



ELSEVIER

Journal of Financial Economics 55 (2000) 425–459

JOURNAL OF
Financial
ECONOMICS

www.elsevier.com/locate/econbase

Can transparent markets survive?[☆]

Robert Bloomfield, Maureen O'Hara*

Johnson Graduate School of Management, Cornell University, Ithaca, NY 14853-4201 USA

Received 11 March 1998; received in revised form 19 April 1999; accepted 23 November 1999

Abstract

This paper investigates whether transparent markets can survive when faced with direct competition from less transparent markets. We first construct a game-theoretic model in which in equilibrium the low-transparency dealers capture early order flow, and use the resulting informational advantage to quote narrower spreads and earn more profits than their more transparent competitors. We then conduct a laboratory experiment that tests and supports all of these predictions. A second experiment shows that most dealers choose to be of lower transparency when they are allowed to do so. However, the informational advantage of low-transparency decreases as there are more such dealers, while the high-transparency dealers get increasing benefit from informed traders who attempt to broadcast deceptive trades. As a result, a small number of transparent dealers persist in our markets. © 2000 Elsevier Science S.A. All rights reserved.

JEL classification: G14; G12; D82

Keywords: Market microstructure; Experimental economics; Stock market competition

[☆]We would like to thank the editor (G. William Schwert), an anonymous referee, David Easley, Charles Schnitzlein, and seminar participants at the Utah Winter Finance Conference, Hebrew University, Tel Aviv University, University of Wisconsin-Madison, and the International Competition for Order Flow Conference for helpful comments. We gratefully acknowledge financial support from the Johnson Graduate School of Management and the National Science Foundation.

*Corresponding author. Tel.: + 1-607-255-3645; fax: + 1-607-255-5993.

E-mail address: mo19@cornell.edu (M. O'Hara)

1. Introduction

Where do you go to trade a stock? The answer for decades was simple: where it was listed. Today, however, this tautological response confronts a new reality due to the globalization of trading venues. Traders can choose to trade an emerging market security on its home market, or on a global market such as the New York Stock Exchange (NYSE) or the London Stock Exchange (LSE). This dual trading of securities is typified by Telefonos de Mexico (Tel Mex) which, while originally and still listed on the Bolsa in Mexico City, now has the bulk of its trading volume occurring on the NYSE.

Exchange-listed stocks routinely trade via crossing networks such as the Arizona Stock Exchange (AZX) or POSIT, while the proprietary trading systems Instinet and Island now clear an estimated 20% of the daily volume in NASDAQ listed stocks. Even the largest markets are not immune to this competitive pressure. Large blocks of NYSE listed stocks are now routinely traded 'off-board' by intermediaries who clear trades outside of the US after the market closes. Such trades are typically handled by large investment banks who buy the entire block as a principal, rather than as an agent. The trades are done after the New York market closes, and are executed outside of the US to avoid regulatory reporting requirements. The investment bank then resells the stock in a similar manner, with neither trade being reported on the consolidated tape. (For more discussion, see "More Big Stock Sales are Handled 'Off Board'" in the *Wall Street Journal*, April 9, 1997.)

An important feature of this new market competition is that it often involves differences in market transparency. While all US exchanges and the NASDAQ are required to report immediately all trades and quotes, many of these new markets require much less transparency. Instinet, for example, must display broker-dealer quotes, but need not provide quote information from other institutional traders. A new SEC rule, Regulation ATS, requires Instinet to report the best bid and ask quote even if that quote is from a non-market maker. This change does not apply, however, to smaller ECNs (or electronic communication networks) who are trading an increasingly larger share of NASDAQ volume. Lack of transparency is also cited as the main advantage of off-board trading, where even the existence of a trade need not be acknowledged.

Another example of trade transparency differences is found in the competition for large trades (known as blocks) between the London Stock Exchange and the Paris Bourse. Because many European firms list on both exchanges, traders have the option of trading on either venue. London allowed the reporting of block trades to be delayed in some cases for several days, effectively providing the low-transparency market modeled here. So much of the block trade volume gravitated to London that the Paris Bourse was forced to change their reporting rules. London and Paris have since harmonized rules. For more discussion of this issue, see Board and Sutcliffe (1995) and Gemmill (1996).

As markets compete head-to-head for trading volume, the role of market transparency becomes a crucial variable in their survival. In this paper, we investigate how transparency affects the competition between markets. Our study builds on a large body of research investigating how transparency affects market behavior (see Pagano and Roell, 1996; Madhavan, 1996; Naik et al., 1997; Biais, 1993; Lyons, 1996; Board and Sutcliffe, 1995, 1996; Gemmill, 1996; Porter and Weaver, 1996; Bloomfield and O'Hara, 1999; Flood et al., 1999; Lamoureux and Schnitzlein, 1997). An important difference found in our work here is that we focus on the simultaneous competition between markets of differing transparency. This allows us to investigate how transparency affects the equilibrium trading strategies of agents, and thus how transparency affects the resulting equilibrium.

The specific issue we address is whether markets that must report trades can survive when faced with direct competition from markets that need not report trades. This focus reflects the growing regulatory and market concern that the ability to hide trades in less transparent regimes is undermining the attractiveness of established markets, and thereby reducing their crucial price discovery role. Such hiding of trades is typified by the diversion of orders to off-board trading in London where reporting is not required. It is also consistent with trading in newly developing stock markets such as the Cayman Islands Stock Exchange and the Bermuda Stock Exchange whose trades are not part of the ITS reporting system. Our focus here also reflects our interest in the dynamic question of whether competition will force all markets to adopt the same transparency level, or whether market forces will allow different markets to co-exist.

We use experimental economics to address these issues. We design two laboratory experiments that include multiple dealers operating under varying degrees of transparency, and traders with differing trade motivations. An important feature of our markets is that traders can choose to hit multiple quotes, thereby allowing us to investigate how trading at the inside spread differs from trades executed off the best quote. Our experimental design is motivated by a game theoretic model that predicts how dealers will behave in a world of differing transparency requirements. Our first experiment includes two high-transparency dealers who must report trades and two low-transparency dealers who need not report trades. We use the data to investigate how transparency differences affect dealer spreads, inside spreads, informational efficiency, and the gains and losses of market participants. In our second experiment, we allow each dealer to endogenously select whether to be transparent or not. This allows us to examine the effects of varying numbers of transparent and non-transparent dealers on the market equilibrium, and to address the interesting and important dynamic effects on the population of dealers in transparent markets.

Our first experiment reveals that low-transparency dealers are more aggressive in setting quotes in early rounds of trading. This behavior allows them to

capture more order flow, and the dealers are then able to use their informational advantage both to quote narrower spreads than their more transparent competitors, and to avoid money-losing trades. However, these advantages decline as repeated rounds of trading force the low-transparency dealers to reveal their information through their choices of quotes. The data also show that although traders often choose to trade with low-transparency dealers at outside quotes, effectively paying a premium to keep their trades private, most of these traders are informed, making these trades typically unprofitable for the dealer. This result also underscores the plight of the transparent dealer who, lacking better information, makes losing trades at the outside spread more frequently. Our first experiment shows that concerns regarding transparency differences between markets are well-founded. Transparent markets do not thrive in competition with less-transparent venues.

Our second experiment examines whether the advantages of low transparency lead to the elimination of high-transparency dealers in markets in which transparency is optional. We find that dealers endogenously gravitate towards less transparency. When there are many non-transparent dealers, however, the informational advantage of any one dealer becomes small. Also, the remaining transparent dealers receive a new benefit: traders attempt to mislead the market by sending unprofitable trades to the transparent dealers. As a result, the transparent dealers remain competitive. However, more research will be necessary to determine whether these forces generalize to larger markets.

This paper is organized as follows. In Section 2, we develop a dynamic incomplete information game to analyze the decision problem facing high-and low-transparency dealers. This game theoretic model suggests the strategies that will be played by the dealers, and it outlines several hypotheses regarding the equilibrium outcome. We also develop hypotheses regarding how the strategies of informed and uninformed traders may differ with respect to transparency. Section 3 then details our first experiment, in which given numbers of transparent and non-transparent dealers compete. In this section we describe our experimental methodology, and we discuss how our experimental framework relates to our economic model and to the testing of our hypotheses. In Section 4, we test our hypotheses regarding dealer and trader behavior, and we present our experimental results. In Section 5, we present the design and results of our second experiment, in which dealers endogenously choose whether to be transparent or not. The paper's final section summarizes our results and discusses their implications for the competitiveness of markets.

2. Transparency and market competition

We begin our analysis by considering how market behavior will change when trade transparency differs between trading venues. We focus on two market

settings: a transparent market, in which trades are immediately reported and a low-transparency market, in which trades are not reported. Two types of forces seem most important in determining how transparent markets compete with low-transparency markets. Informational forces apply directly to dealers. These forces arise because different trade reporting rules allow dealers in low-transparency markets to gain monopoly power over information by capturing order flow that need not be disclosed. To the extent that this information can be profitably exploited, this advantage should affect the pricing strategy and profitability of low-transparency dealers, and, by extension, the pricing strategy and profitability of transparent dealers as well. Strategic preference forces apply indirectly to dealers through the strategic actions of traders. If liquidity traders and informed traders show different preferences for different types of dealers, then dealers may face different degrees of adverse selection, which in turn will cause them to set different quotes and earn different returns. Over time, one or both of these forces may result in one type of dealer being able to outdo the other in direct competition.

