Imperfect competition in financial markets: ISLAND vs NASDAQ

Bruno Biais[1] Christophe Bisière[2] Chester Spatt[3]

Revised November 15, 2002

[1]Toulouse University (IDEI-CRG-GREMAQ)
[2]Toulouse University (IDEI-CRG)
[3]Carnegie Mellon University

Many thanks to conference participants at the Banque de France, the 2002 European Finance Association Meetings, the 2002 European Econometric Society Meeting, the 6^{th} Toulouse Finance Workshop, and the 2002 Rutgers University Market Microstructure Conference in Honor of David Whitcomb. We are particularly grateful for insightful comments from Thierry Foucault, Richard Roll, Gideon Saar and Patrik Sandas. The financial support of the Fondation Banque de France is gratefully acknowledged.

Abstract

The Internet technology reduces the cost of transmitting and exchanging information. ECNs exploit this opportunity to enable investors to place quotes at very little cost and compete with incumbent stock exchanges. Does this quasi–free entry situation lead to competitive liquidity supply? We analyze trades and order book dynamics on Nasdaq and Island. The Nasdaq touch is frequently undercut by Island limit orders, using the finer tick size prevailing on that ECN. Before decimalization, the coarse tick size constrained Nasdaq spreads, and undercutting Island limit order traders earned oligopoly rents. After decimalization, the hypothesis that liquidity suppliers do not earn rents cannot be rejected.

1 Introduction

Internet technology reduces communication and information processing costs and reduces the cost of entry in financial markets for market organizers as well as investors and traders. It can be expected to enhance competition for liquidity supply. Does it? How? To what extent? Our analysis of financial market trading suggests that the Internet can enhance competition, and correspondingly, increase liquidity. It enables efficient modern market architectures to be set up. These compete with the incumbent and force them to catch up. But, perfect competition does not emerge immediately, even on an Internet-based transparent easily accessible limit order book, although Adam Smith's condition for perfect competition seems to hold. Competition between markets is necessary to make up for imperfect competition among traders.

The role of financial markets is to bring together buyers and sellers and discover the prices at which they can trade. A crucial ingredient in this process is the information about the extent to which different agents are willing to buy or sell, and at what price. When information on prices, bids and orders is available only to a small group of market professionals, they enjoy an informational advantage, and correspondingly earn rents. They are also well placed to earn rents when they have privileged ability to produce and disseminate information by posting prices and making offers. These rents are the mirror image of the trading costs borne by less informed parties. Until recently, access to information on prices and offers in financial markets was very unevenly distributed. For example, agents physically present on trading floors had faster and better access to information about the trading process than investors remote from the market center. In the words of Matthew Andresen (2000) (Chairman of Island):

> "The so called time and place advantage has allowed professionals to make very handsome living. Seats on the NYSE sell for million of dollars thanks to the informational advantage of those physically present on the trading floor."

Admittedly, M. Andersen has some vested interest in these matters, which may well be correlated with his views on the relative merits of floors and electronic markets. The academic literature, however, also offers systematic evidence on the rent–making potential of the small group of financial intermediaries with privileged access to information collection and dissemination (see e.g., Christie and Schultz (1994) or Chen and Ritter (2000)).

The advent of the internet technology brought about a dramatic reduction in the cost of transmitting and exchanging information rapidly among a large number of people. This makes it possible to design more open and transparent market mechanisms, based on widely–disseminated information, and the generalized ability to make offers and post prices. In conjunction with this technological revolution, the regulatory change brought about by the 1997 SEC order display rule, have made it possible for Electronic Communication Networks (hereafter ECNs) to offer information dissemination, price quotation and order matching mechanisms for Nasdaq securities.

This evolution raises a number of interesting issues for finance scholars, economists at large, and practitioners: What are the economics of the competition between internet-based and more traditional markets? Do the limit orders placed on ECNs compete away the rents of Nasdaq market makers, and how? Does the fact that internet based markets are widely accessible at a very low cost, lead to a free–entry/perfectly competitive market situation? The analysis adds to the debate on centralized vs. fragmented markets (as illustrated by the theoretical studies by Chowdhry and Nanda (1991) and Parlour and Seppi (2002)).

To study these issues, we focus on Island, one of the largest ECNs, which offers a particularly interesting market structure: it is operated as a fully transparent limit order market, and its book is freely observable to all on the Island website. To conduct our analysis, we combine two sources of data for actively traded Nasdaq stocks. On the one hand we have downloaded the sequence of trades and order book dynamics from the Island site, on the other hand we have used the Nastraq dataset from Nasdaq.

Another interesting feature of the competition between Nasdaq and Island is related to the pricing grid used in these two market places. Before April 2001, the pricing grid on Nasdaq was relatively coarse, as the tick size was one-sixteenth, and quite thin on Island, where the tick size was 1/256. Since April 2001, Nasdaq prices are quoted in cents, while the Island tick is one-thousandth of a dollar.

Results: The results of our empirical analysis of order placement and spreads on Island and Nasdaq can be summarized as follows.

Although the average Nasdaq quote is tighter than the average Island inside spread, limit orders placed on Island often undercut the Nasdaq best quotes, and consequently the Island quotes are often better than their Nasdaq counterparts. Yet, the extent to which Island limit order traders compete with one another is somewhat limited, reflecting in particular the fact that Island limit orders frequently undercut Nasdaq quotes by just one Island (very thin) tick.

We also find that, before decimalization, the Nasdaq spread was often constrained by the rather coarse Nasdaq tick. After decimalization, however, this constraint was significantly relaxed, resulting in much narrower Nasdaq spreads. In this context, Island limit orders, undercutting the Nasdaq spread by just one tight tick are less likely to earn rents.

