First draft: July 23, 1997
Current draft: June 6, 2001

To Err is Human
–But Dealers' Abilities to find non-Dominated Strategies Depends on Market Transparency

Christopher G. Lamoureux
and
Charles R. Schnitzlein*

In this paper we compare subjects' abilities to converge to equilibrium strategies across market structures in an information environment suggested by the financial market microstructure literature. In our common value asset markets with asymmetric information, dealers quickly converge to non-dominated strategies in a centralized transparent market. In a fragmented, opaque market we observe persistent dealer losses amongst experienced subjects. Dealers attempt (unsuccessfully) to learn about intrinsic value by trading aggressively. We thus observe mechanism dependent behavioral departures from rationality in the setting where dealers cannot observe the behavior of other dealers. With sufficient experience and feedback agents do manage to uncover non-dominated strategies by gradually refining their approach, but not by experimenting with dramatically new approaches.

* Department of Finance, College of Business and Public Administration The University of Arizona, Tucson, AZ 85721, phone 520-621 7488, e-mail: lam-oureu@lamfin.bpa.arizona.edu crsl@u.arizona.edu, phone 520-615 0389. We would like to express our appreciation to The Business Law and Economics Center of Washington University, as well as The John M. Olin School of Business for generous financial support. Thanks, too to Ed Dyl, Ronald Huisman, Ron King, Terrance Odean, Mike Weisbach, and seminar participants at the University of Arizona, for comments on earlier drafts.

1 Introduction

Transparency is a fundamentally important aspect of trading mechanism design. Although recent research has focussed primarily on the implications of transparency for market quality, it has also uncovered implications for welfare by trader type. For example Pagano and Röell (1996) build a one-shot model that shows noise traders benefit from more transparent mechanisms, although the assumption of Bertrand competition among dealers leaves dealer profits unaffected. Madhavan (1995) shows in a sequential trade model that dealers benefit from the nondisclosure of trades because they can use private information from past trades to profit on future trades. Bloomfield and O'Hara (1999) use a controlled experiment to examine this issue in a richer setting and find that trade disclosure increases spreads while enhancing dealer welfare.

In this paper our focus shifts to the influence of transparency on the ability of dealers to find equilibrium strategies in a standard market microstructure setting – using the informational structure of the seminal work of Kyle (1985) - and Glosten and Milgrom (1985). We report the results of experimental markets employing two trading mechanisms that differ primarily on the degree to which trades and quotes are disclosed in an information environment suggested by the market microstructure literature. In both settings an insider knows the liquidation value of a risky asset, liquidity traders have exogenous, inelastic demands, and dealers provide liquidity in the pursuit of trading profits. In the first setting all trade is initially funneled through competing dealers with complete quote and post-trade transparency. The dealer market is then shut down, and a period of bilateral trade is permitted that both bypasses and is hidden from the centralized dealer market. Subsequent to this bilateral search market the dealer market is then reopened for an additional interval of centralized transparent trade. The second setting is completely fragmented and opaque. All trade (including trades involving dealers) is the result of bilateral negotiations, with no general market activity reported. We thus compare markets that differ dramatically on the degree to which dealers can observe the actions of other dealers.

In the setting where dealer behavior is transparent, dealers quickly adopt non-dominated strategies in

the sense that although dealers compete aggressively, the dealer industry does not incur losses.[1] In sharp contrast, dealers in the opaque setting (who are not compelled to trade) incur large losses that persist over many sessions with experienced subjects. We conjecture – and provide evidence to support the notion - that dealer losses are due to a bias toward action motivated by an attempt to learn the identity of the insider and the direction of his/her information, with the hope of exploiting it later.[2] Dealers thus display overconfidence with respect to their ability to extract valuable information from other agents who also act strategically. Importantly, dealer behavior only changes very gradually as dealer losses mount. Dealers do not experiment with dramatically different strategies in the face of losses, but rather gradually demand greater price concessions for taking trades, and become less willing to take imbalanced inventory positions. Ultimately, with sufficient feedback and experience, dealers give up their attempts to compete with the insider and dealer losses disappear.

We believe these results are interesting for several reasons. First, given the rapid evolution of trading systems around the world, facility of use is an important aspect of trading mechanism design. A second related point is that historical mechanism choice and evolution may reflect facility of use. Our findings may help explain the historical preponderance of centralized, transparent markets such as those that prevail in large economies.[3] Third, financial economists are increasingly interested in behavior that deviates from that of a fully rational agent using standard choice axioms. We show – in a decentralized search market - how agents learn to play undominated strategies; the manifest behavioral biases are not hard-wired. This point is also relevant for financial economists using experimental methods to learn about the effects of mechanism design on *equilibrium* outcomes.

The behavioral lesson from this paper is similar to that of Friedman (1998). Friedman shows that in

[1] Since theoretical results are unavailable in even this relatively stylized setting we focus on the attainment of undominated strategies. A necessary condition for equilibrium is procedural rationality: no traders are experiencing sustained avoidable losses.

[2] Agents explain their strategies in a questionnaire at the end of each session.

[3] Commenting on the distinctions between economics and psychology, Smith (1991) suggests that laboratory behavior which violates standard axioms of choice may arise because of the artificiality of the setting or context: "I want to suggest that perhaps the structures we observe have survived because of their merit in coaxing Pareto-

an unfamiliar (non-market) setting, subjects behave sub optimally, but that with additional training and feedback, the behavior converges to optimal. Friedman goes so far as to suggest that all behavioral anomalies are not the result of hard-wired human foibles or biases, but rather result from inadequate experience and feedback. This paper may be seen as providing microstructure support for Friedman's premise. But there is a bigger issue: How much experience and how much feedback are required to achieve convergence to optimality? A related question is the role that the mechanism itself plays in facilitating or obstructing this convergence.

Indeed, the study of behavior does not resemble finance or economics in that psychologists have amassed a body of anecdotal evidence, but there is no behavioral paradigm or unifying theory to organize this evidence.[4] As such it is difficult to discuss behavior outside of the context of a particular mechanism or setting. This paper adds to the anecdotal evidence on behavior within the microstructure context by isolating the effect of market transparency on the ability of dealers to find non-dominated strategies.

Another example of context-dependent behavior comes from the auction literature. It is well known that auction mechanisms that are strategically equivalent generate significant behavioral differences.[5] In addition, subjects' ability to change from dominated to non-dominated strategies is mechanism dependent. For example, in common value auctions there is evidence that bidders subject to the winner's curse make significant within-session reductions to their bids (Lind and Plott (1991)), while buyers facing adverse selection in bilateral bargaining games do not significantly adjust their behavior until they return for a second session or gain experience in the role of a seller (Ball, et al. 1991). Part of this difference may be explained by the market-level signals available to bidders in the common-value

---

efficient behavior out of agents who do not know what that means." (Smith, p. 894)

[4] We asked a colleague who is an esteemed applied psychologist about whether this is a fair characterization of modern psychology. His response: "Your statement is correct and is so much a part of modern psychology that nobody would cite a source, it just is.'' (Lee Beach, May 24, 2001) When pressed, Lee Beach asked Terry Connolly (another esteemed colleague in decision theory), and "We concluded that the concept is so presumed that it is like gravity, there really isn't a citation newer than Newton." (Lee Beach, May 25, 2001)

[5] For example, in a private value setting, first-price auctions consistently generate higher revenues than Dutch auctions and revenues in second-price auctions are higher than in English. It is also well known that some auction mechanisms attenuate (but do not eliminate) the problems of winner's curse and adverse selection in common value auctions with inexperienced subjects. Kagel (1995) provides an excellent survey.

auctions (Lind and Plott reported all bids and private signals in addition to the winning bid). In the bilateral bargaining game, by contrast, subjects only learn from their own actions. We also find evidence that the availability of market level signals plays an important role in helping subjects avoid dominated strategies.