In this section, we construct a simple model to capture the informational forces affecting dealers. Because the problem we are addressing involves a multi-dealer, multi-stage pricing game with differential information, we develop a dynamic incomplete information game. The equilibrium concept we use is that of perfect Bayesian equilibrium, allowing us to develop hypotheses regarding dealer pricing strategies that are consistent with rational behavior on the part of the dealers. In the model, we do not explicitly include the strategic preference forces noted above. Instead, we take the traders' strategies as given, and then hypothesize how these strategies may differ given the dealers' behavior. Our goal in this section is to develop hypotheses consistent with rational behavior that we can then test using our experimental markets framework.

We consider a model in which an asset trades in two rounds, denoted $i = 1, 2$. The asset has an intrinsic value of $V \in \{0, 1\}$, where V takes on the values 0 and 1 with equal probability. At the beginning of each round, each of the four dealers sets a price or prices at which he is willing to buy the asset (the bid) and to sell the asset (the ask). Bids and asks are drawn from the closed interval $[0, 1]$. A trader, chosen at random, then executes a buy at the low ask or a sell at the high bid. With probability ε , the trader is informed, and buys whenever the security value is 1 and sells whenever the security value is 0. With probability $1 - \varepsilon$, the trade is from an uninformed trader who buys and sells with equal probability.

There are two markets in which trades may occur, and there are two dealers in each market.¹ In a high-transparency market, all trades are observable by all

¹ Note that computerized trade need not take place in a central location. One could, therefore, characterize our setting as two markets in which one can trade the same security, or one market that includes dealers with different reporting rules.

dealers in both markets. We denote the two dealers in the high-transparency market by t_1, t_2 . In a low-transparency market, trades are only observed by the dealer making the trade. We denote the two dealers in the low-transparency market by n_1, n_2 . More formally, we describe the trade at Time 1 as the vector $Z^1 = \{z, k\}$, where $z \in \{B, S\}$ denotes whether the trader bought or sold, and $k \in \{t_1, t_2, n_1, n_2\}$ indicates the dealer who traded. Once Z^1 is determined, all dealers observe the appropriate information from trades at Time 1, and set their trading prices for Round 2. Orders arrive for Round 2 and are executed, and the game ends.

To construct an equilibrium for the dynamic game, we work backwards from Round 2. That is, given the information sets the dealers bring into the second round, we solve for the dealers' optimal pricing strategies for trades at Round 2. Given these optimal strategies at Round 2, we then solve the game for the optimal strategies at Round 1. The following proposition characterizes any perfect Bayesian equilibrium of the game.

Proposition. Let $\pi^1(k)$ and $\pi^2(k)$ represent the unconditional expected profits earned by dealer $k \in \{t_1, t_2, n_1, n_2\}$ at Rounds 1 and 2, while $b^1(k)$ and $a^1(k)$ represent the dealer k 's bid and ask in Round 1. In any perfect Bayesian equilibrium:

Result 1: If $Z^1 = \{z, t_i\}$, then $\pi^2(t_j) = \pi^2(n_j) = 0$;

Result 2: If $Z^1 = \{z, n_i\}$, then $\pi^2(t_j) = 0$ for all j , $\pi^2(n_j) = 0$ for $j \neq i$ and $\pi^2(n_i) > 0$.

Result 3: $b^1(n_i) > b^1(t_i)$ and $a^1(n_i) < a^1(t_i)$.

Proof. A proof of the proposition appears in the Appendix.

The intuition underlying the results of the Proposition is as follows. If a high-transparency dealer trades in Round 1, then all dealers have identical information in Round 2. As in the model of Glosten and Milgrom (1995), Bertrand competition among the identically informed risk-neutral dealers eliminates any possible gains to trade (Result 1). However, if a low-transparency dealer trades in Round 1, that dealer becomes informed, gaining a monopolistic informational advantage over the uninformed dealers. The newly informed dealer can now earn greater profits (Result 2). Because the low-transparency dealers enjoy this potential gain in Round 2, they are willing to set higher bids and lower asks in Round 1 in order to secure an informational advantage (Result 3).

To see how the informed dealer extracts his informational rents in Round 2, keep in mind that the expected value of the security must be conditioned on the sequence of observed trades. Each observed buy order indicates a higher likelihood of a high value, because the order may come from an informed trader who would buy only when the security has a high value.

To reflect these conditioned expectations, let $EV_{hi} = E[V|B^1, S_2]$ be the expected value given that the informed dealer observed a buy in Round 1 and observes a sell in Round 2, and let $EV_{lo} = E[V|S^1, S^2]$ be the expected value given that the informed dealer observed a sell in round 1 and observes a sell in Round 2. Note that an uninformed dealer can never afford to bid greater than $b^m = (EV_{hi} - EV_{lo})/2$. Executing a trade at b^m yields an expected gain of EV_{hi} when the first trade was a buy, but yields an equal expected loss when the first trade was a sell. Because these events are equally likely, b^m is the highest bid that an uninformed dealer could make without loss. This situation implies that the informed dealer can guarantee a positive expected gain by bidding slightly more than b^m whenever he observes a buy order in Round 1.

Of course, such a strategy cannot be a perfect Bayesian equilibrium, because the uninformed dealers would respond by always bidding EV_{lo} , which would leave the informed dealer regretting his unnecessarily high bid. However, bids lower than b^m would allow an uninformed dealer to bid a slightly higher amount that is still less than b^m , and capture gains. As shown in the proof of the proposition, any perfect Bayesian equilibrium must involve randomized strategy on the part of both the informed and uninformed dealers, with both types of dealers choosing b^m with a strictly positive probability.

The proposition suggests several hypotheses regarding the behavior of transparent and low-transparency dealers. Formally, these hypotheses are as follows:

H1. Low-transparency dealers are more likely than high-transparency dealers to set market (inside) spreads in the first round of trading.

H2. Low-transparency dealers will clear more of the order flow in the first round of trading.

H3. Low-transparency dealers are more likely to trade profitably and less likely to trade unprofitably after the first round of trading.

H4. Low-transparency dealers will earn lower profits than high-transparency dealers in the first round of trading, and greater profits after the first round of trading.

H5. Low-transparency dealers estimate security values more accurately than high-transparency dealers after the first round of trading.

In particular, we would expect to find that the low-transparency dealers set initial inside spreads, reflecting the increased informational value of order flow to the low-transparency dealer. These better prices should also result in the low-transparency dealer clearing more trades in Round 1, but earning lower, or even negative, profits in that round. However, executing the trade results in more accurate beliefs and more profitable trading in future rounds.

An important feature of our model is that all dealers learn from watching market information. As a result, we hypothesize the following:

H6. The low-transparency dealers' advantages identified by H3–H5 decline as trading progresses.

While reporting rules may limit trade information, all dealers get to see the quotes that were posted by their competitors, and changes in these quotes should signal, at least partially, the information each dealer knows. This outcome suggests that the informational advantage enjoyed by low-transparency dealers should diminish as trading progresses. We test this claim by allowing trading to continue for many trading rounds.

Finally, the model predicts that high-transparency dealers never capture order flow in the first round of trading. However, experiments rarely yield such perfect separation of behavior across agents, even when predicted by economic models. If high-transparency dealers ever do capture order flow, Result 1 of the proposition implies that no dealer will capture any advantage, because all information about order flow is given to all dealers. The following hypothesis captures this belief:

H7. The low-transparency dealers' advantages identified by H3–H5 are greater when they capture more order flow in the first round of trade.

We predict that the differences identified above will be greatest when low-transparency dealers capture the most order flow.

Our analysis thus far has focused on the informational forces affecting dealers. As noted earlier, there are also strategic forces affecting dealers that arise from the trading decisions of the informed and uninformed traders. If traders are rational, then their trading decisions should minimize their trading costs. In a one-round game, this rationality would result in traders opting to trade with the dealer quoting the highest bid and the lowest ask. In our model above, this trading behavior reinforces our results on the benefits of capturing order flow by aggressive pricing.

In a multi-round framework, however, traders may engage in more complicated strategies. Madhavan (1995) argues that informed traders may wish to hide their information by trading in a low-transparency market. Uninformed traders with large trading needs may also prefer opacity if it allows them to trade with fewer immediate price effects. Madhavan's analysis predicts that more traders will opt to trade with low-transparency dealers, and all markets will eventually become low-transparency. Chowdhry and Nanda (1991) argue the opposite case. They contend that uninformed traders prefer to trade in transparent markets, while informed agents prefer to trade in low-transparency markets. This divergence would result in an adverse selection problem for the low-transparency dealers, and the worse prices in the

low-transparency markets would eventually drive all traders to the transparent markets.

Because it is so difficult to predict the strategies used by traders, we do not formulate specific hypotheses about their behavior. However, we do conduct a variety of analyses in order to identify any regularities in trader behavior that might affect dealers' incentives. The next section describes the experiments in which we examine trader and dealer behavior.

3. Experimental design

We now describe the nature of our experimental design and the specific features of our markets. The market is most similar to that developed in Bloomfield (1996) and Bloomfield and O'Hara (1999). Flood et al. (1999) also examine transparency issues in a similar setting, while Lamoureux and Schnitzlein (1997) examine such effects in a bilateral search framework.

