is the first screen that appears once you have successfully logged into the market. It is important to be familiar with its elements before beginning to trade. However, if at any time during trading you cannot remember what a particular object on the screen does, simply place your mouse over the object and a mouse-over description will appear.

Below is a description of the different features you will find on the trading screen (refer to the figure below):

1. **Bid** – it displays the current best (highest) bid (listing the price/quantity). If there is an asterisk (*) next to the price/quantity it means that your bid is the current best bid.
2. **Ask** – it displays the current best (lowest) ask (listing the price/quantity). If there is an asterisk (*) next to the price/quantity it means that your ask is the current best ask.
3. **Security Name** – this is the name of the security that you will be trading on each trial. There could be multiple security names listed. By left-clicking on a security name, you will access the security's market screen.
4. **Last** – it displays the last traded price. This box changes color according to how the last traded price changed relative to the previous last traded price.
5. **Buy VWAP** – it specifies the trader's volume-weighted average buying price. It is simply the average price the trader has paid for all shares purchased. It updates in real time.
6. **Sell VWAP** – it specifies the trader's volume-weighted average selling price. It is simply the average price the trader has received for all shares sold. It updates in real time.
7. **Time Left** – each trial specifies a particular amount of time. Once the market begins, Time Left will begin to count down. The time will reset after the completion of each trial (a bigger screen with the time left can be found on the upper right hand corner of the screen).
8. **Information Window** – here you will see information regarding the security value and the target number of shares that you are required to trade. This window will also display the type of trader you are.
9. **Position** – it specifies your current position in a security (i.e., the number of shares that you currently hold in inventory). If this number is positive, then you have a long position. If the number is negative this means you have a short position.
10. **Realized P&L** – it specifies the profit (or loss) of a trader based only on the shares that have been both purchased and then sold (or vice versa). It ignores any shares in the trader's inventory. It updates in real time.
11. **Trading Controls –** this section contains the buttons/fields that will allow you to trade. The buttons may be different for different traders. Please refer to the "Types of Traders" section (above) for a detailed explanation of the trading controls. Remember that to enter a bid and an ask on the book, you must first enter a price and a quantity:
    a. *Price* – here you can enter a specified price for your order (i.e., bid price or ask price).
    b. *Quantity* – here you can enter a specified quantity for your order. In this case, however, the quantity can only be equal to 1.
12. **Order Book** – this book collects all orders entered by all traders in the market. The *bid book* (left half) displays all bids with the highest-priced bid listed first. The *ask book* (right half) displays all asks with the lowest-price ask listed first. The top-listed bid and ask are the *market BBO* (best bid/offer, where offer and ask are equivalent terms).
13. **Bid/Ask/Price Graph –** it provides a graph of the evolution of the best bid, best ask, and traded prices over the trading trial. It updates in real time.