This paper had its genesis in our attempt to construct a laboratory environment that could be used to study the effect of market fragmentation and transparency on market performance using an information structure widely employed in (and thus interpretable within the context of) the large market microstructure literature. Toward that end we ran experimental markets in three settings: the combination centralized dealer and bilateral search market and the pure bilateral search market that we compare in this paper, and a pure dealer market where all trade was funneled through competing dealers. In the two mechanisms that had all dealer activity centralized and transparent, dealers quickly adopted non-dominated strategies despite the complexity of the adverse selection problem. We attribute this to the use of sophisticated subjects, an intensive training session that involved training in all three mechanisms, and the use of performance in the training session to select the actual subjects for the experiment (we eliminated almost 40% of the potential subjects in the training session because of their relative difficulty in mastering the series of experimental environments). We followed these procedures because of our interest in making equilibrium comparisons across mechanisms. We therefore excluded the pure bilateral search markets from the initial analysis (Lamoureux and Schnitzlein (1997))[6] since the bulk of the data in these markets was inconsistent with equilibrium behavior.[7]

Other recent experimental work has examined the influence of transparency on market outcomes. As

[6] Comparing the first and second settings allowed us to investigate the impact of fourth market trading activities on the quality of the primary market. Our primary finding was that market quality improved in the presence of the alternative trading venue: dealers competed more aggressively against the alternative trading venue than they did with each other when it was not available. Increased competition lead to lower dealer rents and more efficient transaction prices.

[7] The reason for comparing the two dealer consolidated mechanisms with one another is because both are equilibrium outcomes – or at least it is clear that no one is playing an obviously dominated strategy. We choose to compare the fragmented market –where it is clear that dealers took much longer to arrive at a non-dominated strategy - with the hybrid mechanism because in both there is a heightened adverse selection problem, relative to a pure

noted, Bloomfield and O'Hara (1999) find that transparency has important implications for dealer welfare: in the opaque setting dealer losses exceed those of strategic liquidity traders (despite there being only two competing dealers), while in the transparent setting dealer profits exceed those of informed traders. Flood, Huisman, Koedijk, and Mahieu (1999) examine the effect of market transparency in a setting characterized primarily by interdealer (zero-sum) trading and find that dealers in an opaque setting employ more aggressive pricing strategies than dealers in a transparent setting. Neither study is designed nor intended to examine the influence of experience on dealer strategies in the opaque setting. In a similar vein to this paper, Pouget (2000) finds that the greater pre-trade transparency of the Walrasian tâtonnement facilitates the playing of equilibrium strategies relative to a call auction. He employs simple mechanisms that permit the derivation of game-theoretic equilibrium strategies, and shows that transparency helps subjects find equilibrium strategies.

The remainder of this paper is organized as follows. Section 2 describes the experimental design. In Section 3 we demonstrate that in the transparent mechanism dealers quickly settle on undominated strategies, while in the opaque setting dealers aggressively trade in order to learn about intrinsic value, they incur persistent losses that diminish very gradually overtime, and that finally (in a sixth experimental session) with additional feedback about the profitability of the dealer industry, they attain undominated strategies. We conclude in Section 4.

2 Experimental Design, Subjects, and Procedures

The experimental markets are organized into two treatments or mechanisms, with each mechanism distinguished by the trading rules governing the trading of a risky asset. Since we are interested in studying the effect of trading mechanism design on market outcomes, the information structures and the distributions of all random variables are identical across both treatments. We limit all inter-treatment differences to those due directly to differences in the trading mechanisms we examine. We explain below

---

dealer consolidated mechanism. In both of the mechanisms compared in this paper the insider may trade directly with liquidity traders – an impossibility in the pure dealer mechanism.

the features that are common to all settings, and then explain in detail in Section 2.4 the rules governing exchange under each trading mechanism.

There are three types of agents: a single informed trader who learns the end of period value of a single risky asset, three competing market makers, and four uninformed traders who trade to satisfy an exogenously determined liquidity motive. The information structure is modeled on Kyle (1985). Although in the experiments, while uninformed traders have an exogenously determined liquidity demand, they choose the timing of trades themselves: Order arrival is endogenous – not specified exogenously as in many microstructure models.

2.1 Communications and Computer Displays

All interactions between subjects were conducted on a series of networked personal computers with custom designed software. The computer screen in front of each dealer and trader contains continuously updated market information that (as will be explained below) is trading mechanism specific, but includes a record of each trader's own trading activity and cash balances, and the time remaining in the current trading interval, in both treatments. In addition, writing materials were provided to all subjects so they could track any additional information that they felt was relevant.

2.2 Agents, Trading Profits, Parameter Values and Variable Distributions

Currency: All price and profit information throughout the paper is expressed in lab dollars (L$). The exchange rate is 20 lab dollars to US$1.0.

Risky **Asset Value:** The asset value each market period is drawn from an approximate normal distribution with mean of L$100, standard deviation of L$8.7, and support on integer values between L$70 and L$130, both inclusive. There is a riskless asset (cash) with an interest rate of zero.

Informed Trader: At the beginning of each market period, one of the five traders is chosen randomly to be the informed trader. The informed trader learns the end-of-period asset value prior to the start of trade, and may engage in trades over the course of the market period to earn trading profits. There are no constraints on informed trader activity. Trading profits are carried forward to subsequent periods.

Dealers: Each of the three dealers is free to engage in trades in pursuit of profits. At the beginning of each market period, dealers neither learn the end of period asset value nor receive required positions. Each period, trading profits or losses of each dealer are added to starting cash balances and carried forward to the next period.

Traders' Required Positions: Each of the four remaining traders is required to buy or sell a randomly determined net number of assets. Each trader's required position is determined by an independent draw from a uniform distribution of the non-zero integers between -5 and 5 (endpoints included). If a trader does not meet this requirement, a penalty is assessed at the end of the period equal to one dollar times the absolute value of the deviation between the required position and the actual end of period position. The magnitude of the penalty ensures that demand is inelastic at the required position. The trading profits or losses of each trader are added to starting cash balances each period, and carried forward to the next period, so traders have an incentive to minimize trading costs.

Trading Profits: All transactions are for a single unit of the single risky asset. Trades are a zero sum game. The profitability of each trade is determined by the difference between the trading price and the end of period asset value. Trading profits on trade t (Π_t) are defined as: $\Pi_t = Q_t(V - P_t)$; where Q_t is +1 for a transaction in which an agent buys and -1 for a sale, V is the end of period asset value, and P_t is the transaction price.

Endowments and Payments: All agents begin the first market period with a cash balance of L$520 or US$26. Trading profits (losses) are carried forward to subsequent periods. At the end of the final market period, the current cash balance is multiplied by 5%, and subjects are paid that amount. Given the zero sum nature of the trading game, cash payments (net of any penalties incurred by traders) average US$26 per subject.