To investigate further the trade–offs underlying the placement of limit orders on Island we develop a structural econometric analysis relying on a standard market microstructure theoretical model, in the line of Glosten (1994), Bernhardt and Hughson (1997), and Biais, Martimort and Rochet (2000). Our estimates of the parameters of the model lead us to reject the hypothesis that limit orders placed on

Island face no adverse selection cost. In addition, they suggest that these orders earned oligopoly rents before decimalization, but not after.

Literature: Our GMM analysis of limit-order schedules and adverse selection builds on the insightful structural econometric analysis of Sandas (2001), though we take a somewhat different viewpoint. While our approach directly stems from the theoretical model (Glosten, 1994, Biais, Martimort and Rochet, 2000) and does not require parametric assumptions, Sandas (2001) relies on some parametric assumptions (to specify the link between the size of trades and their information content) and on some simplifications of the theoretical model (e.g., by postulating an exogenous joint distribution of orders and values). Also, while Sandas (2001) assumes competitive liquidity supply, we investigate market power.

Our empirical analysis is also related to Hasbrouck and Harris (1996), in particular to their "ex–post performance measure," which computes the profitability of different types of orders, such as limit orders at the best quotes for example, when they are executed. Our focus differs from theirs, however. While, they compare the performance of order placement strategies, we focus on the competition to supply liquidity. Furthermore our econometric analysis relies on a theoretical model, while they take a more descriptive approach. Finally, we take advantage in our identification strategy of the fact that we have two different samples, over which the order–handling cost is likely to remain constant while other parameters, such as the oligopoly mark–up are likely to change.

Interesting empirical analyses of ECNs are offered by Simaan, Weaver and Whitcomb (1998), Huang (2002), Barclay, Hendershott and McCormick (2001) and Hasbrouck and Saar (2001). Our focus on the competition for liquidity supply differs from Huang (2002), who analyzes price discovery, Barclay, Hendershott and McCormick (2001), who focus upon market quality, and Hasbrouck and Saar (2001), who analyze the relationship between volatility and the order flow as well as fleeting orders. We complement Simaan, Weaver and Whitcomb (1998) in the following dimensions.

- They rely on Nastraq data, where Island quotes are rounded to match the Island grid, which precludes observing situations when Island quotes are strictly better than Nasdaq quotes by less than one Nasdaq tick. In contrast, by using unrounded Island data, we can shed light on the extent to which Island limit orders undercut Nasdaq quotes.

- We show that the competition between Nasdaq and Island has been significantly changed by decimalization.

- We develop a new structural econometric approach to analyze quantitatively the trade–offs faced by liquidity suppliers between adverse selection costs and oligopoly rents. A by–product of this analysis is to offer a new method to disentangle components of the spread without relying on parametric assumptions on distributions, while building directly from a standard microstructure model.

Structure of the paper: The next section presents the institutional environment. Section 3 presents the data. Section 4 presents our empirical analysis of the Island and Nasdaq quotes and spreads. Section 5 extends this analysis by offering a structural investigation of profits and spreads on Island. Section 6 concludes.

2 Institutional environment

2.1 ECNs

ECNs are e–brokers, relying on web–based platforms, which collect limit and market orders, and match them or display them on internet–based order books. In 2002 they have been estimated to capture 39.3% of the dollar volume of Nasdaq trading (source Market Data).[1] The largest ECN, Instinet, was estimated to represent 12% of the trading volume on Nasdaq in February 2002, while Island amounted to 9.6%, Redi Book to 6.5%, and Archipelago to 10.5%.

While ECNs compete with the traditional source of liquidity on Nasdaq (i.e., market makers), they are brokers: they do not take proprietary positions, but simply handle and display their customer's orders. Since they are regulated as brokers, they are subject to the best execution rule, which means that they cannot conduct trades away from the current best market prices. This best execution rule implies that ECNs, as all brokers, must allocate orders according to price priority. However, time priority is not enforced between ECNs and Nasdaq market makers.

Interestingly, empirical evidence so far does not suggest that the trading process managed by ECNs is systematically free riding on price discovery achieved by the traditional market participants (the Nasdaq dealers). To the contrary, Huang (2002) shows empirically that ECNs are important contributors to the price discovery process.

2.2 Island

Island is a web–based transparent limit order book.[2] Trades and the best 15 quotes on each side of the book can freely be viewed in real time through the internet. Island subscribers can freely place orders. Trades can take place from 7:00 a.m. to 8:00 p.m. Immediately executed orders are charged .25 cent per share traded. Non-marketable limit orders posted in the book receive .1 cent per share when executed, as compensation for supplying liquidity. When an order is transmitted to Island, if it is not immediately marketable, it is stored and displayed anonymously in the

[1] In 2000, this proportion was 26% (McAndrews and Stefanadis, 2000).

[2] Hasbrouck and Saar (2001) offer a good description of the market structure.

Island order book.[3] If the order is marketable it is executed at the best market price. This can be set by an Island quote, or correspond to another quote in the Nasdaq system, in which case the Island order trades with an order outside Island.

Until April 2001, the Nasdaq tick size was 1/16 and the Island tick size was 1/256. Since April 2001, the Nasdaq tick size has been reduced to one cent ($0.01), while on Island prices are quoted with a three digit precision, i.e., the tick size is one-thousandth of a dollar. The thinner price grid on Island makes it easier for traders placing orders on that ECN to undercut Nasdaq market makers' quotes.

