



ELSEVIER

Journal of Financial Markets 7 (2004) 117–143

Journal of
FINANCIAL
MARKETS

www.elsevier.com/locate/econbase

Microstructure with multiple assets: an experimental investigation into direct and indirect dealer competition[☆]

Christopher G. Lamoureux^{a,*}, Charles R. Schnitzlein^b

^a *Department of Finance, The University of Arizona, Eller College of Business Administration,
Tucson, 85721, USA*

^b *Department of Finance, College of Business, UCF, Orlando, 32816, USA*

Abstract

This paper uses the economic laboratory to isolate the effects of direct and indirect competition on dealer profitability. We compare these two settings: (1) three competing dealers in a single asset (direct competition) with (2) three assets with a monopoly dealer in each (indirect competition). We find that: bid–ask spreads are wider, prices are less responsive to order flow (so there is less price discovery), and per-trade dealer profits are larger in the single-asset setting. Important economic differences between these two settings include a heightened adverse selection problem in the three-asset setting and a public good nature of price discovery in the one-asset setting.

© 2003 Elsevier B.V. All rights reserved.

JEL: C90; D40; G10

Keywords: Market microstructure; Dealer competition

1. Introduction

A large body of research in market microstructure studies how alternative forms of market organization that differ on dimensions such as transparency,

[☆] We would like to thank Bruno Biais, Michael Brennan, an anonymous referee, and participants at the 2002 American Finance Association meetings for helpful comments. Any errors are our responsibility.

*Corresponding author. Tel.: +1-520-621-7488.

E-mail addresses: lamoureu@lamfin.bpa.arizona.edu (C.G. Lamoureux),
chuck.schnitzlein@bus.ucf.edu (C.R. Schnitzlein).

fragmentation, or temporal consolidation affect the time series properties of asset prices. A characteristic of most of this literature is that the effect of market structure on dealer competition per se is not modeled. Rather, the focus is on how the interactions of asymmetrically informed agents determine the evolution of the conditional expectation of the asset price. By assumption, dealer competition results in informationally efficient pricing and zero dealer profits. Furthermore, in this literature the implications of market structure for market performance are analyzed in a setting with a single asset, with the parameters affecting the liquidity trading process exogenous to prices in other perhaps closely related substitute assets. As we discuss below, evidence—from both empirical field studies and experimental markets—that dealers in financial markets earn monopoly rents is mounting. In this paper we attempt to learn about the industrial organization of dealership to help shed light on the sources of these profits.

We ask whether direct or indirect competition is a more powerful force for reducing bid–ask spreads using an economic experiment. Experimental work is ideally suited to address this question because the control that it affords in isolating this narrow question far exceeds that from a “natural experiment.” So, while we could look at the consequences of a regulatorily-induced change in market structure on dealer profitability, so many aspects of the environment (including entry and exit of potential competitors) may change that it could be difficult to know what is truly exogenous. Theoretical analysis of this area is also hindered by the problem of where to start. Microstructure models highlight asymmetric information and the assumed existence of uninformed traders to avoid a no-trade result. Given current modeling capabilities, if we want to allow for the extent of dealer competition to arise endogenously, we have to restrict other aspects of the model. Experimental work can help to identify what might be important in that modeling process.

We examine dealer competition directly by comparing the following asset market settings in the lab: (1) three competing dealers in a single asset (direct competition) with (2) three assets with a monopoly dealer in each (indirect competition). A key distinction between these two settings is that when three dealers are competing in a single asset, the overall liquidity demand *in the asset* is fixed or exogenous—nothing the dealers do will change the volume of liquidity trade. Another implication of this is that in the one-asset setting, the adverse selection problem can be controlled as a dealer who sets a wider bid–ask spread than her competitors has no risk of losing money since she will not be involved in any trades. By contrast, in the three-asset setting, each dealer can hope to capture a larger share of the liquidity trade in her asset by tightening her spread—giving rise to a more tangible adverse selection problem than in the one-asset case. Transparency is constant across the two settings.¹ Finally, in the three-asset setting, where each dealer is a monopolist in her asset, this dealer can internalize all the gains from price discovery. By contrast, when

¹This is important because Lamoureux and Schnitzlein (2001) show that market transparency has a large effect on dealers' ability to play non-dominated strategies. So by keeping the transparency the same across the two settings we avoid possibly confounding the equilibrium effects with behavioral effects (such as the differences in learning that are inherent with different transparency regimes).

three dealers are competing in a single asset, all three update their beliefs about the asset's value when any one trades. This public good nature of price discovery may reduce incentives to compete.

We find that in the absence of indirect competition, direct competition is not as useful in encouraging a competitive outcome as is indirect competition in the absence of direct competition. This result is manifest in the experiment in the following ways: Spreads are wider in the single-asset case than in the three-asset case. *Per asset* (or per-trade) dealer profits are larger in the single-asset case than in the three-asset case. Transaction prices are more efficient in the three-asset case than in the single-asset case. Price pressure—measured by the impact of a trade on the inside spread—is smaller in the single asset case than in the three-asset case. Per-asset insider profits are the same in the two settings, so that the allocative effects of the heightened competition in the three-asset setting entail a transfer from dealers to liquidity traders.

In both settings dealers recognize that they need not behave as price takers. Research in industrial organization recognizes the importance of inventory constraints,² technology,³ and number of plays and feedback⁴ to the equilibrium construct. But this offers little guidance to explain how dealers in an asymmetric information environment settle on a non-competitive outcome. This paper helps to provide such guidance. Our experimental results indicate *monopoly dealers*, or specialists, in similar securities have more incentives to compete than multiple dealers within a single security. Alternatively, our results suggest that having multiple dealers in a single security may reduce incentives to compete. They also may shed some light on why trading costs are lower on the New York Stock Exchange, with monopoly specialists than on NASDAQ with multiple dealers in a single security.

Empirical evidence on strategic dealer behavior and profitability includes Christie and Schultz (1994a, b), and Barclay et al. (1999). Weston (2000) provides evidence consistent with the notion that dealers earned monopoly rents prior to recent regulatory changes and that the Limit Order Display Rule has intensified competition among NASDAQ dealers. Weston (2000) also empirically documents that the number of dealers declined after recent regulatory changes. Also, Dutta and Madhavan (1997), Kandel and Marx (1997) and Bessembinder (1999) show that institutional details such as minimum tick sizes and preferencing arrangements may inhibit competition.⁵

²In these markets any single dealer could absorb the liquidity demands of the entire market, so there are no capacity constraints. The point is important in analyzing equilibria in oligopolistic markets (Tirole, 1989, p. 211). As noted by Glosten and Milgrom (1985, p. 73), binding inventory constraints would preclude Bertrand competition.

³In our markets, dealers are endowed with a franchise which requires no additional investment. In general there are constant returns to scale, and there is no entry or exit of dealers.

⁴In the experimental markets, there is no pre-commitment; competing dealers have no costs to revising posted quotes, and may do so continuously.

⁵Grossman and Miller (1987) even argue that in light of the externalities provided by dealers, artificial impediments to competition (such as restricted entry and minimum ticks) may be necessary to prevent spreads being driven to zero, which would result in market failure.

In the meantime, microstructure experiments have identified that dealers make large profits in cases where all trade is routed to competing dealers. Friedman (1993) and Theissen (2000) find that the profits of competing dealers are large.⁶ Cason (2000) looks at non-market communication in a dealership market and finds that overt collusion is not a necessary condition for a non-competitive outcome. Lamoureux and Schnitzlein (1997) find that competing dealers earn large profits, but that when traders are allowed to trade directly with one another (in a bilateral search market), dealer profits vanish, dealer bid–ask spreads are tighter, and prices more efficient than in the benchmark setting. Schnitzlein (1996) compares call markets on the basis of temporal consolidation. He finds that dealer profits are significantly higher in the continuous setting, where dealers have many more mutual interactions and opportunities to observe each others' behavior than in the one-shot setting.

Bloomfield and O'Hara (1999, 2000) and Flood et al. (1999) use the economic laboratory to compare mechanisms that differ on degree of transparency. Bloomfield and O'Hara (2000, pp. 443–444) have a market where both transparent and opaque dealers are present. In this setting, they find that “Low-transparency dealers ... act as information 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 prices more efficiently.”