In this section, we describe our first experiment, in which we investigate the competition between markets when there are a fixed number of transparent and nontransparent dealers. Our second experiment allows the mix of transparent versus nontransparent dealers to be endogenous, and we describe that experiment in a later section of the paper. We conclude this section with a discussion of how our experimental controls improve the power of the experiments and minimize the chance that nuisance variables might interfere with the testing of our hypotheses and the interpretation of the results.

3.1. Overview

Our analysis relies on a short list of conceptual terms, defined as follows. A cohort is a group of six subjects who always trade together. A security is a claim on a terminal dividend, and is identified by the value of the security and the traders' liquidity needs (described below). A market is a sequence of trading rounds for a single security. A session is a 3-h period during which a cohort participates in a series of markets.

In our markets, participants trade securities with values chosen randomly from a uniform distribution over the interval [200, 300]. All prices and values are denominated in a laboratory currency called 'francs'. Each security is traded for eight rounds. In every round, four dealers estimate the value of the security, and then set both bid quotes, at which they will buy shares in that round, and ask quotes, at which they will sell shares. The only restrictions on a dealer's bids and asks are that they must lie in the interval [190, 310], and that the bid price can be no greater than the ask price. The dealers see the quotes reported by the other dealers, but they do not necessarily see the trades made by the other

dealers. Dealers may revise their quotes, but only after they have executed any orders submitted to them at their current bid and ask.²

Traders observe the dealers' quotes, estimate the value of the security, and choose to buy from 0 to 10 shares at each of the four asks and to sell from 0 to 10 shares at each of the four bids. There are no wealth constraints or short-selling constraints in the market. After each round of trading, investors receive some information about actual trades. We manipulate market transparency by altering this information. Consistent with the structure of our theoretical model, the trades of the two 'high-transparency' dealers are revealed to all participants, and the trades of the two 'low-transparency' dealers are not revealed to the market. Dealers, of course, always know their own trades. This means that a low-transparency dealer knows both his trades and those of the two high-transparency dealers, but the high-transparency dealers know nothing about the trades made by the low transparency dealers.³

3.2. *Trader types*

The markets include one informed trader, one active trader, and two computerized liquidity traders.⁴ The informed trader learns the true value of each security at the beginning of each market, and she is the only participant to receive this information. The active trader is given a target number of 15,000 francs to raise or invest over the eight rounds of trading, and he is penalized for falling short of that target. Whether the active trader raises funds or invests funds is determined randomly for each security.

The computerized liquidity traders induce smaller bid–ask spreads by providing a significant flow of uninformed trade, which dealers can use to offset losses incurred by trading with informed traders. One computerized liquidity trader always buys 20 shares in each round at the low ask, while the other always sells 20 shares in each round at the high bid. We use non-random liquidity traders to keep the setting as simple as possible, and we assume that these traders have no opportunity to delay trading. This restriction, similar to that imposed in many theoretical models, reflects the notion that liquidity concerns can force traders to trade at a specific point in time, regardless of the prevailing market price.

² This structure captures the trading protocols found on the London Stock Exchange and on the NASDAQ SOES system whereby a dealer must be willing to trade a specified quantity at his quotes.

³ Note that no dealer in the market ever knows all trade information because the low-transparency dealers do not know each others trades.

⁴ This market structure is standard in microstructure models (see O'Hara, 1995) and is similar to that analyzed in Chowdhry and Nanda (1991).

3.3. *Subjects and incentives*

The participants in the experiments are primarily MBA students in the Johnson Graduate School of Management at Cornell University, and most have participated in previous laboratory market studies. Each subject attends an initial training session, and each receives detailed written instructions, a copy of which is given in Appendix B

Both dealers and active traders earn francs through their trading decisions and their estimates of security value. Every share bought results in a gain or loss of $(\text{Price} - \text{Value})$ and every share sold results in a gain or loss of $(\text{Value} - \text{Price})$, where Value is the true value of the security, and Price is the price at which the trade occurred. Every value estimate earns the participant a bonus of $50 - 2 \times |\text{Estimate} - \text{Value}|$, to a minimum of 0 francs. Thus, participants have an incentive to estimate the value of the security accurately in each round. However, the bonus is small enough, relative to trading gains and losses, that participants are unlikely to hedge their estimates and quotes against one another to reduce payoff risk.

Active traders also lose francs by failing to raise or invest the required number of francs over the five rounds of trading. Every franc by which an active trader fails to meet the target results in a penalty of three francs. This penalty is large enough that in most cases the active trader is better off trading to avoid the penalty than not trading to avoid a trading loss. For example, a trader who sells a share for 190 francs when it is worth 290 reduces his or her penalty by 570 francs, while incurring a trading loss of only 100 francs.

Cash winnings for each session are determined by subtracting a floor from each subject's franc winnings, and then multiplying the result by an exchange rate that converts excess francs into dollars. The floor and exchange rate are derived from pilot experiments conducted separately for each type of participant, and are designed so that each participant type would receive average winnings of approximately \$35 per session.

3.4. *Controlling for nuisance variables*

A primary benefit of experimental methodology is the ability to exercise control over features of the experiment that might influence behavior, but are not the focus of the study. Our experimental design controls for several such features, including differences across subjects. In particular, dealers who were high-transparency in Session 1 are low-transparency in Session 2, and vice versa. Active and informed traders switch roles in a similar manner. This aspect of the design minimizes the possibility that observed differences between high- and low-transparency dealers, or active and informed traders, could be driven by differences in the subjects playing those roles.

Our design also controls for the effects of characteristics of the financial markets that affect behavior, which are unknown but similar across treatments. For example, the wording of the instructions, the incentive system, and discrete security values may affect behavior in unknown ways. However, these unknown effects are unlikely to vary systematically across treatments. As a result, they cannot account for the differences across market treatments that are the focus of our analysis. To the extent that such effects increase the amount of random error in the data, they may make it more difficult to observe difference across cells, but they are unlikely to create such differences.

3.5. *Relating the experiment to the model*

The primary goal of our theoretical model is to characterize how transparency differences in markets affect the optimal strategies of dealers. As in any experimental setting, it is important to confirm that our results reflect a response to generalizable economic forces modeled in the laboratory setting, and do not represent idiosyncratic behavior that reflects misunderstandings of, or unfamiliarity with, the setting. We address these issues in several ways. First, we use only experienced subjects who are familiar with the experimental setting. All subjects in our experiments also participated in the experiments described in Bloomfield and O'Hara (1998, 1999). Second, we examine only the data collected in the second day of trading, when pre-equilibrium behavior is likely to be largely eliminated. Finally, we compare our results to our formal model. To the extent that laboratory data conform to equilibrium predictions, they seem likely to be generalizable.

There are a number of differences between the model and the experimental environment that might generate behavior that is consistent with the economic forces in the laboratory, but inconsistent with the predictions of our model. The model allows only one trade in each round, while the experiment allows every trader to trade with multiple dealers in each round. The model includes only two security values, while the experiment allows many. The model includes only two trading Rounds, while the experiment includes eight trading Rounds per security. However, we believe that these differences, while undoubtedly able to create some differences between laboratory and predicted behavior, are unlikely to generate results that are qualitatively different. Thus, we believe that we can interpret results that are qualitatively similar to our equilibrium predictions as generalizable support for the forces identified in the model.

A more important difference between the model and experiment is that the model assumes that the traders avoid engaging in strategic behavior. The experiment imposes no such restriction on strategic behavior. This difference actually represents one of the primary motivations for conducting an experiment. Experimental results that are qualitatively different from the predictions of the model might arise through strategic behavior, while results that are

qualitatively similar would indicate that strategic trading behavior does not alter the fundamental forces we characterize in the model. Direct analysis of trading behavior can clarify the nature of any strategic behavior and its effects.

4. Results

4.1. Behavior in Round 1

Hypothesis 1 posits that low-transparency dealers are more likely to set the inside market spread in Round 1. To test this hypothesis, we identify the number of inside quotes set by each dealer in each round. Ties between dealers may allow more than one inside ask or bid in a trading round. However, in the absence of ties, the Proposition would predict that low-transparency dealers set both inside quotes in each round, while high-transparency dealers set neither. To reduce the likelihood that our results are diluted by the subject being unfamiliar with the strategic forces they face, we analyze data from only the second day of trading. To compute statistical significance for this test and all subsequent hypothesis tests, we compute the average of the dependent variable within each of the seven cohorts, and then perform a *t*-test with these seven observations. This 'repeated measures' design reduces the problem, common in experimental economics, of overstating statistical significance by assuming that repetitions of the same actions by the same subject or group of subjects are independent events.

Panel A of Table 1 shows that in Round 1, low-transparency dealers are more than twice as likely to set the inside quote than are high-transparency dealers. Low-transparency dealers set the inside quote, on average 1.36 times per round, while high-transparency dealers set the inside quote an average of only 0.57 times per round. This difference is statistically significant ($p = 0.034$) using a one-tailed *t*-test. Thus, while low-transparency dealers set fewer than the 2 quotes/round predicted by our model, the data nonetheless provide strong evidence supporting Hypothesis 1. Panel B of Table 1 shows that differences in bid–ask spreads are less striking. High-transparency dealers' average opening spreads in Round 1 are 61, while low-transparency dealers' spreads are 51.67. This difference is only marginally significant ($p = 0.090$) using a one-tailed *t*-test. The results for spreads are weaker than the results for quote-setting, because capturing the inside quote in a given round requires merely bidding slightly more or asking slightly less than other dealers in that round. Variation in spreads across cohorts and securities may make it hard to observe differences in spreads, even when the low-transparency dealers consistently set the inside quotes.