The best Island bid and ask quotes are displayed on the Nasdaq screen, along with the best quotes of ECNs and market makers. Note however that Island quotes displayed and "advertised" on the Nasdaq screen are not shown at their actual price (quoted on a thin grid) but at rounded prices (from the Nasdaq grid). For example, when the Nasdaq tick was one-sixteenth, if the best ask for stock XYZ on Island was one dollar and 1/24, it was displayed on Nasdaq at one dollar and 1/16. Rounding the Island quotes enables NASDAQ to avoid price priority constraints, and reduces the ability of the ECN to advertise good quotes and thus attract orders. This makes it very important for Island to use another vehicle besides NASDAQ screens to disseminate information. This may be one of the reasons for the excellent and easily accessible website Island has developed.

Since the best Island bid and ask quotes are displayed on the Nasdaq screen, if the grid size was the same in the two market segments, the Nasdaq inside quote would, by construction always be at least as good as the Island spread. The grid is thinner on Island than on Nasdaq, however. This raises the possibility that Island quotes, placed on thin ticks inside the relatively coarse grid Nasdaq, could better the Nasdaq quote, at least on one side of the book.

3 Data

We downloaded a continuous record of the Island book from the Island website during 5 trading days in March 2000 (from March 8 to March 14) and 5 trading days in June 2001 (from June 18 to June 22). We collected this data for 7 stocks: COMS, Cisco, Dell, Intel, Microsoft, QCom, and Sun.

We also acquired quotes and trades data from Nasdaq (including the Dealer Quotes (DQ) file, the Inside Quote (IQ) file, and the Trades (TR) file).

We consider data starting at the opening of NASDAQ (at 9:30 or a few minutes before) and ending at the NASDAQ close: 4:00 p.m. (because of a data feed problem, for March 10, we have data only between 9:30 and 2:30).

Our Island data set includes 138432 trades for the March 2000 sample and 139966 trades for the June 2001 sample. The average trade size in the March 2000 sample is 389 shares, while its counterpart for the June 2001 sample is 450 shares. The corresponding average dollar value per trade is $ 37720 for the March 2000

[3] Traders also can place hidden orders.

sample, and $ 13290 for the June sample. The decrease in the dollar value of the average trade reflects the decline in the general share price level.

One advantage of the Island data (downloaded from the Island website) is that it is not rounded to sixteenths (unlike the ECN quotes reported in the NASDAQ DQ file). Hence we can study the use of fine ticks by Island traders.

4 The Island and Nasdaq spreads

In this section we compare liquidity supply by Island limit order traders and by other Nasdaq participants, in particular market makers. To do so we compare the best Island quotes to the best prices quoted on Nasdaq by the other liquidity suppliers. To document the consequences of decimalization, we conduct this analysis separately for the March 2000 and June 2001 samples.

4.1 Before decimalization

The inside spread on Nasdaq in March 2000 First consider the best quotes posted on Nasdaq by market makers and other ECNs than Island. The corresponding time-weighted average spread in March 2000 was 1.27 Nasdaq ticks, corresponding to $.079. The mode and the median spreads were equal to exactly one Nasdaq tick: one-sixteenth of a dollar ($.0625). In fact, the spread was exactly equal to one tick 81.5% of the time. The two other most frequent values of the spread were 2 ticks (13.25% of the time) and 3 ticks (3.36% of the time).

These observations suggest that before decimalization, the relatively coarse tick size prevailing on Nasdaq was likely to be rather constraining (as illustrated in other contexts by Harris (1991, 1994)). Potentially, this could have resulted in excessively wide Nasdaq spreads. We investigate this point further in the remainder of the paper.

The Island inside spread in March 2000 During this period, the tick size was much finer on Island. Hence, it was possible to offer liquidity much more aggressively on that market. This could be achieved by undercutting the Nasdaq quotes, using the fine grid prevailing on Island. We document this empirically.

On Island the time-weighted average spread for our 7 stocks and our sample period was equal to 50.70/256, which is approximately $ 0.198. The mode and the median were 32/256, i.e., $.125. Thus, the average Island spread was on average larger than the Nasdaq inside spread resulting from Island's competitors.

Figure 1, Panel A, depicts the frequency of different values of the spread for the seven stocks in our sample, on Island and on Nasdaq in March 2000.[4] As illustrated in the figure there is marked clustering in the data.[5] As discussed above, the most

[4]This graphical representation of the empirical frequency of different values of the spread is similar to Figure 2 in Barclay et al, 1999.

[5]The pervasiveness of clustering is documented in Harris (1991).

frequent spread size both on Island and on Nasdaq was equal to one Nasdaq tick. But this occurred much less often for the Island spread than for the Nasdaq inside spread. Interestingly, the spread on Island was quite often just one Island tick below these levels, e.g. 15/256, 31/256, etc. This is likely to reflect undercutting, by a fine increment, of the Nasdaq spread by Island liquidity providers. By following this strategy, they acquire price priority relatively cheaply. Note however that, since Island prices were rounded before being represented on the Nasdaq system, the Island liquidity suppliers benefit from this priority advantage only with respect to the Island order flow.

As illustrated in Figure 2, the average depth was 924.53 shares at the best ask quote, 952.59 at the second best ask quote, and 844.86 at the third best ask quote. The corresponding average depths on the three first best levels in the book on the bid side were 790.09, 757.59, and 692.98.

Comparing the Island and Nasdaq spreads To obtain more insights on the comparison between the Island and Nasdaq spreads we merged the Island data file and the Nasdaq IQ data file.[6]

For March 2000, for the seven stocks in our sample we obtain the following results. First compare the best Island bid quote and the best bid posted on Nasdaq by its competitors. 35.89% of the time the Island bid was strictly highest, 43.16% of the time it was lower, and 20.94% of the time the two bid quotes were equal. On the ask side, the best Island quote was better than that of its competitors 26.71% of the time, it was higher 43.43% of the time, and the two quotes were equal 29.87% of the time.