Bloomfield and O'Hara (1999) compare markets with different levels of transparency. They find that bid–ask spreads are wider and prices are less informationally efficient in more transparent markets, “because of market makers' reduced incentive to compete for order flow” (Bloomfield and O'Hara, 1999, p. 28). Flood et al. (1999) find that dealers earn larger profits in a transparent setting than in an opaque one. However, they also find that dealers respond more aggressively to order flow, and hence prices are more efficient, in the opaque setting.⁷

Theoretical papers that relate to our investigation include Glosten (1989) and Leach and Madhavan (1992, 1993) where specialists who have a monopoly franchise in dealership in a particular asset may dynamically optimize, while competing dealers do not have this luxury. In these papers, this does not result in increased specialist monopoly profits, since a (dynamic) competitive equilibrium is assumed. While such assumptions are made for tractability, in most markets dealers face direct competition from other dealers and even the limit order book in the asset. They also face indirect competition from markets in other assets. Papers that model dealership in a multiple asset setting include: Hagerty (1991), Bhushan (1991), and

⁶Krahnén and Weber (2001) is a similar experiment that finds in contrast that competing market makers lose money. Subject selection and training procedures may explain the different results. Friedman, for example, uses only experienced subjects while Krahnén and Weber use subjects with no prior exposure to financial market laboratory experiments. Growing experimental evidence indicates the important role that experience plays in the convergence to equilibrium outcomes.

⁷Gemmill (1996) describes a “natural experiment” wherein he looks at the effect of delayed publication of a trade on the London Stock Exchange—comparing three timing regimes. The exchange delayed the publication of prices—but not volume—of large trades. He finds no evidence that delaying publication of price data affects market quality or liquidity.

Caballé and Krishnan (1994). All identify a diversification motive that may give a dealer in one asset some monopoly power.

The remainder of this paper is organized as follows. The experimental design is detailed in Section 2. Results are presented in Section 3, and Section 4 concludes the paper. Appendix A looks at replicated experimental sessions to evaluate the robustness of the results.

2. Experimental design, subjects, and procedures

We address a question that is the converse of the theoretical papers cited in the preceding section. Specifically, we do not take it as given that three competing dealers achieve a competitive Bertrand equilibrium. Instead, we are interested in isolating the effect that the nature of competition (direct or indirect) has on dealers' incentives to compete. We would like to consider a setting where the three types of traders are not restricted arbitrarily in terms of their strategies—which is often not possible in theoretical models. In our markets, liquidity traders and informed traders are free to trade whenever they wish. Many of the surveyed models' results hinge on artificial restrictions on the strategy sets of informed traders or liquidity traders.

Trading in these experimental markets is based on the Kyle (1985) model information structure extended to admit multiple assets and a quote-driven protocol. In both the single asset and multiple asset settings there are three types of agents: a single insider in each risky asset who learns the end-of-period asset value prior to the commencement of trade, strategic liquidity traders that receive exogenous liquidity shocks, but attempt to minimize trading costs, and three competing dealers. In both the single-asset and three-asset settings, liquidity traders are free to choose the timing of trades. In both settings, liquidity trader demands/supplies are made perfectly inelastic at the required positions by large end-of-period penalties that are proportional to the difference between the required position and the ending inventory.

Trading is conducted via a transparent quote-driven protocol: the three dealers are required to maintain standing bids and asks at all times and at any time an insider or a liquidity trader is free to hit the inside bid or ask. All agents observe all trades, although trader identities are not disclosed.

In the single-asset case the dealers' quotes are displayed in a centralized limit-order book. A dealer is free to change a bid or ask at any time, and time priority is maintained. In the three asset setting, each dealer is a monopolist market maker in a single asset. Each of the three dealers has a single bid and ask displayed in a separate limit order book. Insiders and liquidity traders are free to trade at any time in any of the three assets by hitting the bid or ask for that asset.

Since we want to isolate the effects of direct and indirect competition, we attempt to maintain as much direct comparability between the single- and multiple-asset cases as possible. In some cases this requires changing parameters from the benchmark single-asset case. For example, in the three-asset case we maintain the same expected per asset liquidity volume by increasing the scale of the distribution

from which liquidity shocks are drawn. We explain below in detail these and other features of the experimental design.

2.1. Trading mechanisms and protocols

Single asset setting. Before the trading period begins, the monopolist insider learns the end-of-period liquidation value of the corresponding asset and each dealer (D1, D2, D3) must submit a single bid and ask.⁸ A dealer's bid represents the price at which he/she is willing to buy a single unit of the risky asset; a dealer's ask represents the price at which he/she is willing to sell a single unit of the risky asset. After all dealers have entered their initial quotes, the markets open and the trading interval clock begins a 120-s countdown. Each dealer does not observe the other dealers' quotes until the market opens, however after the market opens a dealer is free to increase or decrease a bid or ask at any time. Each dealer is required to maintain a single bid and ask at all times. Over the course of the trading period the insider, the liquidity traders, and the dealers are free to trade at the "inside" bid or ask at any time. Since the trading period clock is paused whenever dealers are in the process of submitting new quotes (either due to a transaction or a quote revision) the actual trading interval requires considerably more than 120 s: most took between 8 and 12 min. No restrictions are imposed on the number of trades in which any agent may participate. Selling and buying are symmetric. There are also no restrictions on (or costs to) quote revisions, and no costs or inducements to trading activity other than the pursuit of trading profits, and the incentive to meet required positions. A mini-session consists of a number of periods where the roles of the participants are fixed. A period is defined by an independent set of random draws (on the asset value and liquidity shocks). A session consists of two mini-sessions, and is conducted in an uninterrupted block of time.⁹

Three asset setting. In general, trade proceeds as in the single asset case, with differences noted below. Before the trading period begins, each of the three monopolist insiders learns the liquidation value of one of the assets and each dealer (DAsset1, DAsset2, DAsset3) must submit a single bid and ask in his/her respective asset. Again, a dealer is free to change a bid or ask at any time, but each dealer is required to maintain an outstanding bid or ask in his/her respective asset at all times. In order to accommodate larger liquidity shocks that equalize average liquidity trading per asset relative to the single asset setting, the length of the trading interval is increased to 300 s: including pauses most trading periods took between 10 and 15 min. Each agent is notified of a transaction in any of the three assets as in the single asset case, and the trading period clock is paused whenever there is a transaction or quote revision in any of the three assets.

⁸ Initial bids and asks in the stock must bracket the expected value of L\$100.

⁹ In the single-asset treatment, we ran 3 sessions, and the number of periods within each mini-session are: 5, 5, 5, 6, 6, and 6. We ran 4 sessions under the three-asset treatment. The number of periods within each mini-session are: 5, 3, 4, 4, 4, 4, and 4.

2.2. Communication and computer displays

All interactions among subjects are conducted on a series of networked personal computers with custom software. The computer screen in front of each subject contains continuously updated market information, including trade history, cash balance, asset positions, net market order imbalance (buyer initiated trades less seller initiated trades), current bids and asks, and the time remaining in the current trading period. The insider's trading screen indicates the end-of-period liquidation value of the asset. Trades are displayed on a "ticker," with trades scrolling across the bottom of the screen. Transactions in which an agent participated are marked with an asterisk on his/her screen. The trade history shows the price of each trade that has occurred, and whether it was buyer or seller initiated.

Single asset setting. There is a single limit order book with the dealers' three bids and asks displayed to all participants (one bid and ask per dealer). The quotes are displayed on each agent's computer screen so that bids are in ascending order and asks are in descending order, with the inside quotes highlighted. A dealer can move to the "inside" on either side of the spread by improving on the current inside quote. When a trade is initiated all agents observe the occurrence of the transaction and the price via an information box that pops up on all screens, but they do not learn which agent initiated the trade, or the identity of the dealer who took the other side. The information box stays on the screen during the market pause while the dealer that took the other side of the trade resubmits a bid or ask. After the dealer completes his or her book the trading period resumes.

Three asset setting. There are three limit order books with a single bid and ask for each asset. The records of all trades are maintained and displayed on three tickers, one for each asset. Each asset has a single insider and each insider's trading screen indicates the end-of-period liquidation value of this specific asset. When a trade occurs, the information box indicates both the price and the asset in which the transaction occurred. Again, all bids and asks and the entire trade history for the three assets are displayed to all participants.

2.3. Trading profits, parameter values, and variable distributions

Endowments, payments, and monetary units. The monetary unit employed in the experiment is the laboratory dollar (L\$). To convert laboratory dollars to US\$ multiply by 0.04. Agents begin the first set of trading periods with a cash balance that differs across the type of agent. The insider has a starting cash balance of L\$330, the dealers of L\$370, and the liquidity traders of L\$500. The differences in starting cash balances are intended to minimize differences in profits by trader type. The starting cash balances average L\$420 (\$16.80) in the single-asset setting and L\$400 (\$16.00) in the three-asset setting.¹⁰ Trading profits (losses) are carried forward to the subsequent periods of that mini-session. At the end of the last period of the first mini-session there is a short break, subjects are reassigned to roles, and re-endowed

¹⁰This difference is due to the larger number of insiders in the three-asset setting.

with the above cash balances to start the second mini-session. At the end of the final market period of the second mini-session, each subject is privately paid his earnings from the two mini-sessions. Given the zero-sum nature of the trading game, average cash payments (net of any penalties incurred by traders) are determined by starting cash balances. These payments amounted to roughly \$33 per subject, per 2-h session.