Information Sets: The informed trader learns the asset liquidation value prior to the start of trading, and the other agents do not know the identity of the informed trader. At the end of each market period, all agents learn the asset liquidation value. Only each individual trader learns his/her required position. All information pertaining to distributions, parameters, payments, and the rules governing trade is common knowledge.

2.3 Subjects and Procedures

The 24 subjects who participated in the experimental markets were recruited from a preliminary training round. The training round consisted of training in the information structure, variable distributions, and parameter variables, and two market periods of trade under each of three trading mechanisms (only two of the three mechanisms are used in this paper). It lasted, on average, slightly more than three and a half hours. There were eight subjects in each of the five identical training sessions. Of the forty subjects that participated, eight were undergraduates recruited from a senior-level finance class, 28 were MBA students recruited from a second year finance class, and four were business Ph.D. students.

We used performance in the training session as a criterion to select 24 subjects for the actual experimental markets. Of the 24 students (from the training sessions) initially invited to participate in the experimental markets, 23 accepted. From the pool of 24 subjects, we randomly formed three cohorts of eight subjects (each of the three classes of students was represented in each of the cohorts). We then assigned each of the cohorts to a single mechanism. The experimental data we report below is drawn from five (six in the ease of the fragmented/opaque market) additional two-hour trading sessions under each trading mechanism.

We provided all subjects with detailed written instructions on the trading rules, parameter values, and distributions. Each session was preceded by a review of these instructions. Since all sessions were completed in a maximum of eleven days, the review became progressively shorter as the subjects became more experienced, and averaged a half-hour prior to the first session under each treatment, and five minutes prior to the fifth session. In the first session, subjects were randomly assigned roles. In the second and third sessions, random role assignment was maintained with the exception that there was one randomly selected dealer in each session with previous experience as a dealer. In the fourth, fifth and

sixth sessions, role assignments were purely random.[8] At the end of each session, each subject was paid in private. None of the subjects ever went bankrupt - all left the session with a cash payment, and none had any interim adjustments to their cash balances. In each of the experimental sessions we set the number of market periods to the number that could be completed in a two-hour session after preliminary instructions. We thus have a laboratory database consisting of five or six experimental sessions under each trading mechanism, each of which includes from seven to ten market periods. We exclude the first four market periods from the first session, and the first market period from subsequent sessions to allow for re-familiarization with the market process. Session 6 is constructed by using the same draws as in Session 3, after informing the eight agents that dealers had experienced losses (at the industry level) in all previous sessions. This leaves a total of 27 market periods for the transparent dealer-cum-search market and 35 for the opaque pure search market. The numbers of market periods reported from the six sessions are, respectively: 3, 5, 5, 7, 7, and 8. The asset values each period are the same for both mechanisms.

2.4 Trading Mechanisms

The trading rules pertaining to each trading mechanism are summarized below. As explained above, all random variable distributions and other procedures not specific to the trading mechanism are constant across the treatments.

2.4.1 Fragmented/Opaque (Search) Market

A bilateral search market period is exactly twelve minutes long with no trading halts once the market period begins. During the market period, each trader or dealer may attempt to initiate a trade by sending a trade proposal via computer. A trade proposal indicates the type of trade desired (buy or sell), the trader or dealer to whom the proposed trade is to be directed (T1,T2,T3,T4,T5,D1,D2,D3), and the proposed

[8] The focus of this experiment is learning: how agents learn to play a market making game, and how the market structure influences this process. Therefore the experimental design does not control for "cohort effects." We have a group of 8 subjects under each mechanism each of which requires a total time commitment of at least 14 hours. Meta-analysis of experimental data suggests that cohort effects are more pronounced when subjects are inexperienced. Outcomes and actions are more homogeneous the more experienced the subjects.

trading price. Upon receiving a trade proposal, the recipient may accept the proposal or allow it to expire. If not accepted, all proposed trades expire within 25 seconds of being sent. In addition, the sender may revise trade proposals at any time. Trade proposals can be made to any combination of potential counterparties simultaneously, and there is no limit on the number of trade proposals that may be sent over the trading interval. There are also no costs (or subsidies) associated with proposing or accepting trades.

If a trade proposal is accepted, each participant's screen reflects current positions and trading prices along with who initiated the trade. Only those two agents who participated in that transaction observe trades completed via the direct trade process. No general market activity is reported on participants' screens.

2.4.2 Consolidated/Transparent Dealer Market

The transparent dealer market opens as a pure quote driven mechanism. In this setting, all trades are routed to the dealer with the highest bid (for a sell order) or the lowest ask (for a buy order). The trading clock is not started until all three dealers have entered a bid and ask quote. The computer screens display the three bids and asks, and highlight the inside quote. The identity of the dealer behind each quote is not revealed. Dealers are free to revise their quotes once the market is open. After the first 120-second trading interval there is an eight minute bilateral search market identical to the mechanism described above with the exception that dealers are prohibited from participating. During this search period, dealers' screens are blank, and only the two participants to a trade know that that trade took place (and therefore the terms of that trade). A second 120-second pure quote driven market follows the bilateral trade search market. Trading activity tracked on traders' computer screens includes all transactions completed during the quote driven portions of the market period. We thus refer to this mechanism as the transparent dealer market to highlight the contrast from the fragmented market – as concerns the activity of dealers, (it is of course more complex than a pure dealer market).

Dealers have a temporal monopoly franchise in the transparent dealer market. In this setting, when

the dealers can trade, all trade is routed through the dealers. During these periods, all agents - including those dealers not on the inside of the bid-ask spread, observe all trades and quotes, as well as changes in quotes. In the opaque market, by contrast, at any point in time, a dealer is simply a trader whose identity warrants to all other agents that she has neither a liquidity shock nor an informational advantage The fragmentation means that dealers have to compete against one another at all times to collect information (i.e., to trade). In the transparent market, dealers as a group must compete with the bilateral search mechanism - prior to that period - but all dealers observe any trade with any dealer.

2.4.3 Economic Comparisons of the Mechanisms

The primary difference between the two mechanisms is that when dealers post quotes and trade in the transparent dealer market, they do so openly. In the opaque market no dealer can observe any of the interactions between the other dealers and traders. The message space is larger in the opaque market than it is in the transparent market. In the transparent dealer market, traders can either accept the best dealer quote (by initiating a trade), or do nothing. In the opaque market, traders and dealers can initiate contact with one another by sending a bid or ask to a counterparty, and then entering into a series of proposals and counterproposals. The features of transparency and fragmentation are closely related, and these two mechanisms differ from one another on both dimensions. We make no attempt to disentangle these two dimensions in this paper.

The adverse selection confronted by the dealers is essentially the same in the two mechanisms. Dealers must compete against the liquidity traders and the informed trader to provide liquidity. Post-search in the transparent market, and throughout the opaque market (especially after the first few minutes), the dealer confronts a heightened adverse selection problem relative to the standard dealer market model (of Kyle (1985) or Glosten and Milgrom (1985)). Any trade that a dealer takes is one that the informed trader could have taken, but chose not to, or it involves the informed trader. This is because in both mechanisms the traders can trade amongst themselves. Dealers in the two mechanisms have the

same prior beliefs about the asset value and the liquidity shocks. Dealers in the opaque market may collect different information than dealers in the transparent market. In the transparent market, all dealers know that they all observe the same phenomena, whereas in the opaque market dealers may collect private information.