Fig. 1 shows the overall pattern of bid–ask spreads over time. Spreads for all dealers decline as trading progresses, consistent with the reduction in adverse

Table 1
Quote setting behavior

Data for this and all subsequent tables are generated from a laboratory experiment involving groups to participants trading securities for profit. A dealer sets an inside quote if his bid equals the highest bid in that round or his ask equals the lowest ask in that round. If both are true, the dealer sets two inside quotes. Ties may make the total number of inside quotes per round greater than two. Averaging across cohorts trading different numbers of securities may make the average greater than or less than two. For this and all subsequent tables, *t*-statistics and *p*-values are computed by averaging the dependent variable over each cohort, and then using the resulting 7 means for statistical tests. '.NS' indicates *p*-values greater than 0.20. In Panel B, the spread is defined as the dealer's ask minus his bid. In Panel C, central values are defined as those in the interval [225, 275]; all other values are extreme. Profitable (unprofitable) quotes are those for which the dealer risks executing a profitable (unprofitable) trade, given the value, the quote and the direction of trade.

	Low-transparency mean (Std. dev)	High-transparency mean (Std. dev)	Difference mean (<i>t</i> -stat) <i>p</i> -value
<i>Panel A. Average number of inside quotes set by dealers</i>			
Round 1	1.36 (0.34)	0.571 (0.56)	0.985 (2.39) 0.034
Round 2	1.31 (0.16)	0.728 (0.17)	0.379 (4.70) 0.003
Rounds 3–5	1.13 (0.32)	1.05 (0.349)	0.079 (0.32) .NS
Rounds 6–8	1.17 (0.29)	1.13 (0.29)	– 0.04 (– 0.083) .NS
<i>Panel B. Average spreads</i>			
Round 1	51.67 (15.19)	61.00 (13.31)	– 9.33 (– 1.48) 0.090
Round 2	39.14 (9.72)	46.57 (8.20)	– 7.43 (– 1.83) 0.058
Rounds 3–5	27.27 (6.70)	32.41 (10.63)	– 5.14 (– 1.58) 0.083
Rounds 6–8	20.69 (7.47)	23.00 (12.96)	– 2.31 (– 0.52) .NS

Table 1 (continued)

	Low-transparency mean (Std. dev)	High-transparency mean (Std. dev)	Difference mean (<i>t</i> -stat) <i>p</i> -value
<i>Panel C. Inside quotes in round 2, segregated by extremity of security value</i>			
Central values, Profitable	0.90 (0.13)	0.10 (0.13)	0.80 (7.86) 0.001
Central values, Unprofitable	0.71 (0.26)	0.29 (0.26)	0.43 (2.13) 0.077
Extreme values, Profitable	0.75 (0.14)	0.27 (0.15)	0.48 (4.31) 0.005
Extreme values, Unprofitable	0.27 (0.19)	0.73 (0.20)	− 0.45 (− 3.01) 0.024




Fig. 1. Bid-ask spread by round of trading. This figure shows the mean level of bid-ask spreads (denominated in francs) over the eight trading rounds in each security. The light line indicates the spreads of the low-transparency dealers, while the heavy line indicates the spreads of the high-transparency dealers. The dashed line indicates the market spread, defined as the low ask minus the high bid.

selection the dealers face as they learn from trade outcomes. This reduction in spreads is predicted by microstructure theories (see O'Hara, 1995), and it has been found in previous experimental studies of dealer markets (Bloomfield, 1996; Bloomfield and O'Hara, 1999; Lamoureux and Schnitzlein, 1997). As Fig. 1 shows, the market spread, determined by the difference between the lowest ask and the highest bid in each round, is significantly narrower than the average spreads set by individual dealers.

Table 2
Trading volume per round by trader type

Active traders are human subjects trading to meet liquidity targets. Informed traders are human traders who know the true value of the security. Liquidity traders are computerized traders who always trade 20 shares at the inside bid and inside ask in each round. Statistical significance at the 0.1, 0.05 and 0.01 level in 1-sided tests are denoted by *, ** and ***, respectively.

	Low-transparency dealers	High-transparency dealers	Difference
Round 1			
Active traders	1.86	0.84	1.02
Informed traders	4.88	3.93	0.94
Liquidity traders	29.65	10.35	19.29**
Total	36.39	15.13	21.36**
Round 2			
Active traders	4.69	4.31	0.37
Informed traders	5.52	4.70	0.82
Liquidity traders	25.69	14.31	11.38***
Total	35.90	23.33	12.57***
Rounds 3–5			
Active traders	6.61	5.20	1.41
Informed traders	7.47	6.86	0.61
Liquidity traders	20.79	19.26	1.53
Total	34.87	31.32	3.55
Rounds 6–8			
Active traders	8.59	6.89	1.70
Informed traders	9.82	7.73	2.09*
Liquidity traders	19.44	20.48	– 1.04
Total	37.85	35.11	2.75
All rounds combined			
Active traders	6.50	5.17	1.33
Informed traders	7.78	6.55	1.23*
Liquidity traders	22.0	17.99	4.01
Total	36.28	29.72	6.57*

The dealers' pricing policies should naturally affect the order flow they receive. Hypothesis 2 states that low-transparency dealers will clear more of the order flow in the first round of trading than transparent dealers clear. To address this possibility, we calculated the average total order flow for low-transparency and high-transparency dealers by round. This data is given in Table 2. The data provide strong support for our Hypothesis 2. In Round 1, low-transparency dealers trade an average of 36 shares, while high-transparency dealers trade approximately 15 shares. As trading progresses, the imbalance

between the two dealer types narrows, and in later rounds the trading is approximately equal across the dealer types. Table 2 also stratifies the trades by trader type, revealing a number of differences in the composition of each dealer's order flow. We return to these composition effects later in this section.

4.2. Advantages of low-transparency in later rounds

If low-transparency dealers are learning from their trades in Round 1, then they should be able to set price on the profitable, or right, side of the market spread more frequently in later rounds (Hypothesis 3). To examine this prediction, we calculate the average number of inside quotes set by high- and low-transparency dealers for Round 2. The results of this calculation appear in Table 1 Panel C. The data are segmented into central values, between 225 and 275, and extreme values, which are less than 225 or greater than 275. This division is important, because both types of dealers may find it more difficult to learn about the security value when the true value is centrally located than when the true value takes on an extreme outcome.

Consistent with Hypothesis 3, low-transparency dealers set the spread on the profitable side of the market, buying when the true value is high or selling when it is low, more frequently than they do on the unprofitable side. This result is reversed for transparent dealers, suggesting that those dealers are more likely to be making losing trades than are low-transparency dealers. These effects are most pronounced when security values take on extreme values. For central values, the low-transparency dealers will only know the value is central, but may not tighten their spreads enough to induce trading by the informed trader. This behavior leads low-transparency dealers to set inside quotes about three times as often as the high-transparency dealers do, even for quotes on the wrong side of the market.

Hypothesis 4 concerns the profits earned by transparent and low-transparency dealers. Our theoretical model predicts that low-transparency dealers will earn less than transparent dealers in Round 1, but will earn greater profits than transparent dealers in later rounds. This profit outcome is due to the predicted greater aggressiveness of low-transparency dealers in soliciting order flow, and the exploitation of their resultant informational advantage in later rounds.

Table 3 presents dealer profits, segmented by dealer type and by trading round. The table presents strong evidence that low-transparency dealers earn substantially greater profits than transparent dealers earn. Calculated over all rounds of trading, low-transparency dealers earn an average of 164 francs, while transparent dealers average only 26. Our model predicts that low-transparency dealers will earn less in Round one, potentially earning negative profits, but the data show that both dealer types earn first-round profits that are positive and are not significantly different from each other. We know that low-transparency dealers execute more trades in Round 1, but since they do so at better prices they

Table 3
Average dealer gains

Dealer gains in each round are computed as $(\text{Value} - \text{Price}) \times \text{Trade}$, where Trade is positive for purchases and negative for sales. As in previous tables, t -statistics and p -values are computed by averaging the dependent variable over each cohort, and then using the resulting 7 means for statistical tests. '.NS' indicates p -values greater than 0.20.

	Low- transparency mean (t -stat) p -value	High-transparency mean (t -stat) p -value	Difference mean (t -stat) p -value
Round 1	390.3 (5.34) 0.002	262.4 (2.85) 0.029	127.9 (0.85) .NS
Round 2	370.7 (4.50) 0.004	47.05 (0.65) .NS	323.6 (2.32) 0.039
Rounds 3–5	140.2 (2.58) 0.041	– 19.0 (– 0.37) .NS	139.24 (1.88) 0.109
Rounds 6–8	42.4 (1.18) .NS	– 14.1 (– 0.90) .NS	56.5 (1.02) .NS

actually make less profit on each trade. Why low-transparency dealers do not show even more aggressiveness in pricing is puzzling. This result might reflect risk aversion, an ability to collude in price-setting, or simply a form of irrationality.