Risky asset values. At the beginning of each market period, each asset value is known to be drawn from an approximate normal distribution with mean of L\$100, standard deviation of L\$8.7, and support on the whole laboratory dollars between L\$70 and L\$130, both inclusive. In the three-asset case the asset values are independent. A riskless asset (cash) is included that pays an interest rate of zero.

Liquidity traders' required positions. Each of the three liquidity traders is required to buy or sell a randomly determined number of shares of risky assets. In the single-asset setting, each liquidity trader's required position is determined by an independent draw from a discrete uniform distribution with support on the whole integers between -4 and 4 (endpoints included). In the three-asset setting, each liquidity trader's required position is a draw from a discrete uniform distribution with support on the whole integers between -13 and 13 (endpoints included). In both settings, if the liquidity trader does not exactly meet the end-of-period required position, then a penalty of L\$100 times the absolute value of the deviation between the required position and the actual end-of-period position is assessed. The magnitude of the penalty ensures that demand is inelastic at the required position, however liquidity traders can attempt to minimize trading costs by choosing the timing of trades. Each period, total trading profits and losses of each trader are added to starting cash balances and carried forward to the next period.

Trading profits. Trades represent a zero-sum game, with the profitability of each trade determined on the basis of the relation between the trading price and the end-of-period asset value. Trading profits on a trade are equal to the signed trade ($+1$ for a buy, -1 for a sell) times the end-of-period asset value less the transaction price. All transactions are for a single unit of the risky asset.

Information sets summary. Prior to the start of trade, only the insider in an asset learns the liquidation value of that stock. Each of the three liquidity traders privately observes his/her liquidity shock. All agents observe all transactions but not the identities of the trade participants. At the end of each market period, all agents learn the liquidation value of the stock(s). All information pertaining to distributions, parameters, and the rules governing trade is common knowledge.

2.4. Subjects and procedures

Fourteen of the 17 subjects who participated in the experimental markets were recruited from an MBA level investment banking class. This ensured a common background of at least one intensive one-semester course in introductory statistics, and a familiarity with trading institutions. The remaining three subjects had at least comparable training in finance and statistics. The single-asset setting required seven subjects and the three-asset setting required nine subjects. Since our intention is to study equilibrium behavior we kept a single cohort together for multiple sessions in

order to verify the effect of experience on market behavior. In the single-asset setting the same seven subjects participated in 33 trading periods over three days. In the three-asset setting, nine subjects participated in 32 trading periods over four days. After the first day, one subject was unable to schedule additional sessions at a time when the lab was available, so he was replaced with a new subject for the next 3 days when sessions were held.

In each setting we set the number of trading periods to the number that could be comfortably completed in approximately two hours. At the beginning of each session, subjects received written instructions that were read aloud by the experimenter.¹¹ After clarifying questions were answered, subjects were given instructions to familiarize them with the operation of the computer and the display of information on the computer screens.¹² They were then seated at computer terminals in individual side rooms that made it impossible to engage in any type of communication other than that afforded by market activity. Subjects were then given a short break and then assigned to new roles according to a scheme designed to give all subjects experience as a dealer as soon as possible. After this was achieved, role assignments were random. Subjects then participated in an additional five or six trading periods. In the three-asset setting a similar procedure was followed, but with a reduced number of trading periods possible during a 2-h session, the number of trading periods was set to four both before and after the break. At the completion of trading, each subject completed a questionnaire, was privately paid his/her earnings in cash, and was dismissed. Subjects indicated in the questionnaire that they had fully understood the instructions, felt adequately compensated in cash for the effort expended in participating in the experiment, and understood how their decisions affected their cash payments.

2.5. *Experimental design issues*

The increase in the magnitude of the liquidity shocks in the three-asset setting equalizes the average level of liquidity trading per asset relative to the single-asset case. We note that there is a bias toward reduced competition in the three-asset case as follows. In the single-asset case, a dealer must improve on the current inside quote to move inside—and become eligible to take the next order. In the three-asset case, *ceteris paribus*, dealers with identical quotes might be equally attractive to liquidity traders. This bias may lead to less aggressive dealer behavior relative to the single-asset case, as equalizing spreads could be a coordinating device. The liquidity traders' diversification motive (as in Hagerty, 1991), would also introduce a bias toward less aggressive dealer competition in the three-asset setting.

We rotate agents through roles to provide the deep understanding of the game that comes from having to grapple with strategic considerations from the perspective of

¹¹ A copy of the instructions is available from the authors upon request.

¹² As subjects became more experienced the length of time necessary to give the instructions decreased, allowing for additional trading periods in a two-hour period in the single asset setting. In the three-asset setting, the longer duration of individual trading periods did not allow for such an increase.

each agent type. This introduces a bias against collusive dealer behavior since the potential for developing a reputation is reduced. The larger number of sessions coupled with the larger cohort in the three-asset setting approximately equalizes the average level of dealer experience across treatments. The smaller number of periods per session in the three-asset setting may introduce a bias against a collusive outcome relative to the single-asset setting, although the larger number of transactions per session may be a countervailing effect.

Since our interest is in observing sophisticated market behavior, we recruited students that had performed well in classroom auction games. While this introduces a bias relative to the performance of a random sample of students, it does not introduce a bias across treatments. Many economic experiments control for cohort effects (the possible influence a specific group may have on outcomes), and experience effects (the influence that experience has on subject behavior), by running multiple groups of subjects through multiple sessions. Since we are interested in learning about the determinants of dealer behavior in field markets, we are concerned that the design be simple enough so that subjects can master the setting and use sophisticated strategies, yet be rich enough to afford external validity. Toward this end we recruited very selectively, as described above (in Section 2.4). We believe that the benefits of this design—subjects who fully understood and were properly motivated—outweigh the costs—a limited subject pool and a single cohort under each setting. In order to mitigate a possible cohort effect, we ran each cohort through a minimum of six different configurations of subject role assignments (recall that there are three distinct roles in these markets). Ball et al. (1991) show that in a bilateral bargaining game with asymmetric information, rotating experimental subjects through roles speeds convergence to equilibrium.¹³

3. Results

We report parameter values and summary statistics in Table 1. Panel A contains results for the 33 trading periods in the single-asset setting. Panel B contains the results from the 32 trading periods in the three-asset setting. Each trading period is distinguished by a distinct set of random draws (from the distributions detailed above) on the asset value(s) and liquidity shocks.

The three columns in Table 1, Panel B labeled Liquidity Trader Profits add the profits across the three liquidity traders from their trades in each asset in that period. Looking at the “Number of Trades” column in Panel B, we see that the averages of the number of trades in each asset in the three-asset case are almost identical. Also, when we compare this to Panel A, we see that the average number of trades in the

¹³ In the experimental literature on bubbles and crashes, cohort effects only manifest with inexperienced subjects. These experiments achieve a stable equilibrium when experienced subjects are used, independent of cohort. See Sunder (1995, esp. pp. 477–479). James and Isaac (2000, p. 999) also note this: “The most important stylized fact that emerges from the previous 15 years of research in experimental asset markets is that repeated, shared trading experience ... promotes convergence toward intrinsic value pricing.”