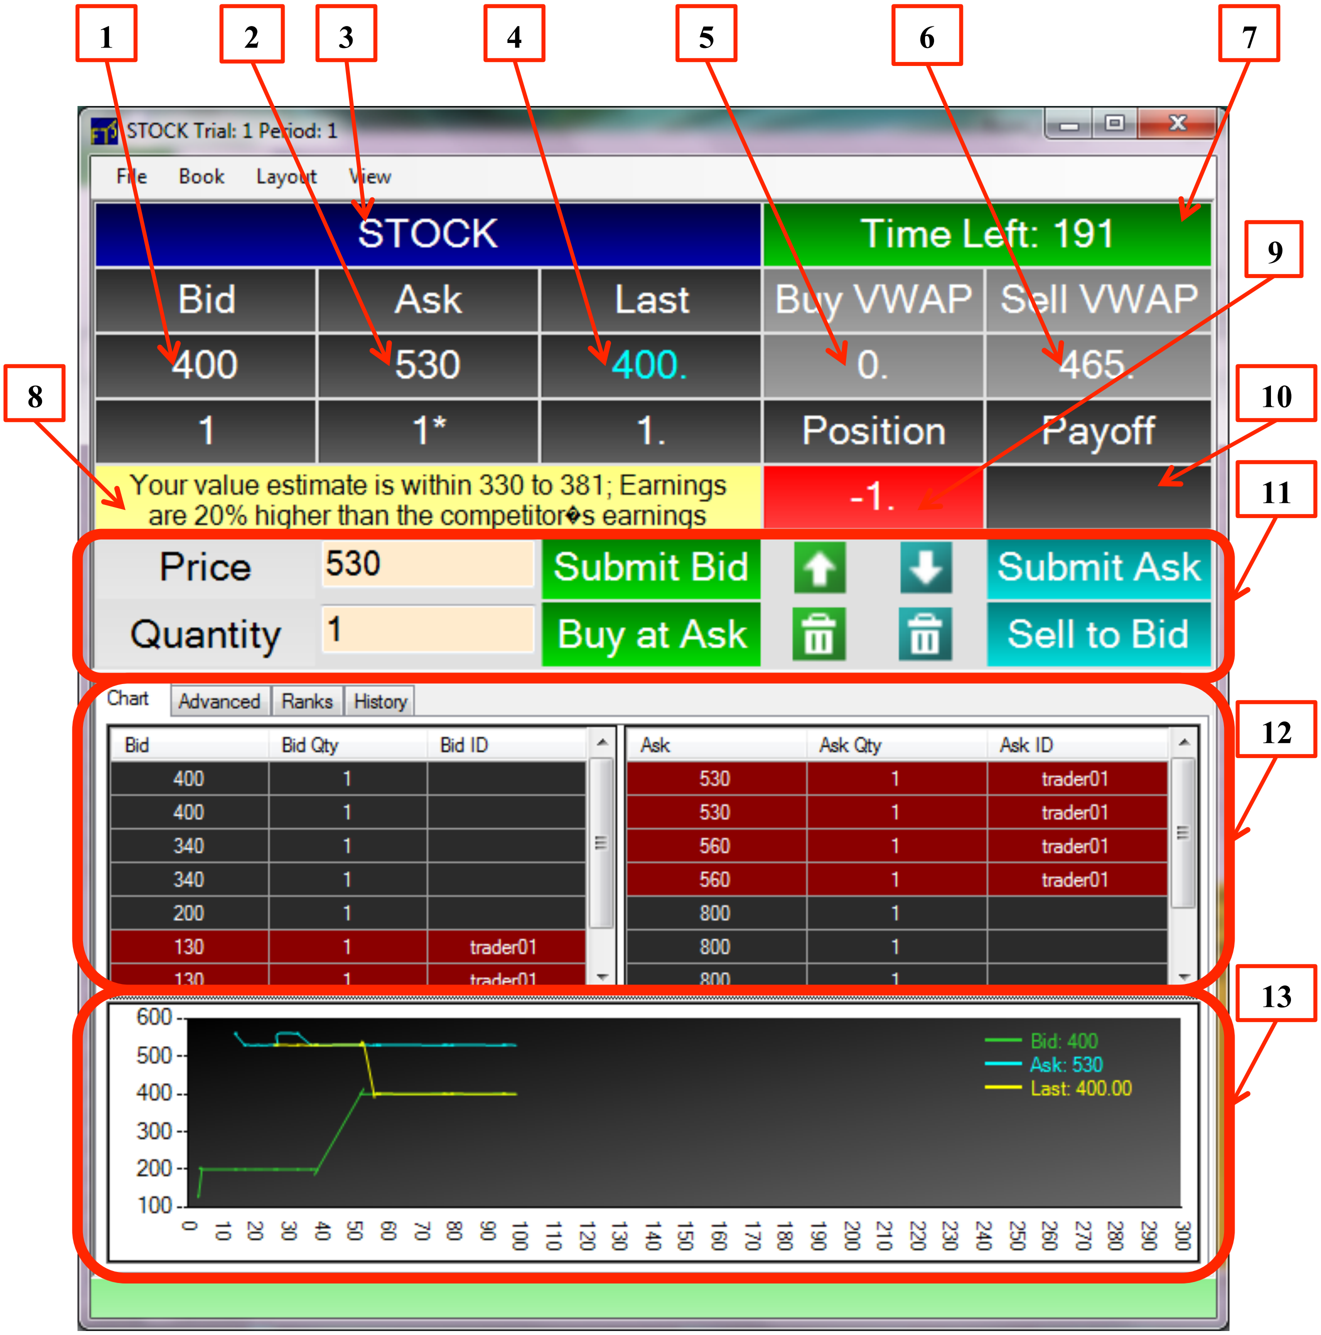
1
2
3
4
5
6
7
8
9
10
11
12
13
STOCK Trial: 1 Period: 1
File
Book
Layout
View
STOCK
Time Left: 191
Bid
Ask
Last
Buy VWAP
Sell VWAP
400
530
400.
0.
465.
1
1*
1.
Position
Payoff
Your value estimate is within 330 to 381; Earnings are 20% higher than the competitor�s earnings
-1.
Price
530
Submit Bid
Submit Ask
Quantity
1
Buy at Ask
Sell to Bid
Chart
Advanced
Ranks
History
Bid
Bid Qty
Bid ID
Ask
Ask Qty
Ask ID
Bid: 400
Ask: 530
Last: 400.00

Appendix II: How to Login

You must wait until the administrator tells you that you can login. Once he/she does, you must follow these steps to login successfully:

1. Double click on the "Launch FTS System Manager" icon on your desktop.
2. Check "Student Applications."
3. Check "Download again before running."
4. On the drop-down menu, select "FTS Trader 2013 Version."
5. Click on "Run Selected Application."
6. The following screen will appear:

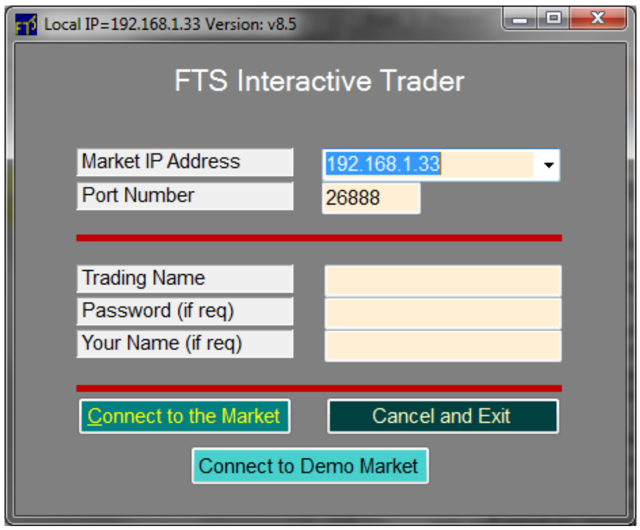

   Enter the information provided to you by the administrator on the "Market IP Address", "Trading Name" and "Password" (if required) cells. Leave all other cells untouched.

7. Click on "Connect to the Market"

   Note: if you click on "Connect to Demo Market" you can see and interact with a demo of the trading screen. You do not need to enter any information to connect to the demo version. In order to get to the trading screen as seen on Appendix I, you first need to left-click on the red/white icon located to the right of the security name. This may be a good tool to become familiar with the software before the trading session).

8. You should now see the main trading screen. Wait for further instruction from the administrator.