3 Results

Dealers in the transparent market immediately settle on undominated strategies, while dealers in the opaque market incur significant losses until the fifth session. Figure 1 gives an overview of dealer profitability by session under the two mechanisms. Differences in dealer profitability are dramatic. While the dealer industry loses money in less than one-half of the first 20 trading periods over the first four sessions under the transparent mechanism, with identical random draws for asset values and liquidity shocks, under the opaque mechanism, the dealer industry losses money in 18 of 20 trading periods. In addition, none of the six individuals that assumed the role of dealers over the first four sessions managed to make a profit. The initially large dealer losses incurred under the opaque mechanism are clearly incompatible with equilibrium as the dealers could always sit out and do nothing.

The results are presented as follows. First we present summary statistics and analysis that explains the determinants of dealer profitability across both mechanisms. Then we analyze how inexperience and a lack of transparency interact to generate dominated dealer strategies. The *possibility* of extracting private information from the order flow plays an important role, and leads to aggressive trading strategies that appear due to overconfidence in their ability to gain useful information through trading. We then show how dealers change their strategies and avoid losses with sufficient experience.

3.1 The Determinants of Dealer Profitability

Table I provides the basic information for each trading period for the five sessions for both the opaque and transparent dealer markets. As noted above, the first 27 trading periods correspond to the first

five sessions. These sessions are identical - across mechanisms in the experience of the subjects and the random draws. The sixth session (last eight periods) under the opaque mechanism was not a part of the original experimental design (which called for five sessions under each of the three mechanisms). Before this session we announced to all subjects that the dealer industry had lost money in each of the previous five sessions. The random draws are a repeat of the third session, and there is no counterpart of this session from the transparent market.[9]

The most striking feature is the fact that dealers made an average profit of 3.27 (with standard error of 8.54) in the transparent market, but lost an average of 29.21 (8.35) in the analogous market periods of the opaque market. Liquidity traders lost on average 48.4 (9.36) in the transparent market, and 22.03 (8.54) in the opaque market. Dealers traded with liquidity traders 152 times in the opaque market, and 212 times in the transparent market. Dealers traded with informed traders 96 times in the opaque market, and 138 times in the transparent market. Dealer losses are higher in the opaque market despite a smaller number of trades with the insider because the insider took a large proportion of profitable trades with liquidity traders: trades where the liquidity trader was trading in the opposite direction of the insider's information. The average root-mean-square-error (of all trades, relative to the true asset value) is 6.25 (0.75) in the opaque market and 6.44 (0.60) in the transparent dealer market. There were 22 intra-dealer trades in the opaque market (an impossibility in the transparent market).

To isolate the determinants of dealer profitability, we regress aggregate dealer (insider) profit each period on the absolute value of the difference between the realization of the asset value in that period and its unconditional expectation (of 100) (Extremeness), the product of the period's realization of the net liquidity shock and asset value less its unconditional expectation (Correlation), and a dummy variable meant to capture low experience vs. high experience. "Correlation" summarizes whether the liquidity traders are tending to trade in the same direction as the informed trader, as well as the magnitude of the

[9] This is an admittedly ad hoc way in which to examine the effect of public information with respect to the dealer industry profitability on the behavior of individual dealers. We of course do not know how dealers would have

shocks. If, for example, the asset value is 120 and the net liquidity shock is +4, then this variable is (120 - 100) × 4. In order to examine learning effects we add an experience indicator that takes on the value of zero in periods 1-20, and one in periods 21 - 35. For comparison, this indicator variable is one in the transparent market in periods 15 - 27. The results from these regressions are presented in Table II. We note that experience has a statistically significant effect on dealer profitability under the opaque mechanism, but that it has no effect under the transparent mechanism. Value has a statistically significant negative effect on dealer profitability in both mechanisms: when the asset value is further from its unconditional expectation dealers make less. The correlation variable has a statistically significant negative effect on dealer profits in the opaque market. The effect of this variable on dealer profitability in the transparent market is positive and statistically significant. This finding suggests a fundamental difference between the two mechanisms.

In the transparent market, when liquidity traders tend to trade in the same direction as the insider, the insider's information is worth less than otherwise. This also tends to facilitate the dealers' inference problem. There is an intermediate period in the transparent market where traders can provide liquidity to one another, and in the case where the liquidity traders want to buy and the asset value is high, then the informed trader will not provide liquidity to the liquidity traders (who would want to trade close to 100 - the unconditional expectation - or the middle of the dealer's spread at the start of the bilateral trade period). There will therefore be more trades for dealers after the search period, when spreads are wide, and volume elasticity is high. In the opaque market, on the other hand, the inference problem is much harder, as dealers can be picked off individually. When the liquidity traders are trading in the same direction as the informed trader it is as if there are more informed traders. Dealers will tend to trade more (see below) because the informed trader does not want to provide liquidity to the liquidity traders around the unconditional expectation of the asset value. This negative effect of the correlation between the two random shocks in the opaque market does not diminish with high experience. When we introduce an

---

behaved in the sixth session without this information. From Figure 1 we see that by the time a dealer has had two

interaction between the high experience dummy variable and the correlation variable, the coefficient is statistically insignificant, and there is no change in the results. One possible explanation for this set of results is that a highly experienced dealer in the opaque market is not willing to trade close to the unconditional expectation. She revises her beliefs about the asset's value in response to the order flow that she can observe, but recognizes that the information in that order flow is very coarse. The dealers therefore set wide enough "spreads" to profit - averaging over the independent distributions for the asset value and the liquidity shocks. What is different between this mechanism and the transparent mechanism is that the inference process in the latter mechanism is precise enough to enable a positive correlation between the net liquidity draw and the asset draw to enhance dealer profitability. This inference is coarser in the opaque market, to the point where such a positive correlation actually hurts dealer profitability.[10]

Unlike its effect on dealer profitability, the effect of this correlation variable on insider profits is the same in the two mechanisms. As expected, when liquidity traders want to trade in the same direction as the informed trader, the monopolistic access to information is worth less, independent of the mechanism. It is telling that just the asset value and this correlation variable explain over 72% of the variation in insider profitability in the transparent market. Also, the effect of experience on insider profitability is not statistically significant under either mechanism.

3.2 Why are dealer losses initially high under the opaque mechanism?

Dealers initially adopt dominated strategies in the opaque market, where we observe four sessions with significant dealer losses, whereas they behave consistent with equilibrium behavior in the transparent markets. In the opaque market there is what we refer to as a two-sided winner's curse as the

---

previous experiences in that role (prior to session 6), dealer losses disappear.

[10] While not the focus of this paper (since we do not claim to compare equilibrium outcomes across institutions), this finding is consistent with some theoretical results concerning market transparency. In Pagano and Röell's (1996) model for example, bid-ask spreads are wider when dealers cannot observe transactions that involve

dealers make unprofitable trades to obtain private information by trading early in the period, and then assume that this private information is more valuable than it is in later periods. This hypothesis is supported by subject questionnaires, which were completed immediately after each session. A typical comment was: "I wait and try to figure out who the informed trader is and then trade in front of him." This calls to mind that the private information in this setting concerns not only the asset value, and liquidity shocks, but also inferences with respect to who is informed.

Examining inter-dealer trades may test evidence that dealers believe they have gained private information through aggressive trading in the low-experience case. A feature of the opaque mechanism is that dealers may trade with one another. Inter-dealer trade is an important feature in dealer-driven stock markets, such as Tokyo, London, and even NASDAQ, foreign exchange markets (e.g., Lyons 1996),[11] and even used car markets (Genesove 1995).