Table 1

Parameter values and summary statistics by trading period. We report below summary statistics by trading periods. The asset value is the realized draw from the distribution governing asset value (approximately Gaussian, with unconditional expectation of 100, and standard deviation of 8.7, and support on the whole integers). Each of the three liquidity trader's required positions is determined each period by an independent draw from a discrete uniform distribution with support on the whole integers between -4 and 4 (endpoints included) in the single-asset setting and -13 and 13 in the three-asset setting. Liquidity trader buy and sell volume is the sum of required buy and/or sell orders. RMSE is the root mean square pricing error relative to intrinsic value. All monetary amounts (value, profits and RMSE) are reported in laboratory dollars. To convert to US\$ multiply by 0.04

Panel A. Single-asset trading periods

Trading period	Asset value	Liquidity trader buy volume	Liquidity trader sell volume	Number of trades	Insider profits	Dealer profits	Liquidity trader profits	RMSE
1	112	3	3	13	26.00	-21.00	-5.00	5.08
2	84	1	3	12	97.00	-107.00	10.00	13.69
3	98	3	8	16	10.00	111.00	-121.00	10.91
4	108	4	2	6	0.00	16.00	-16.00	4.24
5	99	2	4	6	0.00	42.00	-42.00	7.16
6	114	2	4	17	47.00	7.00	-54.00	8.86
7	97	3	3	8	9.00	13.00	-22.00	6.11
8	99	1	4	5	0.00	20.00	-20.00	5.18
9	104	4	3	8	4.00	16.00	-20.00	6.04
10	100	1	7	12	14.50	44.00	-58.50	7.86
11	96	4	4	13	9.01	14.16	-23.17	4.91
12	112	5	3	22	122.35	-85.204	-37.15	10.41
13	105	4	1	9	5.67	-6.22	0.55	4.36
14	101	11	0	17	17.50	94.55	-112.05	9.25
15	96	3	3	6	0.00	25.95	-25.95	5.05
16	88	5	2	19	77.50	-12.00	-65.50	10.39
17	96	0	8	8	0.00	88.00	-88.00	12.90
18	108	1	5	11	28.10	75.40	-103.50	13.42
19	90	3	3	11	5.60	16.90	-22.50	7.58
20	102	4	5	9	0.00	76.0	-76.00	9.07
20	83	1	2	8	35.00	-27.50	-7.50	9.36
22	89	2	5	7	0.00	28.15	-28.15	8.86
23	105	5	2	7	0.00	55.25	-55.25	14.25
24	112	0	8	20	73.75	116.00	-189.75	17.97

Table 1 (continued)

<i>Panel A. Single-asset trading periods</i>								
Trading period	Asset value	Liquidity trader buy volume	Liquidity trader sell volume	Number of trades	Insider profits	Dealer profits	Liquidity trader profits	RMSE
25	91	4	7	14	23.50	51.00	-74.50	13.34
26	103	7	0	7	0.00	23.85	-23.85	9.01
27	100	6	2	9	0.00	132.50	-132.50	16.63
28	111	1	4	5	0.00	143.50	-143.50	31.55
29	89	2	5	7	0.00	36.5	-36.5	9.17
30	97	4	2	6	0.00	59.65	-59.65	11.25
31	102	5	3	8	0.00	30.90	-30.90	4.86
32	111	0	10	28	142.5	41.00	-183.50	13.06
33	109	2	1	6	3.00	0.50	-3.50	2.70
Mean	99.9	3.12	3.44	10.91	22.76	33.93	-56.69	9.83
Std. dev.	8.88	2.18	2.97	0.95	(6.55)	(9.47)	(9.04)	(0.94)
(Std. err.)								

Panel B. Three-asset trading periods

Period	Asset values			Liq. trader buy volume	Liq. trader sell volume	Number of trades			Insider profits			Dealer profits			Liquidity trader profits			RMSE		
	I	II	III			I	II	III	I	II	III	I	II	III	I	II	III	I	II	III
1	114	112	106	19	8	15	18	13	21.0	38.0	6.0	-15.5	-24.0	8.0	-5.5	-14.0	-14.0	6.31	4.36	3.14
2	105	89	101	0	20	18	12	14	41.0	34.0	33.0	20.0	-32.0	5.0	-61.0	-2.0	-38.0	5.96	4.62	5.70
3	94	104	102	5	7	5	10	8	4.0	5.0	9.0	1.0	17.0	3.0	-5.0	-22.0	-12.0	2.05	3.45	4.57
4	104	91	87	0	23	18	19	16	51.0	32.0	37.0	28.0	6.0	-20.0	-79.0	-38.0	-17.0	7.45	4.24	5.70
5	102	99	115	9	8	11	3	7	0.0	2.0	15.0	65.0	10.0	14.0	-65.0	-12.0	-29.0	7.33	6.98	15.11
6	87	117	99	0	14	12	6	13	24.0	19.5	4.0	-33.0	-10.5	53.0	9.0	-9.0	-57.0	3.25	5.16	5.24
7	102	88	102	18	0	15	14	19	22.0	51.5	34.0	18.0	12.0	18.0	-40.0	-63.5	-52.0	4.56	9.71	5.38
8	113	90	94	13	6	11	15	16	51.0	59.0	26.0	-32.0	1.0	33.0	-19.0	-60.0	-59.0	8.60	9.93	7.36
9	104	100	101	5	9	2	7	5	0.0	0.0	0.0	6.0	19.0	18.0	-6.0	-19.0	-18.0	5.00	4.29	3.90
10	96	91	85	2	15	6	4	12	-8.0	0.0	6.0	24.0	-6.0	-27.0	-16.0	6.0	21.0	4.04	2.24	5.81
11	102	115	90	6	19	12	15	20	15.0	121.0	0.0	18.0	15.0	26.0	-33.0	-136.0	-26.0	7.34	18.45	4.48

one-asset case is also close to the three-asset case. The fact that the average number of trades is lower in the one-asset case reflects insider trading activity, since the average number of liquidity trades in the one-asset case is 6.56, and the per-asset average number of liquidity trades in the three-asset case is 6.21.¹⁴

3.1. *Liquidity trader choices*

Our goal is to compare the equilibrium in the one-asset case, where dealers have only direct competition, to that in the three-asset case, where there is only indirect competition. The strategy space for liquidity traders is larger in the latter than in the former. In our single asset markets, liquidity traders choose the timing of their trades. In the three-asset setting, they also choose which asset to trade. Since we have argued that a key aspect of the indirect competition in the three-asset setting is that each dealer may hope to capture a disproportionate share of liquidity trades, the factors that influence liquidity traders' asset selection are important. If, for example, liquidity traders were largely insensitive to differences in bid–ask spreads—seeking diversification at almost any cost—then this motive would vanish. In this section, we use multinomial logit to identify the significant determinants of this choice. The results of fitting this model are shown in [Table 2](#). The statistical model is that the liquidity trader has decided to trade, evaluates the current set of spreads, the history of trades, and her own inventory and required positions, and then chooses which of the three assets to trade.

To this end, we regress the asset choice (1, 2, or 3) on the liquidity trader's absolute value of inventory in each of the three assets, the standing bid–ask spread in each of the three assets, and the price drift interacted with the direction of trade, in each of the three assets. The coefficients per se, in the logit regression are not of direct interest. Instead, we are interested in the marginal effects of the exogenous variables on the probability that the liquidity trader chooses either asset 1, 2 or 3. There are 619 liquidity trades in all periods. This is the sample size of the multinomial logit regressions. With three choices and nine explanatory variables, estimation uses up 18 degrees of freedom.

The multinomial logit regressions are statistically able to explain asset choice in the experiment. The overall model is statistically significant at the 0.0001% level, using a likelihood ratio test. The spread effect is universal and very significant. A liquidity trader is more likely to trade in an asset, the smaller the spread on that asset, and the wider the spread on the other two assets. All nine of these effects are statistically significant at the 0.1% confidence level. Spreads are measured in Lab Dollars in these regressions, so *ceteris paribus*, a one dollar decrease in an asset's spread makes it 6.5–8.4% more likely that the trader would choose to trade in that asset.

The drift/direction-of-trade effect is also statistically significant and symmetric. If an asset has drifted above (below) the unconditional expectation by L\$1, it is

¹⁴The experimental design does not imply that there would be more informed trading in the three-asset case. The amount of informed trading in each setting is an experimental outcome.

Table 2
Marginal effects on asset choice by liquidity traders

Variable	Mean	On Asset 1		On Asset 2		On Asset 3	
		Coefficient	<i>t</i> -value	Coefficient	<i>t</i> -value	Coefficient	<i>t</i> -value
Constant	1.0	−0.13	−1.99	0.06	1.00	0.07	0.88
<i>Inv</i> _{<i>t</i>,1}	1.02	−0.045	−2.08	0.026	1.34	0.019	0.79
<i>Inv</i> _{<i>t</i>,2}	0.98	0.061	2.73	−0.066	−3.28	0.005	0.22
<i>Inv</i> _{<i>t</i>,3}	1.44	0.012	0.80	−0.006	−0.42	−0.006	−0.37
<i>Sprd</i> _{<i>t</i>,1}	9.41	−0.065	−8.55	0.024	4.14	0.041	6.30
<i>Sprd</i> _{<i>t</i>,2}	10.38	0.037	6.32	−0.072	−10.82	0.035	5.76
<i>Sprd</i> _{<i>t</i>,3}	8.62	0.039	5.52	0.045	6.78	−0.084	−8.72
$\Delta p_{t,1} \cdot I_t$	1.50	−0.026	−5.41	0.018	4.22	0.008	1.71
$\Delta p_{t,2} \cdot I_t$	0.97	0.010	2.08	−0.028	−6.12	0.018	3.44
$\Delta p_{t,3} \cdot I_t$	2.12	0.010	2.12	0.026	5.63	−0.036	−6.42