Table 4 also shows that low-transparency dealers are able to earn substantially greater profits than transparent dealers. This dominance is most pronounced in Round 2 where low-transparency dealers earn 370 francs on average while transparent dealers average only 47 francs, a result statistically indistinguishable from 0. In later rounds, the profits of low-transparency dealers decline, but remain strictly positive through Round 5, while the profits of transparent dealers continue to be approximately zero. As we predict, low-transparency dealers are able to earn greater profits than their transparent counterparts, but the advantage declines as trading progresses.

Our model attributes this higher profitability to an information advantage. This idea is captured by Hypothesis 5, which predicts that low-transparency dealers will estimate security values more accurately than transparent dealers after Round 1. Table 4 provides evidence in support of this hypothesis. In Round 1, both transparent and low-transparency dealers have the same estimation error. This outcome is to be expected, since both dealer types have the same information sets prior to any trading. In Round 2, however, the low-transparency

Table 4
Average error in dealer's value estimates

Errors in dealers' estimates in each round are measured as $|\text{Estimate} \times \text{Value}|$. *T*-statistics and *p*-values are computed by averaging the dependent variable over each cohort, and then using the resulting 7 means for statistical tests. '.NS' indicates *p*-values greater than 0.20.

	Low-transparency mean (<i>t</i> -stat) <i>p</i> -value	High-transparency mean (<i>t</i> -stat) <i>p</i> -value	Difference mean (<i>t</i> -stat) <i>p</i> -value
Round 1	87.8 (9.76)	91.71 (13.3)	3.8 (0.88) .NS
Round 2	84.7 (12.2)	96.7 (18.8)	11.99 (2.47) 0.049
Rounds 3–5	74.0 (226)	84.58 (30.94)	10.6 (1.67) 0.147
Rounds 6–8	52.5 (24.9)	57.2 (22.8)	4.68 (0.866) .NS

dealers are able to price much more accurately than transparent dealers, reflecting the information they have learned from their private trades in Round 1. This informational advantage appears to dissipate quickly. As seen in Table 4, the difference in pricing accuracy becomes statistically insignificant during Rounds 3 to 5, and it is virtually gone by Rounds 6 through 8.⁵ Results are very similar when we normalize errors for the dealer types in the first trading round. As a percentage of their first-round errors, the low-transparency dealers' errors are 83.6% in Round 2 and 71.8% in Round 3, while the high-transparency dealers display errors of 94.7% and 84.2%, respectively. The low-transparency dealers have a lower average error, relative to their first-round error, in each of the eight rounds of trading.

In summary, we have found strong support for the hypotheses regarding dealer behavior predicted by our model. Low-transparency dealers do act as informational monopolists by initially pricing aggressively and then exploiting their informational advantage in later rounds. Low-transparency dealers earn significantly greater profits than do transparent dealers, and they are able to set

⁵ Results are very similar when we normalize dealer types' errors by their errors in the first trading round. As a percentage of their first round errors, the low- (high-) transparency dealers' errors are 83.6% (94.7%) in round 2 and 71.8% (84.2%) in round 3. The low-transparency dealers have a lower average error (relative to their first-round error) in each of the eight rounds of trading.

prices more efficiently. These advantages dissipate over time because market participants are able to learn from other market data, eroding the value of private information. What we have not yet investigated are the strategic elements of trading, and, in particular, how the location preferences of traders affect the adverse selection confronting dealers. We now turn to this issue.

4.3. *Conditioning on the first round of trading*

The preceding results show that low-transparency traders enjoy, on average, an advantage over high-transparency dealers. However, Hypothesis 7 predicts that the advantage enjoyed by low-transparency dealers will depend on the outcome of the trade in the first round. In the model presented in Section 1, trade in the first round involves a single trader transacting with a single dealer. If the first-round trade was with a high-transparency dealer in such a limited setting, we could predict that low-transparency dealers would not show any advantages in later quotes, profits, or estimates over other low-transparency dealers, because all dealers would be identically informed.

Conditioning on the first-round trade is more difficult in our experimental setting, because multiple traders can each transact with multiple dealers. To determine how the trade in the first round affects the advantages gained by low-transparency dealers in later rounds, we conduct the correlation analysis shown in Table 5. We measure the extent to which total order flow coming from active and informed traders to the low-transparency dealers exceeds that coming to the high-transparency dealers. We refer to this measure as the trade differential. We then correlate the trade differential with the differences in dealer spreads, winnings, and estimate accuracy in the second round of trading. We define these differences such that we predict a positive correlation in all cases. All p -values are measured using a 1-tailed statistical test.

The results shown in Table 5 are strongly consistent with Hypothesis 7. The total trade differential, including trades with both informed and active traders, is correlated with differences in second-Round spreads, trading gains, and estimate accuracy. The correlations for these second-Round differences, all of which are statistically significant, are 0.30, 0.18, and 0.46, respectively. When we look only at trades with informed traders, we see that capturing more of the informed order flow in the first Round fails to allow low-transparency dealers in the second round to reduce spreads. ($\rho = 0.15$, $p = 0.091$). However, trade with the informed traders drives the dramatic effect of trade differential on gains. No significant correlation exists between active trade differential and gain differences, while the correlation with the informed trader differential is strong ($\rho = 0.19$, $p = 0.035$).

Partitioning the data by first-round trade differentials provides similar results (not reported here). For example, the difference in Round 2 gains is 950.8 francs when the first-round trade differential is 10 or greater, compared to only 560

Table 5
Round 1 correlations of trade differentials with round 2 variables

The informed trade differential in Round 1 measures the number of shares traded by low-transparency dealers with informed traders in Round 1, minus the number of shares traded by high-transparency dealers with informed traders in Round 1. The active trade differential is computed similarly using only trades with active traders. The total trade differential combines the trades of both active and informed traders. The spread difference in Round 2 reflects the average bid-ask spread of high-transparency traders minus the average spread of low-transparency traders. The gain difference in Round 2 reflects the average trading gain of low-transparency traders minus the average gain of high-transparency traders. The accuracy difference in Round 2 reflects the average absolute estimate error of high-transparency traders minus the average absolute estimate error of low-transparency traders. All variables are defined so that the average correlation between the variables is predicted to be positive. The entries indicate the Pearson correlations and one-tailed *p*-values of the trade differential and indicated difference across 81 securities. 'NS' indicates *p*-values greater than 0.20.

	Correlation with informed trade differential (<i>p</i> -value)	Correlation with active trade differential (<i>p</i> -value)	Correlation with total trade differential (<i>p</i> -value)
Spread difference	0.15 (0.091)	0.35 (0.001)	0.30 (0.034)
Gain difference	0.19 (0.035)	0.01 (.NS)	0.18 (0.056)
Accuracy difference	0.41 (0.001)	0.22 (0.024)	0.46 (0.001)

francs when the differential is less than 10. For estimate accuracy, the results are more striking. The difference in Round 2 estimate accuracy is 52.8 francs when the first-round trade differential is 10 or greater, compared to – 3.89 francs when the differential is less than 10.

Taken together, these results confirm that the average benefits in quote setting, gains, and estimate accuracy enjoyed by low-transparency traders are driven by their ability to capture order flow in the first round of trading. When the high-transparency dealers capture that order flow, the advantages gained by low-transparency dealers are strongly curtailed.

4.4. Strategic trading

We derived our hypotheses under the assumption that traders consider only whether they can make a profit in each trade, taken individually, without considering whether a trade in one round might affect prices and increase profits in future rounds. Such considerations might lead traders to pay a premium to conceal their trades by trading with low-transparency dealers (see Madhavan,

1995 or Chowdhry and Nanda, 1991). However, if these traders are informed, low-transparency dealers may face greater adverse selection.

In this section, we present several analyses indicating that strategic trading behavior is not a powerful influence in our markets. As a preliminary analysis, Table 2 presents the numbers of shares traded by each trader type in each round with low-transparency and high-transparency dealers. Liquidity traders trade approximately 30 shares in Round 1 and 26 shares in Round 2 with low-transparency dealers, compared to approximately 10 shares in Round 1 and 14 shares in Round 2 with High-transparency dealers. In each round, the difference for liquidity traders is significant at the $p < 0.05$ level. The informed traders show a slight preference for low-transparency dealers across all rounds, but this preference is never significant at the $p < 0.05$ level. The active traders never show a statistically significant preference for either dealer. The preference results are similar when we examine only securities with extreme values or only securities with central values. Overall, the data show that only liquidity traders have a decided preference for low-transparency dealers. However, these are the only traders that are unable to behave strategically, because they are simply computer algorithms. Liquidity traders simply trade at the best available quote.

We also examine the proportion of total trades that come from informed traders at both inside and outside quotes. We find that this proportion is approximately 10% for both types of dealers when they set inside quotes, but is approximately 70% for both types of dealers when they set outside quotes. The large difference between inside and outside quotes reflects the strong preference of active and computerized liquidity traders for the very best price. However, the similarity across dealer type provides no evidence of strategic trading.

We also examine trading gains and losses at the inside and outside quotes. As shown in Fig. 2, gains at the outside quote are roughly constant over all eight rounds, while gains at the inside quotes decline in concert with the decline in bid-ask spreads. Gains are largest in the early rounds because the wide spreads force liquidity traders to buy at very high prices and sell at very low prices. As spreads decline, liquidity traders transfer less of their wealth to the dealers.