Table I reports the number of inter-dealer trades in each market period under this mechanism. In the low-experience trading periods (1-20), dealers trade with one another 19 times. In only two of these transactions do the inventories of both counterparties to the trade flatten (or move toward zero). In five transactions both inventories move away from zero, and in the remaining 12 transactions one dealer's inventory flattens and the other moves away from zero. In the sessions when the dealers were highly experienced (periods 21 –35), there are six inter-dealer trades. Two of these trades flatten the inventories of both dealers, one moves both dealers' inventories away from zero, and three flatten one dealer's inventory and increase the other's.

These results support the evidence from the questionnaires suggesting low-experience dealers behave as if they posses valuable private information. Risk - or loss - aversion (see footnote 13, below) can explain inter-dealer trade when each dealer is flattening her inventory. The only rationale for a risk-

---

other dealers. What we see in the lab is that the inability to observe such trades hinders the dealer's ability to learn from trades. The dealer in turn protects herself by posting wider spreads: exactly as in Pagano and Röell.

[11] In fact, Lyons (1996) argues that risk-averse dealers would prefer some opaqueness in the price discovery process to facilitate inventory rebalancing in an inter-dealer market.

averse dealer to engage in a trade with another dealer that increases an inventory imbalance is that she believes that she has better information than her counterparty. While such trades are commonplace in the early sessions, they are very rare after dealers gain more experience.

The adverse selection problem is also more severe in the consolidated market period subsequent to the period when the traders trade with one another. This period of inter-trader trade, from which dealers are excluded, affords the opportunity to the informed trader of "picking off" liquidity traders seeking to trade in the opposite direction of the informed trader, conditional on the current conditional expectation. Thus, if the value is 120, we might expect that the informed trader would trade with all of the liquidity traders that had to sell, at close to the conditional expectation at the end of the first period of consolidated trading. As was shown in the paper that compared this mechanism to a pure quote-driven mechanism, dealers had no difficulty recognizing that this is indeed a time of heightened adverse selection, and set much wider spreads and responded much more to the order flow than in the pre-search period.

3.2 How do dealers learn to avoid losses under the opaque mechanism?

While all dealers initially lose money, all dealers also gradually modify their strategies, until dealer industry losses diminish to zero. The largest and most significant reduction in dealer losses occurs between dealers assuming the role for the second and third time ($p<0.01$). When dealers have two previous experiences as dealers, losses become insignificant (Figure 1). With three previous experiences as dealers (or public information that the dealer industry has incurred losses in previous sessions) dealer losses are again insignificantly different from zero. We thus find agents initially sub optimally collect information at too high a price, overweight the information, but then gradually change their behavior as losses mount.

As they become highly experienced, the most important change in dealer behavior in the opaque market is dealers dramatically reduce the number of trades they make close to the unconditional expectation of the asset value early in the period; they give up their attempts to gain private information.

There are two ways to see this. First, in the first three minutes of each period, liquidity traders trade an average of 1.85 times per period with dealers in the first four sessions, and 1.15 times per period in the last two sessions. Similarly, dealers and the informed trader trade 0.67 times per period in the first three minutes of each period in the first four sessions and 0.38 times per period in the last two sessions. The average RMSE relative to the asset value's unconditional expectation (of 100) in the first three minutes of the first four sessions is 3.84 when dealers trade with liquidity traders and 1.65 when dealers trade with the informed trader. The corresponding values in the last two sessions are 9.77 and 13.86. With the heightened experience, the dealers no longer attempt to attract trades by offering to trade at prices close to the unconditional expectation.

In general, as dealers gain experience, they trade less, and stay closer to a flat inventory. Aggregating the information in Table I, we note that the average number of dealer trades, per period, by session declines monotonically over the first five sessions: 13.3, 12.8, 8.8, 8.7, 5.7, and 6.4. Average dealer inventory declines monotonically over the six sessions. Define average inventory as the absolute value of each dealer's inventory at the time a dealer trades – excluding the first two trades. By session, this variable is: 1.91, 1.66, 1.62, 1.58, 1.04, and 0.89.

Table III presents the results of a linear regression of the number of trades involving dealers in a period on the same three exogenous regressors in the preceding regressions, for the opaque market. It is clear that highly experienced dealers trade less. As expected, when the asset value is more extreme, ceteris paribus, dealers will trade more. They also tend to trade more as the liquidity traders want to trade in the same direction as the informed trader, although this effect is marginally significant. We have documented above that as dealers gain experience and widen their spreads and trade less they stop behaving as if they have private information.

From Table I, we see that this increase in dealer profitability comes primarily at the expense of liquidity traders. The average liquidity trader loss per period in Periods 1 - 20 is 16.44. This average loss is 41.58 in the last 15 periods. The insider's average profit per period is 55.45 in the first 20 periods, and

35.23 in the last 15 periods.

The significant losses experienced by the dealers over the first four sessions are not the result of irrationality, but of an evolution to a procedurally rational strategy (i.e., one where they do not expect losses). Figure 2 shows how the dealer profits evolved over the 6 rounds of play, after accounting for the determinants of profits described above. After controlling for the value shock and the correlation between the value shock and the liquidity shock, the per-period profits are positive in rounds 5 and 6, and negative in rounds 1 through 4. The regression is fit over all 6 rounds to isolate the effect of heightened experience. It is likely that the coefficients would be different if we fit the model over each round separately. But as the previous analysis indicates, the losses in the first four sessions are due to the initial context-dependent bias toward action on the part of all dealers.

3.3 Discussion

It is noteworthy that dominated behavior arises and persists in this experiment despite subject selection and training procedures designed to facilitate the attainment of equilibrium strategies, and the considerable experience that the subjects gained over multiple sessions. Importantly, clear deviations from equilibrium strategies arise only in the opaque setting: dealers underestimate the strategic advantage of the insider, and its impact on their ability to compete with other dealers and extract useful information from the order flow. This leads to initially aggressive strategies and high losses that decline monotonically, but remain significant until subjects gain experience as dealers for the third time.

Why does dominated behavior arise only in the opaque setting? We have provided evidence that overconfidence with respect to the possibility of extracting useful information from order flow is a likely explanation for the initial bias toward action in the opaque setting. Thus the possibility of acquiring private information would seem to play an important role. Also potentially important is the impossibility of conditioning on the behavior of other dealers. The overconfidence driven excess entry (into the production of price discovery) seems directly related to the phenomenon uncovered in Camerer and

Lovallo (1999) where optimistic biases lead to excess entry in a coordinated entry game.

As Friedman (1998) stresses, experience and feedback will tend to result in subjects altering sub optimal behavior to behavior that is consistent with rational economic choice. We demonstrate here that the quality of the feedback for this purpose is much higher in the consolidated, transparent market than it is in the opaque, decentralized market.

It is noteworthy that the path of convergence to undominated strategies is similar for all dealers. Rather than experiment with dramatically different approaches in the face of initial losses (as in Osborne and Rubinstein (1998)), the gradual convergence to undominated strategies with experience is consistent with the predictions of learning models developed to explain out-of-equilibrium behavior in game-theoretic settings. For example Erev and Roth (1998) do a meta-analysis of experiments (previously conducted by different researchers) that have a unique equilibrium in mixed strategies. They find even a simple reinforcement learning model (that implicitly assumes that choices that have led to good outcomes are more likely to be repeated, and learning curves tend to be initially steep and then flatten) robustly outperforms the assumption of equilibrium behavior in explaining the results of experiments, or the paths of play of individual players. Importantly, a very simple family of learning theories robustly predicts ex ante, and describes ex post a wide range of experimental data. The application of these models to complex settings where very large strategy spaces preclude the derivation of equilibria is not direct, but it suggests that there are important regularities in the way agents move toward equilibrium behavior; this is exactly what we observe in the opaque setting.