Notes: This table shows marginal effects and asymptotic *t*-values from the following multinomial logit regression:

$$y_{t,j} = \beta_0 + \sum_{j=2}^3 \beta_j \cdot |Inv_{t,j}| + \sum_{j=2}^3 \theta_j \cdot Sprd_{t,j} + \sum_{j=2}^3 \gamma_j \cdot \Delta p_{t,j} \cdot I_t + \varepsilon_t,$$

where $j = 1, 2, 3$ is an index for each asset. Asset 1 is the weighting variable; $y_{t,j}$ is marginal probability that the liquidity trader trading at time t , (conditional on information available at this time), trades in asset j ; $|Inv_{t,j}|$ is the absolute value of inventory the trader who makes the t th trade in asset j ; $Sprd_{t,j}$ is the spread in asset j at the time of the t th trade, $\Delta p_{t,j}$ is the difference between the mid-point of the spread in asset j and 100, at the time of the t th trade; I_t is an indicator variable equal to 1 if the trade is a buy, -1 if a sell. The marginal effect (the probability that the liquidity trader, trading at time t will trade in asset j) is

$$P_{t,j} = \frac{e^{\Psi' X_{t,j}}}{\sum_{k=1}^3 e^{\Psi' X_{t,k}}},$$

where Ψ is the vector of individual coefficients (β, θ, γ), and X is the vector of individual regressors ($|Inv_{t,j}|, Sprd_{t,j}, \Delta p_{t,j}$). This is estimated using all 619 trades initiated by a liquidity trader.

2.6–3.6% more likely that the trader will sell (buy) that asset. Five of the six cross effects are statistically significant at the 1% confidence level. All six of these coefficients are positive meaning that *ceteris paribus*, if the prices of assets two and three have moved up, a buyer (seller) is more (less) likely to trade in asset one.

The inventory effects are less pronounced. The own coefficient is negative in all three cases—statistically significant for assets 1 and 2. Here, a trader is 6.6% less likely to buy (sell) asset 2, for each additional unit of asset 2 that she is long (short).

In general, this analysis of the choice of asset on the part of liquidity traders is consistent with the notion that liquidity volume is endogenous in the sense described above. Liquidity traders seek to trade where the spread is lowest, above other considerations. The drift effect is consistent with liquidity traders beliefs that dealers set spreads asymmetrically around their conditional expectations. They are therefore able to trade closer to the conditional expectation if they alleviate the dealer's inventory risk. Conversely, if the price is drifting up, a liquidity trader who has to

buy wants to avoid having to pay a risk premium to a dealer with a long position in that stock.¹⁵

3.2. Bid–ask spreads

The first comparison between the two settings is on average bid–ask spreads. In the one-asset case, there are 33 periods, and the average across all quotes is 17.80 (standard deviation of period means: 7.5). The analogous mean (standard deviation) from the three-asset case (32 periods) is 9.84 (3.2). When we restrict attention to asset-specific inside spreads only recorded when a trade takes place, the means (standard deviations) in the one- and three-asset cases, respectively, are: 15.65 (7.6) and 8.05 (3.2). Statistical comparison of means indicates a significantly lower bid–ask spread in the three-asset case than in the one-asset case. It should be recalled that in the three-asset case every quote is an inside spread in that an order will hit that quote (this is the heightened adverse selection problem discussed above).

3.3. Price pressure

Price pressure is another aspect of markets that is related to—and may allow measurement of—dealer competition. Table 3 Panel A examines the response of the inside spread to a trade. We see that this response is a statistically significant quote revision in the direction of the trade. The statistically significant coefficient β_1 means that there is 75% less price pressure in the one-asset case than in the three-asset case. As is common in these market experiments we observe behavior that looks like risk aversion.¹⁶ The β_4 and β_5 coefficient estimates imply that when a trade moves a dealer's inventory materially from a flat book, there is more price pressure than otherwise in the three-asset setting, but not in the one-asset setting. However, from β_2 and β_3 we infer that when a trade mitigates the dealer's inventory problem, there is not less price pressure than otherwise in either setting. This asymmetry in quote response as a function of inventory effect is probably due to the fact that in the first case, dealers want to discourage further trades on that side of the spread, while in the latter case, dealers were already trying to encourage trades on that side of the spread.

In Panel B of Table 3, we repeat the analysis of Panel A, but instead of examining the effect of a trade on the inside quote on the same side of the market, we examine the individual dealer's response to the order. For the three-asset case this is the same as in Panel A, as the dealer is always posting the inside quote. However, most of the time another dealer assumes the inside position after a trade in the one-asset market. The effect of the one-asset treatment on quote revisions is statistically insignificant at the dealer level. Similarly, the inventory effect documented above is also not

¹⁵An alternative explanation is that liquidity traders' asset choice could be affected by the price level. While we cannot rule out this behavioral phenomenon we provide additional evidence in Section 3.3, below, of dealer willingness to encourage trades that mitigate inventory imbalances.

¹⁶Rabin (2000) casts doubt on whether risk aversion per se is the source of this behavior. The reason for the behavior is irrelevant for the purposes in this paper.

Table 3
Data analysis

Panel A: price pressure—inside spread

$$\Delta q_{i,t} = \beta_0 \cdot A_{i,t} + \beta_1 \cdot A_{i,t} I_{i,t} + \beta_2 \cdot \Gamma_{i,t} + \beta_3 \cdot \Gamma_{i,t} I_{i,t} + \beta_4 \cdot \Theta_{i,t} + \beta_5 \cdot \Theta_{i,t} I_{i,t}.$$

	Coefficient					
	β_0	β_1	β_2	β_3	β_4	β_5
Estimate	1.66	-1.20	0.21	0.24	0.61	-0.67
GMM <i>t</i> -stat	18.42	-5.02	1.08	0.50	3.71	-2.04
<i>r</i> ² : 16.5%						

Panel B: price pressure—dealer response

$$\Delta q_{i,t} = \beta_0 \cdot A_{i,t} + \beta_1 \cdot A_{i,t} I_{i,t} + \beta_2 \cdot \Gamma_{i,t} + \beta_3 \cdot \Gamma_{i,t} I_{i,t} + \beta_4 \cdot \Theta_{i,t} + \beta_5 \cdot \Theta_{i,t} I_{i,t}.$$

	Coefficient					
	β_0	β_1	β_2	β_3	β_4	β_5
Estimate	1.66	0.11	0.21	1.03	0.61	0.10
GMM <i>t</i> -stat	19.29	0.26	1.10	1.26	3.63	0.17
<i>r</i> ² : 18.1%						

Panel C: mid-point spread dynamics

$$\begin{aligned} SMP_{i,t} - SMP_{i,t-1} = & \beta_0 \cdot A_{i,t-1} + \beta_1 \cdot A_{i,t-1} I_i + \beta_2 \cdot \mathcal{C}_{i,t-1} A_{i,t-1} \\ & + \beta_3 \cdot \mathcal{C}_{i,t-1} A_{i,t-1} I_i + \beta_4 \cdot \mathcal{R}_{i,t-1} A_{i,t-1} + \beta_5 \cdot \mathcal{R}_{i,t-1} I_i A_{i,t-1}. \end{aligned}$$

	Coefficient					
	β_0	β_1	β_2	β_3	β_4	β_5
Estimate	1.47	-0.87	0.45	0.91	-0.48	1.90
GMM <i>t</i> -stat	15.78	-4.69	1.64	0.51	-2.79	3.71
<i>r</i> ² : 31.4%						

Panel D: dealer profits

$$\begin{aligned} \Pi_{i,t} = & \beta_0 + \beta_1 \cdot I_{i,t} + \beta_2 \cdot |V_{i,t} - 100| + \beta_3 \cdot |V_{i,t} - 100| I_{i,t} \\ & + \beta_4 \cdot \mathcal{L}_i(V_{i,t} - 100) + \beta_5 \cdot \mathcal{L}_i(V_{i,t} - 100) I_{i,t}. \end{aligned}$$

	Coefficient					
	β_0	β_1	β_2	β_3	β_4	β_5
Estimate	27.25	46.80	-1.67	-4.26	-0.29	-0.41
GMM <i>t</i> -stat	5.82	3.46	-2.24	-2.32	-2.50	-1.47
<i>r</i> ² : 33.1%						

Table 3 (continued)

Panel E: root mean square error of transaction prices

$$RMSE_{i,t} = \beta_0 + \beta_1 \cdot I_{i,t} + \beta_2 \cdot |V_{i,t} - 100| + \beta_3 \cdot |V_{i,t} - 100|I_{i,t} \\ + \beta_4 \cdot \mathcal{L}_i(V_{i,t} - 100) + \beta_5 \cdot \mathcal{L}_i(V_{i,t} - 100)I_{i,t}.$$