Our results are consistent with the findings by Simaan, Weaver and Whitcomb (1998), that ECNs often establish the inside market, and are less likely to quote odd–sixteenths. Our results differ from, and complement those of Simaan, Weaver and Whitcomb (1998) because we analyze data on unrounded Island quotes, downloaded from their site, rather than rounded quotes from the Nasdaq DQ file. Hence, we find more frequent occurrences of the situation where Island beats the Nasdaq market makers quotes, and we document undercutting by Island orders on a finer grid than the sixteenth grid.[7]

4.2 After decimalization

The inside spread on Nasdaq in June 2001 Again, consider the best quotes posted on Nasdaq by market makers and other ECNs than Island. In June 2001, the Nasdaq tick was one cent, and the time-weighted average spread was 1.48 ticks. The spread was exactly equal to one tick 78.08% of the time. The two other most frequent values of the spread were 2 ticks (12.26% of the time) and 3 ticks (4.23% of the time).

[6]This can raise synchronicity problems. The consequences of these non–synchronicity problems and the way we dealt with them are discussed in the Appendix.

[7]In the Appendix we offer some further discussion of the impact of rounding for the data.

Thus, decimalization led to a dramatic decrease in the average spread on Nasdaq, as well as a reduction in clustering. This is consistent with the view that the coarse tick size before decimalization constrained liquidity supply and led to artificially large spreads.

The spread The average spread on Island for our 7 stocks in June 2001 was equal to $ 0.0648. This is definitely below the corresponding figures for March 2000.

Figure 1, Panel B, presents the histogram of spread sizes on Island and Nasdaq for our June 2001 sample. As in the March 2000 case, there is a lot of clustering on the Nasdaq price grid. Both the Island spread and that of its competitors is most frequently equal to .01, .02 and .03, .04 or .05. There also is clustering for the Island spread just one tick below these values, reflecting undercutting of the Nasdaq quotes.

Comparing the distribution of Island spreads in June 2001 to its March 2000 counterpart (see Figure 1) points to the stark reduction in spreads contemporaneous to the reduction in the Nasdaq tick from 1/16 to 1/100 and the more modest reduction in the Island tick during that period from 1/256 to 1/1000.

Depth While the spread has decreased after decimalization of Nasdaq, this could be offset by a corresponding decrease in the depth at the quotes. To shed light on this point we have computed the average depth at the best quotes as well as at other levels in the Island book. Our results, presented in Figure 2, show that the depth at the best quotes on Island is not lower in the June 2001 than in March 2000.

4.3 Conclusion

Putting together the above results, the following picture of the competition between Nasdaq and Island emerges:

Before decimalization, liqudity supply on Nasdaq was constrained by the coarse tick size, resulting in large spreads. The decrease in spread brought about by decimalization led to a stark decrease in the Nasdaq spread.

Island liquidity suppliers compete for order flow by frequently undercutting the Nasdaq quotes. However, they seem to compete less aggressively within Island, and undercut each other less frequently. The tick size prevailing on Island in 2000 was already extremely thin and very unlikely to constrain liquidity suppliers. Yet, the Island spread was strongly reduced by decimalization. This reduction took place because the Island limit order traders had to react to the decrease of the spread of their Nasdaq competitors. That they were able to engage in this reduction, and yet had not done it before, is suggestive of imperfect competition among Island liquidity suppliers. The next section studies this point further.

5 An econometric analysis of the costs and profits of Island limit order traders

5.1 Theoretical framework

In this section we examine econometrically the costs incurred by Island limit order traders and the profits they earn. Consider a limit order to sell, at time t, at the best quote in the Island limit order book: $A_{1,t}$. If this order is filled its profit is:

$$A_{1,t} - v - (c - f),$$

where v is the fundamental value of the asset, c is the order–handling cost incurred by the limit order trader, and f is the compensation offered by Island to executed limit orders. Note that f is not a parameter to be estimated, but an observable pricing rule set by Island. Similarly, if the sell order was at the ith best quote in the Island book, its profit is:

$$A_{i,t} - v - (c - f).$$

Now consider a standard market microstructure model (as in Glosten, 1994): competitive risk neutral limit order traders face risk–averse investors privately informed about the underlying value of the stock and their own risk–sharing needs. In this case, the marginal limit order just breaks even on average. Denoting the expected profit π_1, this break-even condition can be written as:

$$\pi_1 = E(A_{1,t} - v - (c - f)|H_t, Q_t \geq Q_{A_{1,t}}) = 0,$$

where H_t is the information set of the liquidity suppliers just before receiving the order, Q_t is the size of the market buy order hitting $A_{1,t}$, and $Q_{A_{1,t}}$ is the depth of the order book at the best quote. As first emphasized by Rock (1990) and Glosten (1994), the conditioning set in this upper tail expectation reflects the workings of the limit order book: the marginal limit order at the best ask in the book is executed if and only if the total size of the market buy order is greater than or equal to the depth at the best ask price in the book.[8] The limit order reflects this informational content of trades.

Similarly, in this competitive case, the expected profit of the marginal limit order at the ith best price level in the book is:

$$\pi_i = E(A_{i,t} - v - (c - f)|Q_t \geq Q_{A_{i,t}}) = 0.$$

A symmetric equality holds on the bid side of the book.

[8]This differs from the information structure arising in the signaling trading game analyzed in Kyle (1985). In the latter, the transaction price is equal to the expectation of the value of the asset conditional on the exact size of the trade.

On the other hand, if the limit order traders are strategic, as in Bernhardt and Hughson (1997), and Biais, Martimort and Rochet (2000), their expected profits are not equal to zero. In that case, the relevant condition is:

$$E(A_{i,t} - v - (c + \pi_i - f)|Q_t \geq Q_{A_{i,t}}) = 0,$$

where the expected profit (π_i) is not in general equal to 0, as long as the number of liquidity suppliers, N, is finite. As shown in Biais,, Martimort and Rochet (2000), when N goes to infinity, the oligopolistic mark–up π_i goes to 0, and quotes go to their competitive level.