A striking result is that neither high- nor low-transparency dealers are able to earn money at the outside quotes. Even though trades at outside quotes are executed at more favorable prices, the higher degree of adverse selection eliminates dealer profits. Low-transparency dealers make greater profits at inside quotes than do high-transparency dealers, but this advantage is simply driven by their superior price-setting information, as shown in Table 3. High-transparency dealers actually lose money trading at the inside quote after the second round of trading. This result suggests that high-transparency dealers are insufficiently aware of the adverse selection they face when competing against better-informed dealers. Previous research has shown that agents in laboratory settings




Fig. 2. Dealer gains over trading rounds. This figure shows the mean level of dealer gains, segregated by dealer type. The heavy (light) solid line indicates total gains to high-transparency (low-transparency) dealers. The heavy (light) dashed line indicates the gains to high-transparency (low-transparency) dealers at their outside quotes. Outside quotes are bids that are lower than the highest bid in that dealers' market in that round, and asks that are higher than the lowest ask in that dealers' market in that round.

often underestimate this type of adverse selection risk, often referred to as the 'winner's curse'. For example, assume the low-transparency dealers set high bids and asks because the informed trader has been buying shares from them. If the high-transparency dealers do not set their quotes sufficiently far apart to account for this possibility, they may unintentionally provide the low ask, and lose money to the informed trader. This type of error would appear to account for the losses suffered by the high-transparency dealers when trading at inside quotes.

As a final test for strategic trading behavior, we calculate the francs paid or received from shares traded at outside quotes that could have been traded at more favorable quotes. That is, suppose that a trader buys 5 shares of a security at 150 from a high-transparency dealer, and also buys 5 shares at 155 from a low-transparency dealer. Since the trader could have bought all 10 shares from the high-transparency dealer, this behavior represents a premium of 25 francs paid to the low-transparency dealer for the privilege of participating in a low-transparency market.

We find that the absolute size of the net premiums paid to low-transparency dealers are not statistically different from those paid to high-transparency dealers by both active traders and informed traders. This result suggests that traders are unwilling to pay premiums to have their trades concealed. Alternatively, it may be the case that informed and active traders occasionally attempt to fool the dealers by engaging in losing trades with high-transparency dealers. We calculated the average number of francs lost by informed investors when they buy above value or sell below value, and the average number of francs active traders raise when they need to meet an investment goal, or invest when they need to meet an earnings goal. Both numbers were much larger for trades with high-transparency dealers than for trades with low-transparency dealers,

a result consistent with the expectation that traders might attempt to fool the market by publicizing misleading trades. Thus, it appears that traders engage in two different types of strategic trading. Because these strategies offset one another, strategic trading does not alter the informational advantages of low-transparency markets described in our model.

4.5. *Summary of Experiment 1*

In summary, our first experiment demonstrates that transparency has real effects on dealer and market behavior. We found strong support for our hypotheses that low-transparency dealers are more likely to set initial inside spreads, are more likely to trade on the profitable side of the market spread, clear more of the early order flow, and are able to set prices in later rounds with greater accuracy. We also found strong support for the hypothesis that low-transparency dealers earn greater profits in later trading rounds, but we did not find the predicted result that their first-round profits were negative. Finally, our results provide little evidence that traders behave strategically in our setting.

By virtually every measure that we examine, it appears that low-transparency dealers are able to outdo high-transparency dealers in direct competition. Because low-transparency dealers earn greater profits, most or even all dealers might be expected to opt for less transparency if allowed to do so. In our second experiment, we investigate these issues by allowing dealers to decide whether they wish to have high or low transparency.

5. Experiment 2: Endogenous transparency

Allowing for endogenous transparency requires specifying some algorithm for determining dealers' choices. We use the following protocol. Before trading begins in each security, Dealer 1 chooses his or her dealer type. Dealer 2 then learns Dealer 1's choice, and makes a similar decision. This process is repeated for all four dealers, with each dealer knowing how many dealers previously chose to be of high or low transparency. The dealers then learn the total number of high-transparency and low-transparency dealers before setting quotes in the first round. We use this sequential design because a dealer's choice of transparency might depend on other dealers' choices, and a dealer's choice of quotes may depend on the total number of dealers of each type. Our algorithm allows dealers pricing decisions to depend on the number of high and low transparency competitors, but it does not allow for dealers to change their transparency designation depending on the decisions of subsequent dealers. Although more complex schemes might be interesting, they would greatly complicate determining the equilibrium number of dealers, and may lead to collusion among dealers

(see Cason, 1999). All other aspects of the market setting are identical to those of Experiment 1.

Experiment 2 involved six cohorts of investors, each trading ten securities, except for one group that traded only nine securities. For all groups, a practice session was held in which two securities were traded. To ensure that the results are robust across securities, each group traded a different set of securities, for a total of 59 distinct securities. Panel A of Table 6 shows the distribution of dealer types over the 59 securities. The data provide striking evidence that dealers have a clear preference for being low-transparency. For 25 securities, all four dealers chose to be low-transparency. Three of the four dealers chose to be low-transparency for 20 securities, while two of the four dealers chose low-transparency in 11 securities. In only 3 securities, do three dealers chose to

Table 6
Choices and payoffs in experiment 2

In Panel A, the entries indicate who many dealers chose to be of high transparency for each security in each cohort. In Panel B, * indicates that the gain is statistically significant at the $p < 0.05$ level in a 2-tailed test. No other trading gain averages are significant at the $p < 0.2$ level in a 2-tailed test. The benefit to increasing (decreasing) transparency indicates the expected gain from switching to become a High-transparency (Low-transparency) dealer in a market with one more such dealer.

Panel A. Number of high-transparency dealers by security and cohort

Security	Cohort	1	2	3	4	5	6	Total
1		2	0	0	3	3	0	8
2		2	0	1	2	2	0	7
3		1	0	0	1	2	1	5
4		2	0	1	1	1	1	6
5		2	0	1	2	1	0	6
6		1	0	0	1	1	0	3
7		1	0	0	1	0	0	2
8		1	0	0	2	1	0	4
9		1	0	0	2	1	0	4
10		–	1	0	3	1	0	5

Panel B. Trading gains and benefits to switching transparency levels

High-transparency dealers	Number of securities	Cohorts	Gains of low-transparency dealers	Gains of high-transparency dealers	Benefits to increasing transparency	Benefit to decreasing transparency
0	25	4	23.8		45.4	—
1	20	6	26.8	69.2*	– 61.0	– 45.4
2	11	3	– 59.0	– 34.2*	31.3	61.0
3	3	2	– 17.8	– 27.7	—	9.9
4	0	0	—	—	—	—

be high-transparency, and there is not a single security in which all four dealers chose to be high-transparency.

Several different tests show that the preference among dealers for low-transparency is statistically significant. Treating each security as an independent observation, only 50 out of 236 choices are for high-transparency (21.2%), a proportion that is significantly less than 50% ($p < 0.01$). Treating each cohort as an independent observation, we find that all six cohorts had more low-transparency dealers than high-transparency dealers. A one-tailed binomial test shows that this difference is significant at the $p < 0.02$ level.

The data provide less clear evidence on whether the high-transparency dealers will eventually be driven out of the market entirely. Panel A of Table 6 shows that the total number of high-transparency dealers in each market ranges from 6 to 8 in the first five securities traded by each cohort, and ranges from 2 to 5 in the last five securities traded by each cohort. However, the lowest number of such dealers is found in the sixth and seventh securities traded by each cohort. Thus, there is not a persistent trend toward the elimination of high-transparency dealers. No trend is statistically significant when using a repeated measures analysis that views each cohort as an independent observation. Further tests examining the choices made by dealers yield no significant trends or effects, other than a preference for low transparency over high transparency. Models designed to predict decisions on the basis of decisions in previous markets or decisions by previous dealers in the current market yield no significant associations other than individual and cohort differences.

The persistence of high-transparency dealers under these conditions suggests that there may be an advantage to being a high-transparency dealer when all of the other dealers are low-transparency. Panel B of Table 6 provides support for this possibility by calculating the payoff to each type of dealer in markets with different numbers of high-transparency dealers. The 'no data' cells in Panel B of Table 6 reflect that no markets were composed solely of 4 high-transparency dealers, preventing analysis of that possibility. The table shows that a sole high-transparency dealer in a market earns 69.2 francs. If given the opportunity, this dealer would not want to become a low-transparency dealer, because the low-transparency dealer in a market with no high-transparency dealers earns only 23.8 francs, generating a payoff reduction of 45.4 francs. However, no Low-transparency dealer in a market with only one high-transparency dealer would wish to switch to being a high-transparency dealer, because that choice would entail giving up a gain of 26.8 francs for a loss of 34.2 francs. Thus, no dealer wishes to deviate from an outcome in which exactly one dealer chooses to be transparent.

The benefit of being the sole high-transparency dealer in a market of low-transparency dealers seems to be twofold. First, although the high-transparency dealer still faces an informational disadvantage relative to the low-transparency dealers, when there are several low-transparency dealers participating

in the market, none has too great an informational advantage over the high-transparency dealer. This situation arises because the informed trader can spread trades more evenly among the low-transparency dealers, and therefore hide her identity more effectively. This behavior on the part of the informed trader serves to reduce the losses high-transparency dealers incur due to mispricing. Second, to the extent that the informed traders also intentionally engage in losing trades with the high-transparency dealer in order to mislead all dealers about the true security value, all of these gains are enjoyed by the sole high-transparency dealer. The sole high-transparency dealer earns 15.25 francs per round from such strategic trading, while the three low-transparency dealers earn only 2.65 francs.