	Coefficient					
	β_0	β_1	β_2	β_3	β_4	β_5
Estimate	4.01	4.41	0.35	-0.16	-0.04	-0.01
GMM <i>t</i> -stat	8.22	3.63	4.36	-0.93	-3.62	-0.52
r^2 : 43.6%						

Panel F: spread and time

$$Sprd_{i,t} = \beta_0 + \beta_1 \cdot I_{i,t} + \beta_2 \cdot N_{i,t} + \beta_3 \cdot N_{i,t}I_{i,t}.$$

	Coefficient			
	β_0	β_1	β_2	β_3
Estimate	8.83	6.73	-0.04	0.01
GMM <i>t</i> -stat	10.94	2.99	-1.47	0.05
r^2 : 23.9%				

$$Sprd_{i,t} = \beta_0 + \beta_1 \cdot I_{i,t} + \beta_2 \cdot T_{i,t} + \beta_3 \cdot N_{i,t}I_{i,t}.$$

	Coefficient			
	β_0	β_1	β_2	β_3
Estimate	9.38	6.41	-1.87	1.19
GMM <i>t</i> -stat	11.51	2.82	-2.71	0.76
r^2 : 24.7%				

$\Pi_{i,p}$ is the profit of the dealer(s) in asset i in period p . $I_{i,p}$ is an indicator variable set to 1 in the one-asset case, 0 in the three-asset case. $V_{i,p}$ is the value of asset i in period p ; (known to the informed trader at the beginning of the period). \mathcal{L}_p is the net liquidity shock in period p (divided by three in the three-asset case). $RMSE_{i,p}$ is the root-mean-square-error of all trades in asset i in period p . $\Delta q_{i,t}$ is the change in the quote on the side of the market of trade t in asset i , pursuant to trade t . $A_{i,t}$ is an indicator variable set to 1 if trade t in asset i was a buy order (dealer sold), -1 if it was a sell order (dealer bought). $\Gamma_{i,t}$ is an indicator variable set to 1 if the inventory is less than -3, and the trade was a buy (dealer sold), or -1 if the inventory is greater than 3, and the trade was a sell (dealer bought). $\Theta_{i,t}$ is an indicator variable set to 1 if the inventory in asset i is less than -3, and the trade was a sale (dealer bought), or -1 if the inventory is greater than 3, and the trade was a buy (dealer sold). $Sprd_{i,t}$ is the (inside) spread in asset i , at the time of the t th trade. $SMP_{i,t}$ is the mid-point of the (inside) bid-ask spread in asset i at the time of the t th trade (or at the end of the period). $\mathcal{C}_{i,t}$ is an indicator variable set to 1 if the spread mid-point on trade t in asset i is more than \$6 away from the unconditional expectation (of \$100) and trade t is a trade which continues the price trend. $\mathcal{R}_{i,t}$ is an indicator variable set to 1 if the spread mid-point on trade t in asset i is more than \$6 away from the unconditional expectation (of \$100) and trade t is a trade which reverses the price trend. $N_{i,t}$ is the number of trade t within the market. $T_{i,t}$ is an indicator equal to 1 if trade t is one of the first three trades in the single-asset case (9 trades in the three-asset case); 0 otherwise.

significantly affected by the one-asset treatment. Thus, the individual dealer responses to trades are similar in the two treatments. Of course, in equilibrium the dealers who are outside know that their quotes may become inside quotes pursuant to a trade, so the different competitive environments give rise to similar individual responses, but different market-level behavior (as shown in Panel A).

The previous analysis examined the instantaneous impact of a trade. In Table 3 Panel C, we consider the mid-point of the (inside) bid–ask spread at the time of each trade. This also serves to integrate the analysis at the market versus individual dealer level in Panels A and B above, by looking at the dynamics of the inside bid–ask spread when trades occur. At this level, there is no evidence that the dealers are able to identify informed trades in either setting (which is what we would expect in equilibrium). But, we saw that in the three-asset setting, liquidity traders prefer to buy an asset whose price has fallen and vice-versa. To examine the possible effects of this behavior on market price dynamics we isolate orders submitted when the spread mid-point is materially above or below the unconditional expectation. Of 360 trades in the one-asset case, 47 are sells (buys) when the spread mid-point is above 106 (below 94), 30 of these 47 orders were submitted by the insider. Of the total 360 trades there are 12 sells (buys) when the spread mid-point is below 94 (above 106).

Of the 1100 trades in the three-asset case, 194 are sells (or buys) when the spread mid-point is above 106 (below 94), 112 of these 194 orders were submitted by the insider. Of the total 1100 trades there are 55 sells (buys) when the spread mid-point is below 94 (above 106). This is a regression of the change in the mid-point of the spread on the trade direction variable, and indicator variables that isolate the effect of trades that go against the direction of the market, or that continue a price trend. That is if the price has moved up (and the mid-point of the spread is greater than \$106), and a trader sells, or vice-versa, the first indicator (\mathcal{C}) is set to 1—0 otherwise. If the price has moved up (and the mid-point of the spread is greater than \$106), and a trader buys, or vice-versa, the second indicator (\mathcal{A}) is set to 1—0 otherwise. In Table 3, Panel C, we see that the overall spread moves more in response to the order flow in the three-asset case than in the one-asset case. Orders that continue the price trend have the same effect as other orders. On the other hand, we see that orders that reverse the price direction have a significantly larger effect in the one-asset case than other orders. By contrast, in the three-asset case, such orders have a significantly smaller effect than other orders. When we interact this effect with trader identity (i.e. informed or liquidity traders), the interaction term is statistically insignificant.

These results are consistent with the notion that dealers in the three-asset case are managing inventories more than dealers in the single-asset case (this is examined more closely in Section 3.4 below). Such an effect is to be expected since the spreads are much tighter in the three-asset case than in the single asset case.

3.4. Dealer profits

Next we turn to dealer profits per asset. In the one-asset case, we aggregate the profits across the three dealers, and in the three-asset setting, there are three unique observations per period. We regress dealer profit *per asset* per period, on asset

value extremeness, and the correlation between the value shock and the liquidity shock, and treatment dummies in Table 3, Panel D. Here we see that per asset dealer profits are statistically significantly smaller when the asset value is further from its unconditional expectation (of 100), and when the liquidity traders are more likely to trade in the same direction as the informed trader. Both of these effects are expected and are consistent with the experimental literature referenced in Section 1 above. Also looking at Table 3, Panel D, we note that dealer profits are larger in the one-asset case than in the three-asset case. Dealer profits are statistically lower in the one-asset case—when the asset value is further from 100—suggesting that there are wider spreads that are less responsive to the order flow in the one-asset case than in the three-asset case.

Of course, one explanation for higher dealer profits—not due to softer competition—is increased risk, coupled with risk aversion. In Section 3.3 above, we saw some evidence that there is a non-monotone inventory effect in price pressure, which may be attributable to “risk aversion.” There was no evidence that this effect exists at the market level in the single-asset case. To investigate further whether dealers in the single-asset setting may face more risk, we examine the pattern of individual dealer trading. We would imagine, for example, that if risk aversion were important, then we would tend to see trades consistent with dealers’ attempts to maintain a flat book. To examine this, we do a simulation of the order flow in each period assigning the orders to the three dealers randomly. We construct the distribution of a test statistic of inventory imbalance within each period using 100,000 draws:

$$D_p = \sum_{t=2}^N \sum_{j=1}^3 (I_{j,t} - \bar{I}_t)^2, \quad (1)$$

where D_p is a draw from this statistic, $I_{j,t}$ is dealer j 's inventory after trade t , and \bar{I}_t is average inventory after trade t . Thus the distribution of D_p is conditioned on the actual pattern of trades in each period, p . The corresponding sample statistic \hat{D}_p is also constructed using the actual dealer inventories in the period. A small value of \hat{D}_p relative to the simulated distribution of D_p would imply that there is low inventory dispersion across the three dealers—conditioned on the order flow in that period. A high value of \hat{D}_p would be consistent with a situation where one dealer took all the buy orders and another took all the sell orders—i.e., showing no interest in controlling their inventories. If the sample realization of \hat{D}_p is consistently below the 25%ile of the (simulated) distribution of D_p , this would be an indication that the three dealers were trading to aggressively manage inventory imbalance.