5.2 Econometric approach

5.2.1 Using the spread

In this subsection we show how the bid and ask equations above yield empirical restrictions which can be used to test the model and estimate its parameters. Subtracting the bid from the ask, the spread is:

$$A_{i,t} - B_{i,t} = \alpha_i + 2(c - f + \pi_i),$$

where:

$$\alpha_i = [E(v|Q_t \geq Q_{A_{i,t}}, H_t) - E(v|Q_t \leq -|Q_{B_{i,t}}|, H_t)]$$

denotes the informational component of the spread at the ith level in the book.

5.2.2 Using price changes following trades

Some time after (say at time $t + \Delta t$), the liquidity suppliers have updated their expectation of the fundamental value of the asset to form:

$$E(v|H_{t+\Delta t}).$$

This can be proxied, for example, by the mid–quote say half an hour or an hour after the trade:

$$m_{t+\Delta t} = E(v|H_{t+\Delta t}) + \epsilon_{t+\Delta t}.$$

For simplicity, we assume that $\epsilon_{t+\Delta t}$ is white noise. In this context, we obtain that:

$$m_{t+\Delta t} - A_{i,t} = [E(v|H_{t+\Delta t}) + \epsilon_{t+\Delta t}] - [E(v|Q_t \geq Q_{A_{i,t}}, H_t) + c - f + \pi_i].$$

Taking expectations conditional on the occurrence of the purchase at time t:

$$\begin{aligned} E(m_{t+\Delta t} - A_{i,t}|Q_t &\geq Q_{A_{i,t}}, H_t) = \\ E([E(v|H_{t+\Delta t}) + \epsilon_{t+\Delta t}] - [E(v|Q_t &\geq Q_{A_{i,t}}, H_t) + c - f + \pi_i]|Q_t \geq Q_{A_{i,t}}, H_t). \end{aligned}$$

Applying the law of iterated expectations:

$$E([E(v|H_{t+\Delta t}) + \epsilon_{t+\Delta t}]|Q_t \geq Q_{A_{i,t}}, H_t) = E(v|Q_t \geq Q_{A_{i,t}}, H_t).$$

Hence, the expected difference between the ask price and the subsequent mid–quote simplifies to:

$$E(A_{i,t} - m_{t+\Delta t}|Q_t \geq Q_{A_{i,t}}, H_t) = c - f + \pi_i.$$

A similar equality holds for the bid side:

$$E(m_{t+\Delta t} - B_{i,t}|Q_t \leq -|Q_{B_{i,t}}|, H_t) = c - f + \pi_i.$$

The intuition is that, on average, the informational component of the spread differences out, so that the difference between the transaction price and the subsequent midquote, i.e., the gross trading profit of the liquidity supplier, is equal to his non–informational cost (c net of the compensation offered by Island to liquidity supply, f) plus the oligopoly rent (π_i).

5.2.3 The 2 moment conditions

Building on the above analysis, we obtain two moment conditions for each level i of the book, which can be used to estimate the parameters and test the model:

$$E(A_{i,t} - B_{i,t} - [\alpha_i + 2(c - f + \pi_i)]|H_t) = 0,$$

and:

$$\begin{aligned} E([B_{i,t} - (m_{t+\Delta t} - (c - f + \pi_i))]I(Q_t &\leq -|Q_{B_{i,t}}|) \\ +[A_{i,t} - (m_{t+\Delta t} + (c - f + \pi_i))]I(Q_t &\geq Q_{A_{i,t}})|H_t) = 0. \end{aligned}$$

where $I(.)$is the indicator function equal to 1 when the condition in the argument holds and 0 otherwise.

5.2.4 Identification

The second moment condition enables one to identify $c + \pi_i$. Denote θ_i the sum of these two parameters. Equipped with the estimate of $\theta_i = c + \pi_i$ obtained from the second moment condition, we can estimate α_i from the first moment condition. Hence α_i is identifiable. Our approach to decomposing the spread is potentially robust and avoids relying upon auxiliary parametric assumptions, instead building upon a fundamental microstructure model.

Unfortunately we cannot identify separately the two components of θ_i. Yet, information about these components can be obtained in the two following manners.

First, consider the moment conditions under the hypothesis that limit order traders are competitive. In that case, π_i is equal to 0, $\forall i$. Consequently, $\theta_i = c, \forall i$. To test if this condition holds, we can test if θ_i is constant across levels in the book, i.e., $\forall i$. If we reject this restriction we can reject the hypothesis that limit orders are placed by competitive risk neutral traders.

Second, note that we observe data from two subperiods: period 1 before decimalization, and period 2 after. A priori there is no reason to expect α_i or π_i to be constant across the two periods. Indeed, a change in the oligopoly mark–up π_i is to be expected, since the rules of the game are different in the two samples, reflecting the change in the tick size. The information content of trades α_i also might well have changed. Correspondingly, we carry the estimation separately over the two periods, which yields two sets of parameter estimates: $\{\alpha_{i,1}, \theta_{i,1}\}$ and $\{\alpha_{i,2}, \theta_{i,2}\}$. On the other hand, it is plausible that the order handling cost c is constant across our two subsamples. Under this assumption, we disentangle π_i from c_i by comparing the results obtained for the two periods. Since $\theta_{i,1} = c + \pi_{i,1}$ and $\theta_{i,2} = c + \pi_{i,2}$ we have that: $\theta_{i,1} - \theta_{i,2} = \pi_{i,1} - \pi_{i,2}$. Thus, if we find that $\theta_{i,1} - \theta_{i,2} > 0$ and since $\pi_{i,2}$ must be non–negative, we know that $\pi_{i,1} > 0$.