Together, these two effects seem to suggest an equilibrium in which a small number of dealers willingly choose a high level of transparency, while the rest choose a low level of transparency. This outcome differs from the equilibrium predicted by Chowdhry and Nanda (1991), who argue that all dealers would choose to become transparent. Uninformed investors contribute to this difference in outcomes by being unable or unwilling to engage in strategic trading. Chowdhry and Nanda assume that such behavior will lead all traders to seek out the transparent venue, thereby forcing dealers to be transparent as well. We find little evidence of this effect. Instead, uninformed traders prefer the better prices in the less transparent setting, and gravitate toward those dealers. While some strategic advantages remain for transparent dealers, they are insufficient to offset their pricing disadvantage.

Three important caveats must accompany this evidence that transparent markets can survive. First, the payoffs that sustain the equilibrium are based on a relatively small number of observations. This leads to some payoff averages that are unusual but not statistically significant. For example, when dealers are allowed to choose their level of transparency and there are two dealers of each type, as in our first experiment 1, the low-transparency dealers appear to lose more money than the high-transparency dealers. However, because three cohorts never chose more than one high-transparency dealer per security, these averages are based on the behavior of only three cohorts, and the losses of the low-transparency dealers are not statistically significant. The losses of the high-transparency dealers in these markets are statistically significant, however, thereby providing results consistent with those of our first experiment.

Second, the sole transparent market-maker in a non-transparent market enjoys one distinct benefit in being the beneficiary of 'losing' strategic trades. This benefit seems likely to be less important in larger markets, in which a single trader is less likely to be able to influence the market. Finally, our markets may simply not have lasted long enough. With more chances to change and learn, it is possible that all dealers would eventually settle into a steady-state choice of low transparency. Thus, a full understanding of the long-term survival of low-transparency dealers will require further research.

6. Conclusions

Our research has shown that transparency has dramatic effects on market behavior. We found that low-transparency dealers are able to outdo their more transparent counterparts on a number of important dimensions, including profitability. What drives this dominance is the enhanced ability of low transparency dealers to set prices more efficiently, thereby allowing them greater opportunities to set and trade at the inside spread. This advantage derives from an increased aggressiveness on the part of low-transparency dealers to attract order flow, and from their subsequent ability to reduce their losses due to adverse selection. High-transparency dealers are handicapped both by reduced access to information, and by the strategic trading decisions made by uninformed traders making large trades.

Our result that low-transparency dealers produce better spreads than those created by their high-transparency counterparts accords with behavior observed in actual markets. Spreads on Instinet, for example, are routinely inside those of the Nasdaq, lending credence to those researchers who argue that less transparent venues can undermine transparent venues. Indeed, an intriguing feature of these Instinet spreads is that they are frequently set by NASDAQ market makers, who use the information they learn in the less-transparent market to extract rents in the transparent venue. Our results are also consistent with the experience of the Paris Bourse, which found itself losing trading volume in cross-listed shares to the less-transparent London Stock Exchange. As our results here demonstrate, traders will gravitate to markets with better spreads, and that behavior provides the basis for the greater success of less-transparent venues in comparison to transparent markets.

The trend toward low transparency in our markets provides an interesting contrast with the results of Bloomfield and O'Hara (1999), which showed that dealers earn higher profits in markets in which all dealers are transparent than in markets in which all dealers are of low-transparency. This result suggests that the choice of transparency presents dealers with a version of a social dilemma. All dealers might be better off if they were all transparent, but each individual dealer has an incentive to defect to low transparency. Dealers, therefore, might prefer a world in which transparency is mandated by regulation.

Can transparent markets survive in the absence of regulation? Our research suggests that transparent markets operate at a disadvantage to low-transparency settings largely because transparent markets face an informational disadvantage. In effect, the ability of dealers in the less-transparent market to know their own trades as well as those in the transparent market induces a type of free-riding problem. Not surprisingly, this incentive can lead to a failure of transparent markets to provide a competitive venue.

Despite these disadvantages, transparent dealers still survive in our markets. Apparently, the fragmentation of order flow when there are many

low-transparency dealers may reduce the informational asymmetry between markets to the point that it is inconsequential. Moreover, transparency comes with some advantages, the chief among these being that it provides a means of broadcasting one's trades. As we have seen, this publicity can result in traders purposely choosing the transparent venue in which to trade. Nevertheless, it is difficult to know whether this advantage would be important in very large markets, in which broadcasting may bring traders fewer advantages. As non-transparent trading alternatives proliferate, we must begin to question the dominance of transparent markets, if not their viability.

Appendix A. Proof of proposition

The proof of Result 1 is essentially identical to the proof presented in Glosten and Milgrom (1995), that identically informed dealers compete to eliminate all possible trading gains. We omit the formal proof. We prove Result 2 through a series of two observations. We characterize bidding behavior only. Proofs for asking behavior are analogous, and are omitted. At the end of the proof, we characterize one possible equilibrium outcome. Throughout the appendix, the informed dealer is the non-transparent dealer who captured order flow in the first round of trade; all other dealers are uninformed.

Observation 1. In the second trading round of any perfect Bayesian equilibrium, the informed dealer must randomize when observing a buy order, and at least one uninformed dealer must always randomize.

Proof. Let b_i and b_u denote the bids of the informed and uninformed dealers. Let $R_i[b_u]$ denote the informed dealer's best response when all uninformed dealers bid b_u , and let $R_u[b_i]$ denote an uninformed dealer's best response when the informed dealer bids b_i . In equilibrium, we must have $R_i[R_u[b_i]] = b_i$ and $R_u[R_i[b_u]] = b_u$. The following exhaustive cases show that there exists no equilibrium in which the informed dealer chooses a pure strategy:

- (i) If $b_i < b^m$, $R_i[R_u[b_i]] = R_i[b_u \in (b, b^m)] > b_i$,
- (ii) If $b_i > b^m$, $R_i[R_u[b_i]] = R_i[R_u[EV_{10}]] < b_i$,
- (iii) If $b_i = b^m$, $R_i[R_u[b_i]] = R_i[R_u[EV_{10}]] < b_i$.

To see that at least one uninformed dealer must also randomize, note that after observing a buy order, the informed dealer's expected payoff from a bid b is equal to the product $(EV_{hi} - b)F[b_u]$, where $F[b_u]$ represents the probability that the maximum bid of the uninformed dealers is less than b_u . This product

must remain constant over all of the bids comprising the support of the informed dealer's mixed strategy. If not, the dealer would place full weight on the best-paying strategy. Because $(EV_{hi} - b_u)$ increases as b decreases, $F[b_u]$ must decrease, which means that at least one uninformed dealer must randomize as well.

Observation 2. For both the informed dealer and at least one uninformed dealer, b^m must be the greatest bid that is made with strictly positive probability.

Proof. Simply redefine b_i and b_u as the maximum bid that receives strictly positive probability. Then results (i) and (ii) above carry through unchanged. Result (iii) does not carry through unchanged, because the best response of the uninformed dealers to a maximum bid of $b_i = b^m$ can still be $R[b_u] = b^m$. To see this result, write the expected payoff to the uninformed dealer as $0.5[EV_{hi} - b^m]G[b^m|EV_{hi}] + 0.5[EV_{lo} - b^m]G[b^m|EV_{lo}]$, where $G[b|EV_i]$ represents the probability that the informed dealer's bid is less than b given that the informed dealer expects a value of EV_i . Substituting $G[b^m|EV_{hi}] = 1$ and $G[b^m|EV_{lo}] = 1$ yields an expected payoff of 0.

To see why the informed dealer earns positive expected profit, note that the profit from bid b^m is greater than zero. The uninformed dealers earn zero expected profit because the payoff to a bid of b^m is exactly zero. If any other strategy generated a higher payoff, the dealers would place no weight on these strategies, which is an outcome that does not occur.

The proof of Result 3 is analogous to the proof of Result 1. Non-transparent dealers compete in Round 1 to capture order flow, thereby eliminating the gains from Round 2. This strategy forces non-transparent dealers to bid more and ask less than transparent dealers in Round 1. As in the case of Result 1, we omit the proof.

Finally, we construct an equilibrium for Round 2 as follows:

1. An informed dealer who observes a sell order in Round 1 bids no more than EV_{lo} .
2. Exactly one uninformed dealer bids with a probability distribution to satisfy $F[b] = (EV_{hi} - b^m)/(EV_{hi} - b)$ for $EV_{lo} \leq b \leq b^m$.
3. An informed dealer who observes a buy order in Round 1 bids b with a probability distribution to satisfy $G[b] = (b - EV_{lo})/(EV_{hi} - b)$ for $EV_{lo} \leq b \leq b^m$.
4. All other uninformed dealers bid EV_{lo} with certainty.

In this equilibrium, the informed dealer never bids more than EV_{lo} after observing a sell, because any greater bid will certainly lose money. For the informed dealer to engage in a randomized strategy after observing a buy order,

the gain to all positive-probability bids $(EV_{hi} - b)F[b]$ must be the same. Requirement 2 above shows that $F[b^m] = 1$, so the gain to trade at all bids must be $(EV_{hi} - b)F[b] = (EV_{hi} - b^m)F[b^m] = (EV_{hi} - b^m)$. Solving this equation for $F[b]$ determines the single uninformed dealer's strategy.