A contribution of this paper is to show systematic departures from optimizing behavior can arise under the familiar microstructure information structure, but as a function of transparency. The extent to which mechanisms reduce the propensity for dominated strategies to emerge is an important characteristic of mechanism design related to facility of use.[12]

[12] Are there departures from expected utility maximization that persist in a market microstructure setting when subjects are experienced and have clearly settled into equilibrium patterns of behavior? An example from this experiment is dealer inventory management in a transparent setting. Dealer inventory management (forgoing

4 Conclusion

The first result in this paper is that it took dealers much longer to settle on procedurally rational strategies in an opaque market than in a transparent market. Dealers are critically important in this context, because they are uninformed and have no liquidity shock. If they were to choose to leave the market, then the market could break down.[13] It is not surprising that dealers require significant experience before adopting undominated strategies in the opaque setting since they must make conjectures about the actions of competitors. What is interesting is that all dealers confront this task by trading aggressively (when a cautious approach is viable) in attempts to extract profitable information from order flow. This bias toward action phenomenon is consistent with overconfidence but it eventually disappears as agents gradually update their beliefs.

Markets differ from one another in ways that can affect economic and allocational efficiency. This first result shows that another important difference is the experience needed -within the given market - to determine procedurally rational strategies or at least strategies that do not result in avoidable large average losses. In centralized, transparent markets dealers can observe one another's behavior, and also observe all trades and the prices at which these trades take place. Each dealer can readily infer the profitability of the dealer industry. In opaque markets, a dealer can only observe the trades that she is directly involved in. It is not possible to observe industry profitability, nor is it possible to update beliefs by observing a trade between a trader and another dealer. The private observation of bilateral bids and offers means that dealers have private information in the opaque market, whereas such information is

---

expected return to lock in profitable positions) occurs in a variety of experimental microstructure settings despite significantly different experimental designs (Lamoureux and Schnitzlein (1997), Schnitzlein (1996), Bloomfield (1996)). Rabin (2000) and Rabin and Thaler (2001) make the case that since expected utility theory implies that people are approximately risk-neutral when stakes are moderate, its inability to explain behavior in the face of risk when stakes are small or moderate, and its wrong predictions about the relationship between risk aversion over moderate stakes and large stakes requires a better explanation of risk averse behavior. They propose loss aversion and decision isolation as the most promising candidates that emerge from the large literature on alternatives to expected utility theory.

[13] In any case, the informed trader would be the monopolistic provider of liquidity. This would in turn imply that information about the asset's fundamental value is worth less than otherwise, as the informed trader would find it profitable to extract as many rents as possible by simply providing liquidity at both sides of the market. Furthermore, liquidity traders would likely seek other markets where the provision of liquidity is more competitive.

public in the centralized market. This explains why there is no evidence of overconfidence in most of the microstructure experimental literature, which has tended to focus on transparent dealer markets.

The second result of this paper is that behavior which could be viewed as overconfidence in an opaque asset market changes with heightened experience and feedback. It is not an equilibrium phenomenon in the usual sense. As agents became increasingly experienced and expert in the market, dealer losses diminished, and finally vanished. Also, the phenomenon depends on the market structure. Confronted with exactly the same adverse selection problem (and identical random draws), dealers in a transparent market did not lose money. The critical differences between the two mechanisms are the possibility of private information and the inability to observe the actions of competitors in the fragmented market, which are natural implications of the reduced transparency.

The third result of this paper is that, to the extent we can abstract from the learning going on in the opaque markets, fragmentation results in a diminished ability to extract information from the order flow. Dealers have to compensate by setting wider spreads. This is consistent with previous experimental research (Bloomfield and O'Hara (1999).

In light of the groundswell of interest in behavioral models in finance, financial economists should seek additional insights into human behavior in asset markets. The assumption that people behave overconfidently will generalize to financial markets in a mechanism-dependent way. We see that with adequate feedback and experience the behavior goes away. In a setting where all dealers observe the same phenomena, the behavior is never observed. In actual markets, participants choose what information to collect, how to interpret it, and how to act on it. Our findings suggest that policies that promote common knowledge with respect to the information sets and actions of others will make departures from undominated equilibrium strategies less likely.

References

Ball, S.B., M.H. Bazerman, and J.S. Carroll, 1991, An evaluation of learning in the bilateral winner's curse, *Organizational Behavior and Human Decision Processes* 48, 1-22.

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

Bloomfield, Robert and Maureen O'Hara, 1999, Market transparency: Who wins and who loses? *The Review of Financial Studies 12, 5-35.*

Camerer, Colin and Dan Lovallo, 1999, Overconfidence and excess entry: An experimental approach, *American Economic Review* 89, 306-318.

Erev, Ido and Alvin E. Roth, 1998, Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria, *American Economic Review* 88, 848-881.

Flood, Mark D., Ronald Huisman, Kees G. Koedijk, and Ronald Mahieu, 1999, Quote disclosure and price discovery in multiple-dealer financial markets, *The Review of Financial Studies* 12, 37-59.

Friedman, Daniel, 1998, Monty Hall's Three Doors: Construction and Deconstruction of a Choice Anomaly, *American Economic Review* 88, 933-946.

Genesove, David, 1995, Search at wholesale auto auctions, *The Quarterly Journal of Economics* 110, 23-49.

Glosten, Lawrence R. and Paul R. Milgrom, 1985, Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71-100.

Kagel, John H., 1995, Auctions, In *The Handbook of Experimental Economics*, John H. Kagel and Roth, Alvin E., eds., Princeton, New Jersey, Princeton University Press, 501-585.

Kyle, Albert, 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315-1335.

Lamoureux, Christopher G. and Charles R. Schnitzlein, 1997, When it's not the only game in town: The effect of bilateral search on the quality of a dealer market, *The Journal of Finance* 52, 683-712.

Lind, B., and C. R. Plott, 1991, The winner's curse: Experiments with buyers and sellers, *American Economic Review* 81, 335-46.

Lyons, Richard, 1996, Optimal transparency in a dealer market with an application to foreign exchange, *Journal of Financial Intermediation* 5, 225-254.

Madhavan, Ananth, 1995, Consolidation, fragmentation, and the disclosure of trading information, *The Review of Financial Studies* 8, 579-603.

Osborne, Martin J. and Ariel Rubinstein, 1998, Games with procedurally rational players, *American Economic Review* 88, 834-847.

Pagano, Marco and Ailsa Röell, 1996, Transparency and liquidity: A comparison of auction and dealer

markets with informed trading, *Journal of Finance* 51, 579-611.

Pouget, Sébastien, 2000, The Walrasian tâtonnement to economize on cognitive transactions costs: An experiment, Unpublished Working Paper, Toulouse University.

Rabin, Matthew, 2000, Risk aversion and expected-utility theory: A calibration theorem, *Econometrica* 68, 1281-1292.

Rabin, Matthew and Richard H. Thaler, 2001, Anomalies: Risk aversion, *Journal of Economic Perspectives* 15, 219-232.