5.3 Empirical results

5.3.1 Observations and instruments

To carry the GMM estimation, we use the following 3 instruments (in addition to the constant): an indicator variable taking the value one if the ask quote is on the Nasdaq grid, an indicator variable taking the value one if the bid quote is on the Nasdaq grid, and the current size of the spread.

As discussed above, if limit order traders are competitive and risk neutral, the marginal limit orders break even on average (in line with Glosten, 1994, and Sandas, 2001). In contrast, infra-marginal orders can earn profits, even in the competitive case. Hence, to test the hypothesis that liquidity supply is competitive, we impose the moment conditions only on trades involving marginal limit orders in the book. For example, suppose that, at the best ask quote in the book, 500 shares are offered. A market buy order for 250 shares would not hit the marginal limit order. Hence it would not be included in the data we use to estimate the model and test the competitive hypothesis. In contrast, a market buy order for 500 shares or more would be included.

5.3.2 A first, simple, specification

First we estimate a simple specification, where it is assumed that the parameters are constant across market conditions. The GMM parameter estimates are presented in Table 1.

- As can be seen in the table, the estimates of α are significantly positive in both subperiods, and at both price levels in the book. Negative estimates of

α would have contradicted the model. Significantly positive estimates lead to rejection of the hypothesis that there is no adverse selection.

- The estimate of $c + \pi_1$ is significantly positive in 2000. Hence, we reject the hypothesis that there is no order–handling cost or market power in that period. In contrast, the estimate of $c + \pi_1$ is not significantly different from 0 after decimalization. Under the plausible assumption that c did not vary between the two periods, this suggests that liquidity supply on Island was imperfectly competitive before decimalization, i.e., π_1 was significantly positive during the first period. The estimate of $c + \pi_2$, although positive in the two periods, is not significantly different from 0.

5.3.3 A more flexible specification

While the assumption that c is constant through time and market conditions is reasonable, adverse selection and rent earning opportunities are likely to vary. Consequently, we estimate a more flexible specification allowing the parameters to vary with market conditions. In this specification the adverse selection cost parameter at the ith level in the book and at time t ($\alpha_{i,t}$) is specified as:

$$\alpha_{i,t} = \beta^{\alpha}_{i,1} + \beta^{\alpha}_{i,2}\sigma_t + \beta^{\alpha}_{i,3} r_{m_t} + \beta^{\alpha}_{i,4} I_m,$$

where σ_t is the volatility of the stock during the last half–hour, r_{m_t} is the return on the index computed as the average of the prices of the 7 stocks in our sample during the last half–hour, and I_m is an indicator variable taking the value 1 in the morning and 0 in the afternoon. Similarly,

$$\theta_{i,t} = \beta^{\theta}_{i,1} + \beta^{\theta}_{i,2}\sigma_t + \beta^{\theta}_{i,3} r_{m_t} + \beta^{\theta}_{i,4} I_m.$$

Since the number of parameters to be estimated is greater than for the simpler version of the model studied above, the number of instruments must be greater also, so as to maintain a positive number of degrees of freedom. Hence, to the four instruments used above (constant, indicator variable taking the value one if the ask quote is on the Nasdaq grid, indicator variable taking the value one if the bid quote is on the Nasdaq grid, and spread), we add the number of shares exchanged in the last trade.

The values of the χ^2 statistics and the corresponding p–values for this more flexible specification are shown in Table 2. The null hypthesis that the model is correct is not rejected at the 1% level. This suggests reasonable adequacy of the model to the data.[9]

[9] This reasonable fit of our structural model with a flexible specification is not unlike that obtained in Sandas (2001).

6 Conclusion

This paper is a study in the industrial organization of liquidity supply. We examine the competition between limit order traders and market makers as well as on the competition between markets.

We find that, before decimalization, Nasdaq spreads were constrained by the tick size, and were correspondingly excessively wide. Reacting to this situation, limit order traders used Island as a platform to compete for the supply of liquidity. To do so they often undercut the Nasdaq inside quotes, by using the finer Island grid. Undercutting on Island did not lead to competitive liquidity supply however. In contrast with zero–profit free–entry equilibrium, limit orders placed on Island, before the Nasdaq decimalization, earned positive profits (net of transactions costs). After the Nasdaq decimalization, the Island spread became much tighter (without reduction in the depth at the quotes). In this context, the rents earned by Island limit orders virtually disappeared.

Our results suggest that the wide dissemination of information and the reduction in the costs of accessing markets brought about by the internet technology are important but not sufficient to eliminate market power in financial markets, in particular in the supply of liquidity. In addition to the competition between liquidity suppliers within a marketplace, competition between trading mechanisms plays an important role.

Our findings point at the competitive pressure exerted by ECNs such as Island on the Nasdaq system. In light of our results, decimalization on Nasdaq can be interpreted as reflecting (at least in part) a reaction of the market makers to the competitive pressure from Island.[10] Our empirical findings also point at the reduction in Island spreads, brought about by the reduction of the Nasdaq spread, generated by decimalization which took place on Nasdaq.

[10]This suggest the Nasdaq market organizers faced the following trade–off. On the one hand, keeping a relatively coarse grid size (one sixteenth) can maintain artificially high spreads, at which Nasdaq dealer earn rents. On the other hand, keeping such a coarse grid makes it easier for Island to compete the order flow away from Nasdaq. This is reminiscent of the classical dilemma faced by oligopolists between the benefits, in terms of profit per unit, of charging large prices, and the costs of this strategy, in terms of market share.