For the uninformed dealer to engage in a randomized strategy, the gain to all bids must be $0.5[EV_{hi} - b]G[b|EV_{hi}] + 0.5[EV_{lo} - b]G[b|EV_{lo}] = 0$. Substituting $G[b|EV_{lo}] = 1$, allows us to solve for $G[b|EV_{hi}] = (b - EV_{lo}) / (EV_{hi} - b)$ for $EV_{lo} \leq b \leq b^m$.

All of the non-randomizing uninformed dealers earn zero profit by always bidding EV_{lo} . These uninformed dealers have no incentive to defect, because they cannot improve their expected profit. In addition, an equilibrium exists in which all uninformed dealers choose the same strategy, and $F[b]$ represents the cumulative probability of the greatest bid from an uninformed dealer.

Appendix B. Laboratory stock market instructions to participants

Overview. In this market, you will trade shares of a number of securities. Each security has a value denominated in a laboratory currency called 'francs'. The value of each security is chosen randomly from the interval [200, 300] francs, with each value equally likely. Francs are converted into cash at the end of the session.

Types of Participants. The market includes four types of traders:

Dealers, four participants who set prices at which they will buy and sell shares;

Active Trader, one participant who must raise or invest francs by trading in each security, but can choose *when* to do so;

Informed Trader, one participant who knows the true value of the security;

Constant Traders, two computerized participants, one who buys and one who sells 20 shares in each round of trading at the most favorable prices available.

Sequence of Events. The sequence of events in the market is as follows:

Start Trading a New Security

Active Trader observes the number of francs he or she must raise or invest.

Informed Trader observes true value.

Trading round (8 rounds per security)

Everyone sees information about market activity in the previous round.

Dealers quote prices at which they will buy and sell.

Traders buy or sell at the Dealers' quotes.

End Trading in Security

All participants observe the true security value and their gain or loss from trading in that security. Market prices may change from one trading round to

the next, as may your estimate of the security's true value. However, the true value of a security never changes.

The Dealers' Task. If you are a Dealer, you will set three numbers in each round of trading:

Ask Quote: The price at which you will sell a share

Estimate: Your best estimate of value.

Bid Quote: The price at which you will buy a share

A quote is a promise to buy or sell as many shares as traders wish to trade at that price in that round; at the end of that round, you may set new bid and ask quotes. The bid and ask must lie between 190 and 310, inclusive. The ask must be above the bid—this protects dealers from unintentionally offering to 'buy high and sell low'.

Two Dealers' trades are always reported publicly at the end of each trading round. The other two Dealers' trades are private – they are never seen by anyone but themselves. *All Dealers' quotes* are reported publicly to all participants after all Dealers have set quotes.

The Traders' Task. Traders must enter their estimates of the security's value each Time they observe the Dealers' bid and ask quotes. Traders may trade up to 10 shares in each round at each available quote. Thus, a trader could buy 40 shares and sell 40 shares in each round.

There are two types of traders: Active Traders and Informed Traders. For each security, the Active Trader must either raise or invest at least 15,000 francs. Meeting this 'target' requires a sale or purchase of at least 50 shares over the eight rounds. Whether the Active Trader must raise or invest francs is determined randomly and independently for each security. The Active Trader is penalized 3 francs for each franc he or she fails to raise or invest. There is no penalty for exceeding the target (i.e. raising or investing more than required).

Informed Traders learn the true value of each security before trading begins for that security. Informed traders have no 'target'.

Constant Traders. Constant Traders are computerized traders. One buys 20 shares in each round at the lowest ask, while the other sells 20 shares in each round at the highest bid. If the two Dealers set the same bid, the constant trader who sells shares will sell to the Dealer who entered the bid first. If the two Dealers set the same ask, the constant trader who buys shares will buy from the dealer who entered the ask first.

Incentives. Market participants earn francs in three ways: by trading; by accurately estimating the value of the security; and by avoiding penalties for not meeting liquidity targets.

Trading. Whether you are a Dealer or an Active Trader, every share you buy increases or decreases your wealth in francs according to the formulas:

BUYING: Change in Wealth per share = (Value – Price)

SELLING: Change in Wealth per share = (Price – Value)

Estimates. In each round of trading, if your estimate is exactly right, you will earn a 10 franc bonus. For every franc your estimate deviates from true value, your bonus will decrease by 1 franc; thus, if your estimate is off by 3 francs, you will earn a bonus of $10 - 3 = 7$. If your estimate is off by more than 10 francs, you earn no bonus (but you are not penalized).

Penalties for not meeting trading targets. Active Traders are penalized 3 francs for each franc they fail to raise or invest. This penalty is large enough that it *always* outweighs trading losses incurred in meeting the target, no matter how unfavorable the market prices. Of course, it is still better to minimize trading losses or maximize trading gains, as long as this does not interfere with meeting the target. There is no penalty for buying or selling more than required.

Converting Francs to Dollars. Your cash winnings for each session will be determined by the following equation:

$$\text{CASH} = (\text{francs won} - \text{'floor'}) \times (\text{Exchange Rate}).$$

You will not learn either the 'floor' or the exchange rate until the experiment has ended. However, both are independent of your winnings. Different agent types have different floors and exchange rates, designed to allow each type to earn an average of about \$30/session. *Your cash winnings are not affected by whether other participants do well or poorly.*

Winnings are determined separately for each session. Losses in one session do not offset gains in another session. No matter what happens, you will receive *at least \$10* for completing each session.

I consent to participate in this experiment, and agree to abide by all of the rules determined by the experiment coordinator throughout my participation. I recognize that: (1) if I breach any of the rules governing the market, I forfeit my right to any money I might have earned by participating; (2) I have the right to leave the experiment at any Time, without penalty, but that in doing so I forfeit my right to any money I might have earned by trading; (3) this experiment has been approved by the Cornell University Committee on Human Subjects as research that uses no deception of any kind.

Instruction Supplement for Experiment 2

Overview. In this session, all aspects of trading are the same as in the previous session you attended, with the following exceptions:

1. Instead of having two private and two public dealers, each dealer gets to choose whether he or she wishes to be private or public. Dealers make a new decision for each security. Dealer 1 chooses first, then dealer 2, and so on. Each dealer will know the previous dealers' choices before making their own decision. After all dealers have chosen their types, everyone will learn how many public and private dealers there are.
2. The group will trade two practice securities for 3 rounds each, followed by twelve 'actual' securities for 8 rounds each (a total of 14 securities).
3. There are new floors and ceilings.
4. There is a new software tool that allows the active trader to determine how a given set of trades will affect the remaining trading target.

References

- Biais, B., 1993. Price formation and equilibrium liquidity in fragmented and centralized markets. *Journal of Finance* 48, 157–183.
- Bloomfield, R., 1996. Quotes, prices and estimates in a laboratory market. *Journal of Finance* 51, 1791–1808.
- Bloomfield, R., O'Hara, M., 1998. Does order preferencing matter? *Journal of Financial Economics* 50, 3–37.
- Bloomfield, R., O'Hara, M., 1999. Market transparency: who wins and who loses? *Review of Financial Studies* 12, 5–35.
- Board, J., Sutcliffe, C., 1995. The effects of trade transparency on the London Stock Exchange: A summary. Unpublishing working paper. London School of Economics, Financial Markets Group.
- Board, J., Sutcliffe, C., 1996. The proof of the pudding: The effects of increased trade transparency in the london stock exchange. Unpublished working paper. London School of Economics, Financial Markets Group.
- Cason, T.N., 1999. The opportunities for conspiracy in asset markets organized with dealer intermediaries. *Review of Financial Studies*, forthcoming.
- Chowdhry, B., Nanda, V., 1991. Multimarket trading and market liquidity. *Review of Financial Studies* 4, 483–511.
- Flood, M., Huisman, R., Koedijk, K., Mahieu, R., 1999. Price discovery in financial markets: The effect of alternative communication schemes. *Review of Financial Studies*, forthcoming.
- Gemmill, G., 1996. Transparency and liquidity: a study of block transactions in the london stock exchange under different publication rules. *Journal of Finance*, 1765–1790.
- Glosten, L., Milgrom, P., 1995. Bid, ask and transaction prices in a dealer market with heterogeneously informed investors. *Journal of Financial Economics* 13, 71–100.
- Lamoureux, C., Schnitzlein, C., 1997. When it's not the only game in town: The effect of bilateral search on the quality of a dealer market. *Journal of Finance* 52, 683–712.
- Lyons, R., 1996. Optimal transparency in a dealership market with an application to foreign exchange. *Journal of Financial Intermediation* 5, 225–254.

- Madhavan, A., 1995. Consolidation, fragmentation, and the disclosure of trading information. *Review of Financial Studies* 8, 579–603.
- Madhavan, A., 1996. Security prices and market transparency. *Journal of Financial Intermediation* 5, 255–283.
- Naik, N., Neuberger, A., Viswanathan, S., 1997. Disclosure regulation in competitive dealership markets: Analysis of the London Stock Exchange. Unpublished working paper. London Business School.
- O'Hara, M., 1995. *Market microstructure theory*. Blackwell Publishers, Cambridge, MA.
- Pagano, M., Roell, A., 1996. Transparency and liquidity: A comparison of auction and dealer markets with informed trading. *Journal of Finance* 51, 579–611.
- Porter, D. C., Weaver, D., 1996. Pre-trade transparency and market quality. Unpublished working paper. Marquette University, Milwaukee, WI.