Schnitzlein, Charles R., 1996, Call and continuous trading mechanisms under asymmetric information: An experimental investigation, *Journal of Finance* 51,613-636.

Smith, Vernon L., 1991, Rational choice: The contrast between economics and psychology, *Journal of Political Economy* 99, 877-896.

Table I

Profitability and Trading Activity By Trader Type Under the Opaque and Transparent Mechanisms

This table reports summary statistics for each of the 27 market periods across the five experimental sessions, and an additional sixth session under the opaque mechanism. During the opaque market, each trader may attempt to initiate a trade with any other trader by sending a trade proposal. A trade proposal indicates the type of trade desired (buy or sell), the trader to whom the proposed trade is directed, and the proposed trading price. Upon receiving the trade proposal, the recipient may accept the proposal or allow it to expire. Only the two traders who participated in the transaction observe a trade completed in the opaque market. Under the transparent mechanism, trade opens with a centralized dealer market. This is followed by a bilateral search market that operates as under the opaque mechanism with the exception that dealers are excluded. A second dealer market then follows this. The monetary unit in all tables is laboratory dollars (L$ in the text). To convert to US$, multiply the laboratory dollars by 0.05. Net Liquidity Trade is the sum of the liquidity demands for the liquidity traders, each of whom makes an independent draw from a uniform density from -5 to 5, excluding 0. The asset value is the realized draw from the distribution governing asset value (Gaussian, with unconditional expectation of 100, and standard deviation of 8.66). RMSE is the root mean square error of transaction prices relative to intrinsic value.

			Opaque Mechanism							**Transparent Mechanism**							
Market Period	Net Liquidity Trade	Asset Value	Dealer Industry Profits	Liquidity Traders' Profits	Insider Profits	RMSE (Value) L$	Insider Trades	Dealer Trades	Dealer to Dealer Trades	Dealer Industry Profits	Liquidity Traders' Profits	Insider Profits	RMSE (dealer portion)	RMSE (frag. portion)	Insider Trades (dealer portion)	Dealer Trades	Insider Trades (frag. portion)
1	-7	81	-164.5	55.0	109.5	11.66	10	18	1	-11.9	-35.3	47.7	6.59	17.21	2	11	2
2	1	108	13.8	-15.8	2.0	4.99	4	13	1	11.3	-68.7	57.4	4.32	7.13	11	20	6
3	3	98	-12.5	-0.5	13.0	4.02	4	9	0	-13.5	-26.5	40.0	9.02	2.80	5	8	2
4	3	105	-14.0	7.0	7.0	3.81	3	14	1	42.0	-66.6	24.6	6.30	5.12	2	11	4
5	1	84	-50.0	-27.0	77.0	12.77	8	9	1	-47.0	-48.0	95.0	10.82	15.51	6	12	3
6	5	95	-9.0	-35.0	44.0	5.13	10	11	0	3.0	-30.8	27.8	3.42	5.19	1	3	5
7	11	111	-58.0	39.0	19.0	6.43	6	19	1	74.0	-92.0	18.0	9.71	4.37	4	19	0
8	9	87	-67.0	-129.0	196.0	11.25	18	11	3	31.5	-214.5	183.0	9.04	13.83	5	14	12
9	6	95	-13.5	-58.0	71.5	5.35	13	6	0	16.2	-45.8	29.6	4.72	5.14	2	12	6
10	-3	120	-94.5	-89.0	183.5	16.55	13	14	1	-140.3	-10.6	150.9	9.88	13.27	12	17	10
11	5	93	-9.5	-54.5	64.0	5.65	12	7	3	20.1	-37.9	17.8	5.56	4.16	3	9	4
12	0	107	-4.0	-5.0	9.0	2.04	7	5	1	-4.7	-1.3	6.0	3.28	2.41	1	5	2
13	12	99	-38.0	-37.0	75.0	4.76	22	12	0	40.0	-49.6	9.6	7.30	0.93	3	8	7
14	-1	113	-29.0	-19.0	48.0	7.81	10	9	1	-22.9	-61.2	84.1	9.45	8.00	11	21	1
15	6	97	-3.0	-14.0	17.0	2.28	8	4	0	7.1	-19.1	12.0	2.22	N/A	7	17	0
16	9	98	-8.0	-25.5	33.5	2.33	14	5	2	-21.5	-81.2	102.7	7.71	N/A	15	30	0
17	-4	106	-11.0	-14.0	25.0	5.03	5	7	0	22.8	-53.5	30.7	6.43	4.51	12	27	3
18	-1	98	-29.0	20.0	9.0	3.90	4	13	0	0.5	-9.4	8.9	1.74	0.47	6	13	2
19	7	119	-144.0	51.0	93.0	12.07	9	16	2	-57.2	19.9	37.3	11.64	2.83	11	25	1
20	-4	101	4.0	-17.0	13.0	2.93	7	7	1	-5.0	-3.2	8.2	1.00	0.96	2	7	7
21	-1	100	-2.0	-8.0	10.0	2.10	8	7	0	17.0	-16.2	-0.8	4.94	1.12	1	4	2
22	-8	117	-2.0	-136.5	138.5	13.05	11	11	0	-37.2	-115.3	152.5	10.27	12.70	4	8	10
23	10	110	-48.0	11.0	37.0	5.99	10	8	1	114.5	-146.0	31.5	12.16	3.11	8	21	3
24	5	98	1.0	-26.0	25.0	3.57	8	3	0	22.9	-50.9	28.0	4.50	5.65	1	8	6
25	-9	97	3.0	-20.0	17.0	2.42	11	2	0	17.1	-19.1	2.0	3.31	2.48	0	8	5
26	-2	104	-9.0	4.0	5.0	6.42	3	3	0	10.0	-13.1	3.2	3.52	2.37	1	5	2
27	-1	92	9.0	-51.0	42.0	4.50	9	6	2	-0.4	-10.9	11.3	5.10	8.50	2	7	0
Mean	1.93	101.22	-29.21	-22.03	51.24	6.25	9.15	9.22	0.81	3.27	-48.40	45.15	6.44	5.55	5.11	12.96	3.89
Std Dev (Std Err)	5.14	9.84	(8.35)	(8.54)	(10.02)	(0.75)	4.35	4.49	0.90	(8.54)	(9.36)	(9.51)	(0.60)	(0.93)	4.22	7.22	3.20
28	6	95	3.0	-27.0	24.0	4.57	6	2	0								
29	-3	120	69.0	-104.0	35.0	10.57	10	17	2								
30	5	93	-2.0	-18.0	20.0	5.51	6	1	0								
31	0	107	-12.0	3.0	9.0	4.11	4	10	1								
32	12	99	23.0	-129.5	106.5	11.35	15	5	0								
33	-1	109	-15.5	-9.0	24.5	8.30	4	5	0								
34	0	99	0.0	-4.0	4.0	0.94	5	0	0								
35	6	111	14.0	-45.0	31.0	7.27	6	5	0								
Mean	3.13	104.13	9.94	-41.69	31.75	6.58	7.00	5.63	0.38								

Std.Dev (Std Err)	4.65	9.53	(4.86)	(8.84)	(5.75)	(0.63)	3.50	5.19	0.70

Table II

Dealer Profits and Learning Under the Opaque Mechanism and the Transparent Mechanism

In order to examine the relationship between market experience and dealer industry profitability we estimate the following model of dealer profits with ordinary least squares:

$$PROFITS = b_0 + b_1 VALUE + b_2 CORRELATION + b_3 EXPERIENCE + \varepsilon.$$

Variable definitions are as follows. DEALER PROFITS is the per period dealer industry profit, VALUE is the absolute value of the difference between the asset value and its unconditional expectation, CORRELATION is the product of the signed difference between the asset value and its unconditional expectation and the net liquidity shock, and EXPERIENCE is a dummy variable that takes on the value of 0 for the first 20 market periods under the opaque mechanism and the first 15 market periods under the transparent mechanism, and 1 for periods 21 - 35 under the opaque mechanism and 15 - 27 under the transparent mechanism.