References

Andresen, M., 2000, Don’t CLOBer the ECNs, *Wall Street Journal*, March 27.

Barclay, M., W. Christie, J. Harris, E. Kandel, and P. Schultz, 1999, The effects of market reform on trading costs and depths of Nasdaq stocks, *Journal of Finance*, 54, 1–34.

Barclay, M., T. Hendershott and T. McCormick, 2001, Electronic communications networks and market quality, working paper, University of Rochester.

Bernhardt, D., and E. Hughson, 1997, Splitting orders, *Review of Financial Studies*, 10, 69–102.

Biais, B., P. Hillion and C. Spatt, 1995, An empirical analysis of the order book and order flow in the Paris Bourse, *Journal of Finance*, 50, 1655-1689.

Biais, B., D. Martimort and J.C. Rochet, 2000, Competing mechanisms in a common value environment, *Econometrica*, 68, 799–837.

Chen, H., and J. Ritter, 2000, The seven percent solution, *Journal of Finance*, 55, 1105-1131.

Chowdhry, B., and V. Nanda, Multimarket trading and market liquidity, *Review of Financial Studies*, 4, 483-511.

Christie, W., and P. Schultz, 1994, Why do Nasdaq market makers avoid odd–eight quotes?, *Journal of Finance*, 49, 1813–1840.

Glosten, L., 1994, Is the electronic limit order book inevitable?, *Journal of Finance*, 49, 1127–1161.

Harris, L., 1994, Minimum price variations, discrete bid-ask spreads, and quotation sizes, *Review of Financial Studies*, 7, 149-178.

Harris, L., 1991, Stock price clustering and discreteness, *Review of Financial Studies*, 4, 389-415.

Harris, L., and J. Hasbrouck, 1996, Market versus limit orders: The SuperDOT evidence on order submission strategy, *Journal of Financial and Quantitative Analysis*, 31, 213–231.

Hasbrouck, J., and G. Saar, 2002, Limit orders and volatility in a hybrid market: The Island ECN, working paper, New York University.

Huang, R., 2002, The quality of ECN and Nasdaq market maker quotes, *Journal of Finance*, 57, 1285-1319.

Kyle, P., 1985, Continuous auctions and insider trading, *Econometrica*, 53, 1315-1336.

Market Data, `http://www.marketdata.nasdaq.com.`

McAndrews, J., and C. Stefanadis, The emergence of Electronic Communications Networks in the US equity markets, *Federal Reserve Bank of New York Current Issues in Economics and Finance*, Volume 6, Number 12, pages 1 to 6.

Parlour, C., and D. Seppi, 2002, Liquidity-based competition for order flow, *Review of Financial Studies*, forthcoming.

Rock, K., 1990, The specialist’s order book and price anomalies, working paper, Harvard Business School.

Sandas, P., 2001, Adverse selection and competitive market making: Empirical evidence from a limit order market, *Review of Financial Studies*, 14, 705–734.

Simaan, Y., D. Weaver, and D. Whitcomb, 1998, The quotation behaviour of ECNs and Nasdaq market makers, working paper, Baruch College.

Table 1: GMM estimates when the parameters are assumed to be constant across market conditions

Panel A: First level in the book (inside quotes)

	$c + \pi$	α
2000	0.0184 (2.64)	0.0906 (6.45)
2001	-0.0006 (-0.31)	0.0379 (9.80)

Panel B: Second level in the book

	$c + \pi$	α
2000	0.0304 (1.75)	0.0737 (2.11)
2001	0.0072 (1.56)	0.0241 (2.59)

Table 2: p-level of the flexible model

	first price level in the book	second price level in the book
2000	2.96 %	38.14 %
2001	49.80 %	15.87 %

Appendix

The consequences of rounding To better document the impact of the rounding procedure on the quotes observed on the Nasdaq system, we conducted the following experiment. Using the Island data for March 2000 (from the NASTRAQ DQ file), we computed the mean spread on Island. It was equal to \$59.26/256, which is greater than its Island data counterpart, \$50.70/256. This shows that the rounding procedure made the Island quotes much less attractive than they were actually. This confirms our remark above that Island traders relied on other ticks than sixteenths to quote to a large extent.

Synchronicity problems Since the Island quotes are incorporated in the Nastraq IQ quotes, the former can be better than the latter only when they are on a finer price grid than the Nasdaq grid. This offers an opportunity to assess the magnitude of the problems induced by synchronicity. When the best Island bid (resp. ask) is better than the best Nasdaq bid (resp. ask), it should be on a finer tick than the Nasdaq grid. In our data this is the case for DELL 84.60% (resp. 84.63%) of the time. This suggests that in 15% of the cases synchronicity problems induce mistakes in our best quotes comparisons.

Figure 1, Panel A: Frequency of different values of the spread on Island and Nasdaq (excluding Island) in March 2000 for the 7 stocks in our sample

Note: graph zooms on 0 to 2 Nasdaq ticks. 46% of the time Island spread was greater than 32/256.