In 18 of the 33 periods, the sample statistic lies within the inter-quartile range of the distribution of D , under the null hypothesis of random inventories. Of the remaining 15 periods, the statistic lies within the 90%ile range 13 times. In Period 16, the dispersion coefficient is greater than the 99%ile of the randomized statistic, and in Period 19, it lies between the 95 and 99%ile values.¹⁷ Of equal importance, the

¹⁷The quantitative results are identical whether in the simulated samples trades are randomly assigned to the 3 dealers with equal probability or the actual proportion of the trades made by that dealer in the period.

sample estimate of \hat{D}_p exceeds the mean and median of the distribution of the corresponding sample statistic in 18 of the 33 periods. Finally, the number of times the sample statistic from the period is less than the corresponding 25%ile of the (simulated) distribution (constructed under the assumption that dealers take orders randomly) is seven of the 33 periods—less than 25% of the time. Thus there is little evidence from the pattern of dealer trading that the dealers in the single-asset setting are attempting to manage their inventories in a manner consistent with risk aversion. This, coupled with the lack of inventory effects on the inside spread, suggests that dealer profits are not higher in the single-asset setting as a compensation for increased risk.

3.5. Price efficiency

We compare the price efficiency in the two settings using the regression in Table 3, Panel E. Here we use the root-mean-square-error of all transaction prices as well as spread mid-points at the time of trades within a period, for each asset, as a measure of efficiency.¹⁸

When we use transaction prices, we see that the RMSE is more than twice as high in the single-asset case, *ceteris paribus*. The RMSE is higher in an asset the further is the asset value from its unconditional expectation (of 100). The correlation between the net liquidity shock and the size of the asset shock is negatively correlated with the asset's RMSE. Neither of these effects is statistically different between the two treatments. This confirms the earlier finding that dealers in the one-asset case are less responsive to the order flow. They set wider spreads, but do not trade as much with the informed trader as in the three-asset case, and transaction prices are less informationally efficient than in the three-asset case.

3.6. Timing

In many of the periods under both treatments trading is deferred for a material portion of the period. In the one- (three-) asset case, the average time elapsed before the first trade is 78.8 (42.5) s. In the three-asset case the average spread at the open is 17.42, and at the time of the first trade, this is 14.23. The average spread across the three dealers in the one-asset case is 26.8 when the period opens, and 21.1 at the time of the first trade. The inside spread (in the one-asset case) averages 20.7 on open and 17.0 at the time of the first trade. This pattern has been documented in previous experimental asset markets—in different places and with different experimental designs (see for example Lamoureux and Schnitzlein (1997) and de Jong et al. (2001)).

In the experimental markets, all trades are for one share at a time, so a trade-size separating equilibrium as in Glosten (1989) is not possible. Some of the theoretical papers cited in the literature review suggest that monopoly specialists may attempt

¹⁸ In period 12, there were 8 trades in asset 1, 0 trades in asset 2, and 16 trades in asset 3. We exclude asset 2 in period 12 from this regression, since the RMSE is undefined.

strategies to temporally segment traders by type (as in Leach and Madhavan, 1992, 1993). We evaluate this in Panel F of Table 3. There does not appear to be a linear reduction in spreads as the period progresses. There is evidence that spreads are significantly higher in the first few trades; for example the first three trades in the one-asset case, and the first nine trades in the three-asset case. This effect is not different across the two treatments.

3.7. Experience effects

In this section we look at the effect of increased experience on dealer competition in the two settings. In the three-asset markets there is evidence that as they gain experience, dealers compete more aggressively for order flow. By contrast, there is evidence that in the one-asset case dealers become less competitive as they gain experience. To quantify this phenomenon, we regress dealer profit on an intercept, an experience dummy, and the two explanatory regressors in Table 3 Panel D—asset value and correlation. Here experience is defined as the last four sessions (periods 73–96). The estimates (GMM *t*-statistics) are respectively: 30.5 (6.07); –15.4 (–2.65); –1.64 (–2.56); and 0.31 (3.00). We see a statistically significant drop in dealer profits with experience.

We run the same regression in the one-asset case, where experience is defined as the last three sessions (periods 22 through 33). The analogous estimates (GMM *t*-statistics) are respectively: 61.3 (4.74); 37.1 (2.42); –6.0 (–4.18); and –0.60 (–2.70). Opposite the effect in the three-asset markets, dealer profits increase with experience in the one-asset markets.

3.8. Discussion of data analysis

There is no evidence that any traders were playing dominated strategies. There are important differences in dealer behavior between the one-asset and three-asset cases. These are manifest in the tighter spreads, and lower losses for liquidity traders in the three-asset case than in the one-asset case. Insider profits are actually higher in two of the three assets in the three-asset case than in the one-asset case although overall there is no statistical distinction between insider profits in the two treatments. A regression of insider profits on the same variables used in Table 3, Panel D (for dealer profits) shows that insider profits are statistically larger when the asset value is farther from 100, and statistically smaller when liquidity traders are trading in net in the same direction. However, none of the three coefficients on the treatment dummy is significant. Liquidity traders are the beneficiaries of the heightened (indirect) competition in the three asset setting relative to the (direct) competition in the one-asset setting. This result stands in contrast to Bloomfield and O'Hara (1999) where reduced transparency allows the informed trader to gain at dealers' expense.

From the behavior of the liquidity traders, perhaps the dealers in the three asset setting have an easier time of separating the liquidity trader from the informed trader. But if this were the case then we would expect to see a different dealer

reaction to an insider, especially when the spread has moved away from the unconditional expectation, in the three-asset setting. We have documented that this is not the case. If we use the resulting revision in the mid-point of the bid–ask spread, the dealers were never able to discriminate whether a trade came from a liquidity trader or the informed trader. This result is common across the two settings. Thus, while the multinomial logit results suggest that a liquidity trader is more likely to buy asset i if the price has moved up, it is not the case that the dealer in asset i acts as if she places a different conditional probability on the source of the order.

The experience effects identified in Section 3.7 suggest that the statistical analysis is conservative. If we compared only the last half of the sessions, the qualitative results would be identical, but the quantitative comparisons would be sharper.

4. Conclusions

The industrial organization of asset market makers confronted with adverse selection is a complex problem. We designed an experiment to isolate aspects of the competitive pressures on competing dealers. We confirmed independent experimental results that three dealers in a single-asset make large profits. It is not possible for the dealers in these markets to explicitly collude—they cannot communicate during the experiment, their roles are randomized over time, and this is a zero-sum game. Yet the markets do not converge to a competitive (Bertrand) equilibrium.

The experimental environment allows us to isolate how different market structures affect the degree to which competitors are able to earn monopoly rents. We are able to observe the extent of competition along with other market outcomes in a controlled, canonical microstructure environment suggested by standard models. The setting is richer than the models' because all agents behave strategically and we do not assume a form of equilibrium. Previous experiments have identified that dealer profits are affected in important ways by institutional structures, such as market transparency and the existence of “off-floor” trading. This has also occurred on NASDAQ, where recent regulatory changes have resulted in changes in the profitability of dealers. This paper is an attempt to use the economic laboratory to directly isolate the extent to which indirect competition where a monopoly specialist competes against other assets in attracting order flow is different from direct competition in affecting dealer behavior and market outcomes. We evaluate direct competition in a market with three dealers in a single asset. This is contrasted with indirect competition which we evaluate in a market with three assets (seen as a priori perfect substitutes by liquidity traders) each with a monopoly dealer.

In the three-asset setting, when we analyze the asset choice by liquidity traders, we see that they are very sensitive to the relative bid–ask spreads. In this context, indirect competition is more effective than direct competition in reducing bid–ask spreads and *dealer* monopoly rents. This contrast is sharpened as dealers gain additional experience. Liquidity traders are direct beneficiaries of this increased competition. Informed trader profit is the same in the two settings. Reasons for this

include the fact that per-asset liquidity volume is endogenous with indirect competition, but not with direct competition. Along the same lines, direct competition does not heighten the dealer's adverse selection problem—setting a spread wider than your competitors' cannot result in losses. On the other hand, there is canonical adverse selection with indirect competition. Finally, price discovery creates a public good with direct competition, but not with indirect competition. A unique dealer in her own asset can internalize all the rents from doing a better job of price discovery than her competitors. This is not true in the one-asset setting: with direct competition, even dealers not involved in a trade observe the trade and revise beliefs accordingly.