Mechanism	$\hat{b}_0$	$\hat{b}_1$	$\hat{b}_2$	$\hat{b}_3$	Periods	Adj R^2
Opaque						
Dealer Profits	-11.39	-3.04	-0.32	32.70	35	0.53
	(-1.17)	(-3.57)	(-3.52)	(3.17)		
Insider Profits	16.54	4.78	-0.31	-16.81	35	0.44
	(1.38)	(4.55)	(-2.78)	(-1.32)		
Transparent						
Dealer Profits	37.19	-3.89	0.26	-8.23	27	0.25
	(2.30)	(-2.98)	(2.08)	(-0.52)		
Insider Profits	4.29	5.59	-0.49	-8.22	27	0.72
	(0.47)	(6.55)	(-5.86)	(-0.70)		

Table III

Dealer Trading Activity and Learning Under the Opaque Mechanism

In order to examine the relationship between market experience and dealer trading activity, we estimate the following model with ordinary least squares:

$$NUMBER\ DEALER\ TRADES = b_0 + b_1 VALUE + b_2 CORRELATION + b_3 EXPERIENCE + \varepsilon.$$

Variable definitions are as follows. NUMBER DEALER TRADES is the per period aggregate number of dealer industry trades, VALUE is the absolute value of the difference between the asset value and its unconditional expectation, CORRELATION is the product of the signed difference between the asset value and its unconditional expectation and the net liquidity shock, and EXPERIENCE is a dummy variable that takes on the value of 0 for the first 20 market periods, and 1 for periods 21 - 35.

Mechanism	$\hat{b}_0$	$\hat{b}_1$	$\hat{b}_2$	$\hat{b}_3$	Periods	Adj R^2
Opaque	6.80	0.44	0.02	-4.10	35	0.53
	(6.14)	(4.52)	(1.59)	(-3.48)		

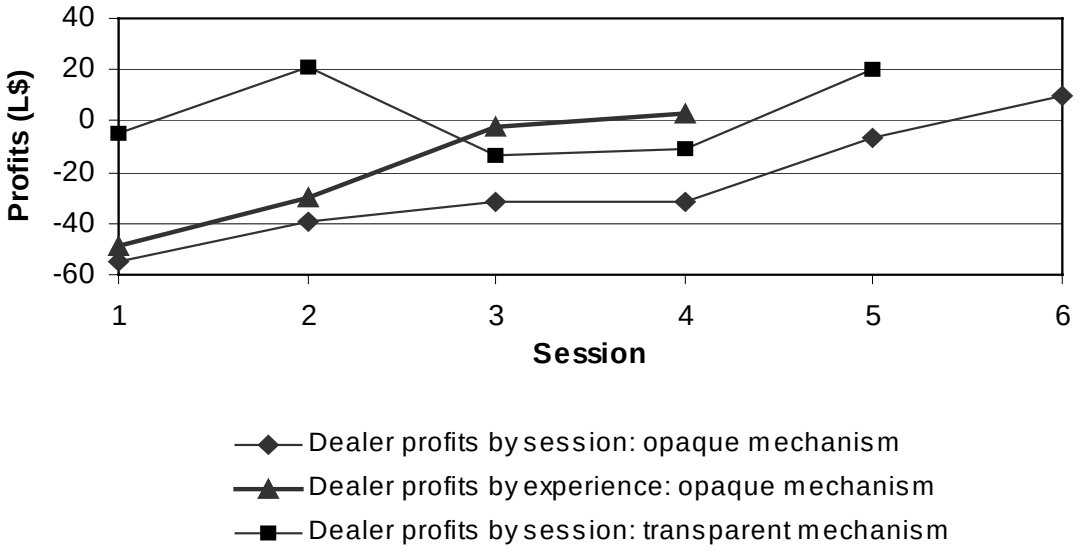

Figure 1. Dealer Industry Profitability by Session and Dealer Experience. This figure gives average per period dealer industry profits for each of the five sessions under the transparent mechanism and each of the six sessions under the opaque mechanism. Under the transparent mechanism, dealer industry profits in each session (or any combination of sessions) are insignificantly different from zero. Under the opaque mechanism, combined dealer profits over the first four sessions are significantly negative ($p<0.01$). Over the last two sessions combined dealer profits are positive and insignificantly different from zero ($p=0.93$). Dealer profits by experience shows profitability for the dealer industry as a function of dealer experience. Here the data point when session equals one gives average dealer profit per session multiplied by three (in order to make the figure comparable with dealer industry profits by session) when dealers are assuming the role for the first time, and so on. Before session six we told all subjects that the dealer industry had losses in the five previous sessions. We therefore include in the data point where session equals four all dealers that participated in session 6. This includes one dealer with only two sessions of experience as a dealer.

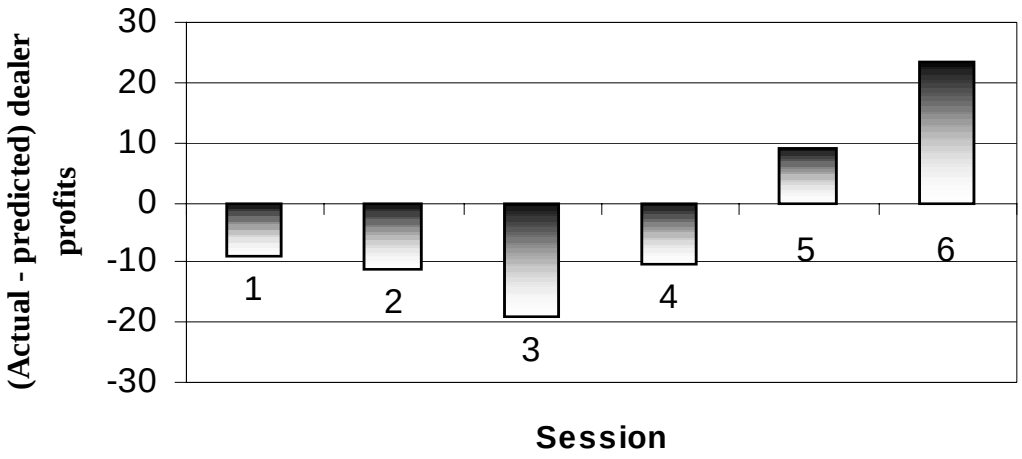

Figure 2. Actual minus predicted dealer profits. This figure reports the difference between predicted per period dealer industry profits (by session) with actual profits. We control for the random draws by estimating the following model of dealer profits: DEALER PROFITS = b_0 + b_1VALUE + b_2CORRELATION + ε. VALUE is the absolute difference between the asset value draw in the period and its unconditional expectation. CORRELATION is the product of the difference between the asset value draw in the period and its unconditional expectation and the (signed) sum of the liquidity shocks.