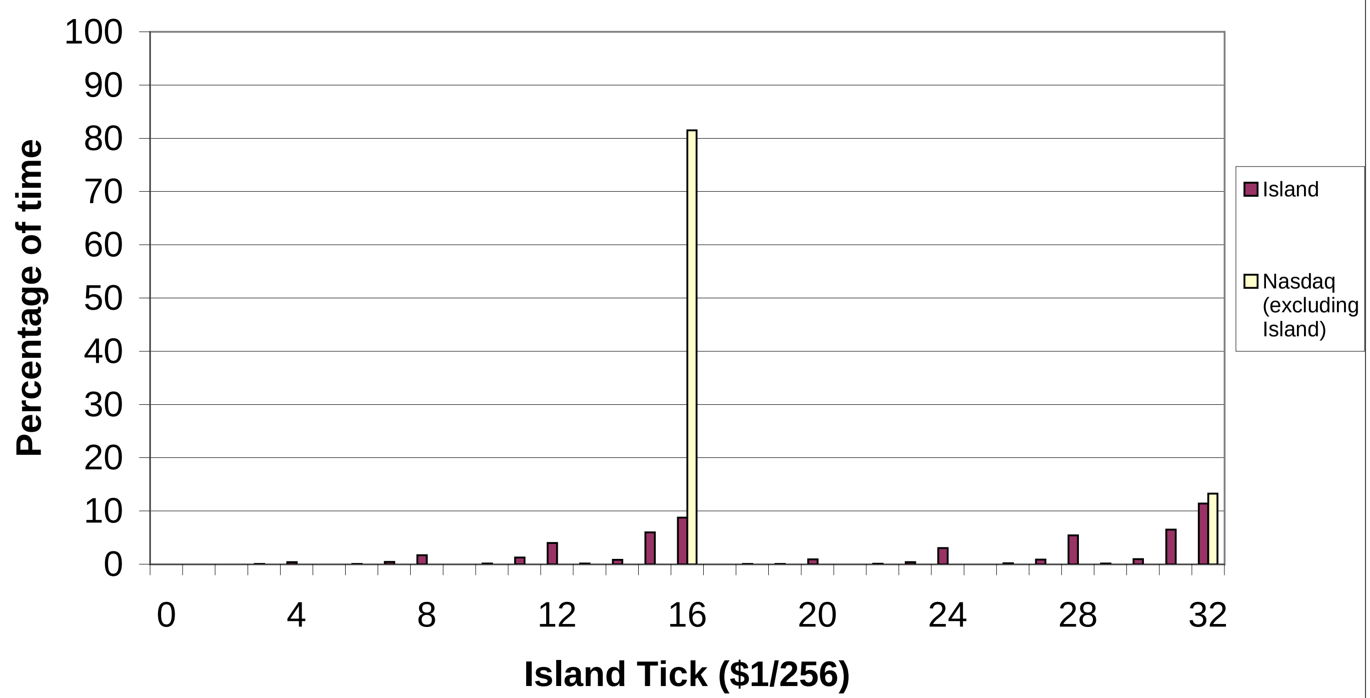

Figure 1, Panel B: Frequency of different values of the spread on Island and Nasdaq (excluding Island) in June 2001 for the 7 stocks in our sample

Note: graph zooms on 0 to 5 Nasdaq ticks. 39% of the time Island spread was greater than 50/1000.

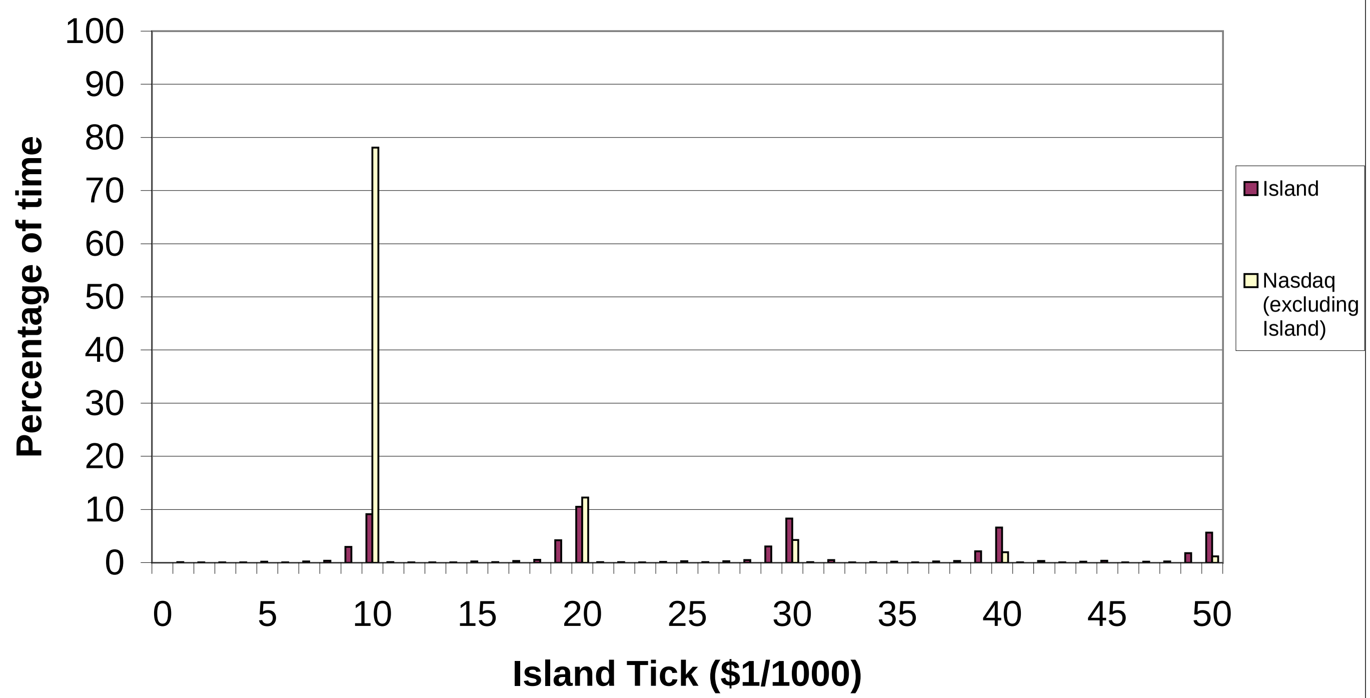

Figure 2: Average depth (in number of shares) in the Island book in March 2000 and June 2001