Appendix A. Experimental robustness

This paper reports the results of an economic experiment designed to identify aspects of the nature of dealer competition. The experiment itself requires motivated subjects who have experience with both the experimental software and the information environment. While there is evidence in experimental economics that *equilibrium* outcomes are robust to cohort effects (as discussed in the text in Section 2.5), readers of this experiment may choose to question the robustness of the results, given that only two groups of subjects (nine in the three-asset markets, and seven in the one-asset markets) were used in the design. In this appendix, we report the results from an additional sequence of three-asset markets that were conducted using an entirely different group of nine subjects, roughly 2 years after the original markets were conducted. We replicate the three-asset markets, because while there is a growing body of evidence suggesting that single-asset dealership markets converge to non-competitive outcomes, this is the first paper that reports results from multiple asset dealer markets. Of particular interest is whether the replicated markets show the same patterns in convergence as the original market—a distinguishing feature of our results comparing the single and three asset markets. Both the replicated data and analysis of experience effects suggest that the experimental design along with the statistical analysis of the results reported in the text are conservative. The *qualitative* conclusions are strengthened by the extended analysis in this appendix.

In October, 2002, we replicated the experimental design described in the paper for the selection and training of subjects for the three-asset markets. Instructions and payments are identical to those described above. The subject pool comprised two MBA students, six Master's in Finance students, and one Master's in Agricultural Economics student, all selected from a Master's level Investment Banking course at the University of Arizona. Subjects reported for three evenings of two four-period sessions each. These subjects were selected based on their performance in a classroom exercise before the experiment. The first of these three evenings was planned to serve as training rounds—to ensure familiarity with the roles, etc. The first three sessions replicate the draws from periods 10–12 from Table 1, Panel B, and the last four sessions replicate the draws from periods 13 through 16 from

Table 1, Panel B. Period 4 does not have a direct counterpart in the original experiment.¹⁹

Robustness checks are conducted by repeating the regressions in Table 3 using the additional 24 asset/periods (from the last two sessions) along with the original 129 asset periods used in that table. Dummy variables are interacted with all of the regressors to allow statistical comparison of possible cohort effects. All such effects are either statistically insignificant, or they provide stronger evidence on the identified contrasts between the one- and three-asset cases. For example expanding Table 3 Panel D, we add three regressors to identify possible cohort effects on the intercept and the asset-value and correlation between asset value and net liquidity shock—as these affect dealer profits. The respective coefficients along with GMM *t*-statistics are: 10.24 (0.72); 0.72 (0.31); and -0.43 (-0.93). Similarly, when we add the 24 asset periods to the root-mean-square-error regressions in Table 3, Panel D, the three interactions on possible cohort effects are statistically insignificant.

As another example, consider the price-pressure regressions from Table 3, Panel A. Here we add the 256 quote revisions (from the last two sessions) to the original 1460. We interact the possible cohort effects with the three regressors—the intercept, the inventory-mitigating, and inventory-exacerbating variables. The estimates and GMM *t*-statistics are: 1.05 (3.91); 1.30 (0.91); and 1.24 (3.40), respectively. Thus, there was even more price pressure in the direction of a trade, and this was even more sensitive to inventory effects than in the original three-asset periods, where the substantive point was that quotes are more responsive to order flow in the three-asset case than in the one-asset case.

Finally, as discussed in Section 2.5 and footnote 13 above, while equilibrium results are expected to be robust across cohorts, this is not necessarily true about the convergence to equilibrium. Thus it is especially important to evaluate the robustness of the experience effects analyzed in Section 3.7 above. To this end, we combine the data from the replicated 3-asset markets with the original 3-asset data. This provides 144 periods. We regress dealer profit on an intercept, an experience dummy, asset value and correlation—replicating the regression reported in Section 3.7. Here experience is defined as the last four sessions in the original data (periods 73–96) and the last two sessions in the replicated data (periods 13–24). The estimates (GMM *t*-statistics) are respectively: 41.3 (6.63); -12.9 (-1.97); -1.64 (-2.56); and 0.30 (2.60). Thus across the two cohorts we see a significant drop in dealer profits with experience.

References

- Ball, S.B., Bazerman, M.H., Carroll, J.S., 1991. An evaluation of learning in the bilateral winner's curse. *Organizational Behavior and Human Decision Processes* 48, 1–22.
- Barclay, M.J., Christie, W., Harris, J., Kandel, E., Schultz, P., 1999. The effects of market reform on the trading costs and depths of NASDAQ stocks. *The Journal of Finance* 54, 1–34.

¹⁹The asset values in this additional period are: 99, 95 and 99, respectively, and the liquidity shocks are: -5 , 5, and 6, respectively.

- Bessembinder, H., 1999. Trade execution costs on NASDAQ and the NYSE: a post-reform comparison. *Journal of Financial and Quantitative Analysis* 34, 387–407.
- Bhushan, R., 1991. Trading costs, liquidity, and asset holdings. *The Review of Financial Studies* 4, 343–360.
- Bloomfield, R., O'Hara, M., 1999. Market transparency: who wins and who loses? *The Review of Financial Studies* 12, 5–35.
- Bloomfield, R., O'Hara, M., 2000. Can transparent markets survive? *Journal of Financial Economics* 55, 425–459.
- Caballé, J., Krishnan, M., 1994. Imperfect competition in a multi-security market with risk neutrality. *Econometrica* 62, 695–704.
- Cason, T.N., 2000. The opportunity for conspiracy in asset markets organized with dealer intermediaries. *The Review of Financial Studies* 13, 385–416.
- Christie, W.G., Schultz, P.H., 1994a. Why do NASDAQ market makers avoid odd-eighth quotes? *The Journal of Finance* 49, 1813–1840.
- Christie, W.G., Schultz, P.H., 1994b. Why did NASDAQ market makers stop avoiding odd-eighth quotes? *The Journal of Finance* 49, 1841–1860.
- de Jong, C., Koedijk, K.G., Schnitzlein, C.R., 2001. Stock market quality in the presence of a traded option. Working paper, UCF.
- Dutta, P.K., Madhavan, A., 1997. Competition and collusion in dealer markets. *The Journal of Finance* 52, 245–276.
- Flood, M.D., Huisman, R., Koedijk, K.G., Mahieu, R., 1999. Quote disclosure and price discovery in multiple-dealer financial markets. *The Review of Financial Studies* 12, 37–59.
- Friedman, D., 1993. Privileged traders and asset market efficiency: a laboratory study. *Journal of Financial and Quantitative Analysis* 28, 515–534.
- Gemmill, G., 1996. Transparency and liquidity: a study of block trades on the London Stock Exchange under different publication rules. *The Journal of Finance* 51, 1765–1790.
- Glosten, L.R., 1989. Insider trading, liquidity, and the role of the monopolist specialist. *Journal of Business* 62, 211–235.
- Glosten, L.R., Milgrom, P.R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14, 71–100.
- Grossman, S., Miller, M., 1987. Liquidity and market structure. *The Journal of Finance* 43, 617–637.
- Hagerty, K., 1991. Equilibrium bid–ask spreads in markets with multiple assets. *Review of Economic Studies* 58, 237–257.
- James, D., Mark Isaac, R., 2000. Asset markets: how are they affected by tournament incentives for individuals? *The American Economic Review* 90, 995–1004.
- Kandel, E., Marx, L.M., 1997. NASDAQ market structure and spread patterns. *Journal of Financial Economics* 45, 61–89.
- Krahnhen, J., Weber, M., 2001. Marketmaking in the laboratory: does competition matter? *Experimental Economics* 4, 55–85.
- Kyle, A.S., 1985. Continuous auctions and insider trading. *Econometrica* 53, 1315–1335.
- Lamoureux, C.G., Schnitzlein, C.R., 1997. When its 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.
- Lamoureux, C.G., Schnitzlein, C.R., 2001. To err is human—but dealers' abilities to find non-dominated strategies depend on market transparency. Working paper, University of Arizona.
- Leach, C., Madhavan, A., 1992. Intertemporal price discovery by market makers: active versus passive learning. *The Journal of Financial Intermediation* 2, 207–235.
- Leach, C., Madhavan, A., 1993. Price experimentation and security market structure. *The Review of Financial Studies* 6, 375–404.
- Rabin, M., 2000. Risk aversion and expected-utility theory: a calibration theorem. *Econometrica* 68, 1281–1292.
- Schnitzlein, C.R., 1996. Call and continuous trading mechanisms under asymmetric information: an experimental investigation. *The Journal of Finance* 51, 613–636.

- Sunder, S., 1995. Experimental asset markets: a survey. In: Kagel, J.H., Roth, A.E. (Eds.), *Handbook of Experimental Economics*. Princeton University Press, Princeton, pp. 445–500.
- Theissen, E., 2000. Market structure, informational efficiency and liquidity: an experimental comparison of auction and dealer markets. *Journal of Financial Markets* 3, 333–364.
- Tirole, J., 1989. *The Theory of Industrial Organization*. MIT Press, Cambridge.
- Weston, J.P., 2000. Competition on the NASDAQ and the impact of recent market reforms. *The Journal of Finance* 55, 2565–2598.