Imperfect competition in financial markets:

ISLAND vs NASDAQ

Bruno Biais (Toulouse University), Christophe Bisière (Perpignan University) and
Chester Spatt (Carnegie Mellon University)

February 2002

Abstract

The Internet technology reduces the cost of transmitting and exchanging information. ECNs exploit this opportunity to compete with incumbant stock exchanges. They enable all investors to place quotes, at very little cost. Does this lead to competitive liquidity supply? We analyze empirically trades and order book dynamics using the Nastraq dataset and data downloaded from the Island website. The Nasdaq touch is frequently undercut by Island limit orders, using the finer tick size prevailing on that ECN. Liquidity supply on Island is not perfectly competitive, however. On average Island limit orders earn rents.

Imperfect competition in financial markets:

ISLAND vs NASDAQ

1 Introduction

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 goup 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 not totally uncorrelated 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 small group of financial intermediaries with priviledged access to information collection and dissemination (see e.g. Christie and Schultz (1994) or Chen and Ritter (2000)).

The advent of the internet technology brings about a dramatic reduction in the cost of transmitting and exchanging information rapidly between 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 ECNs and more traditional market participants? Do the limit orders placed on ECNs compete away the rents of Nasdaq market makers, and how? Does the presence of ECNs lead to a free–entry/perfectly competitive market situation, as in Glosten (1994)?

To study these issues, we focus on Island, the second largest ECN, 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 acquired from Nasdaq the Nastraq dataset.

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. Our study of data generated by the two markets before and after the decimalization enables us to analyze the following issues:

- Did the coarse price grid prevailing on Nasdaq result in excessively large spread on this market ?
- If so, how did liquidity supply react to this situation? Did the Island market offer competitive liquidity supply, and correpondingly tight spreads? Or did the Island liquidity suppliers react strategically to that situation, undercutting the Nasdaq quotes to gain market share, but keeping large spreads to earn rents?

Thus, the setting we analyze enables us to shed light on whether liquidity suppliers in financial markets are competitive or strategic.

Based on the pre–decimalization data we obtain the following results:

- The best Island quotes were quite frequently strictly better than the Nasdaq quotes. Orders placed on Island often bettered the Nasdaq touch by using a finer price grid than the Nasdaq grid. This suggests that the coarseness of the Nasdaq grid (possibly in combination with strategic behaviour), resulted in excessively large Nasdaq inside quotes. Responding to this situation, Island limit order traders undercut the Nasdaq quotes.
- Were these undercutting strategies profitable, or did they result in a zero–profit competitive equilibrium, à la Glosten (1994)? We find that, on average

limit orders placed on Island earned positive profits. In fact, these profits were greater when the Island orders matched the Nasdaq quotes than when they undercut these quotes. These results suggest that the competition to supply liquidity on Nasdaq and Island is closer to the Cournot type of imperfective competition (analyzed by Bernhardt and Hughson (1997) and Biais, Martimort and Rochet (2000)), than to the free–entry zero–profit equilibrium analyzed by Glosten (1994).

While the results based on our March 2000 data are consistent with the conclusion that Island liquidity suppliers are non–competitive, we cannot reject the hypothesis that their average profits simply were the counterpart of their operational costs, including order–handling costs, infrastructure, personnel, clearing and settlement fees, etc... As discussed below, our analysis of the post–decimalization data enables us to disentangle the two hypotheses.

Based on our June 2001 sample, we obtain the following results: after the decimalization, spreads are much tighter; while clustering still occurs, it corresponds to much tighter levels of the spread. Furthermore, the depth at the quotes is not reduced. Under the plausible assumption that the administrative costs of market–making did not decrease substantially between our two sample periods, our results suggest that, before the decimalization, liquidity was strategically supplied on Island. Liquidity providers just slightly undercut the (excessively wide constrained) Nasdaq spreads, while leaving the Island spread sufficiently large to earn rents.

In the next section, we describe the market structure. In Section 3 we present our data. Section 4 presents results on the spread in Island and Nasdaq before the decimalization. Section 5 analyzes the profit of Island liquidity suppliers during that period. Section 6 analyzes the order book on Island after the Nasdaq decimalization. Section 7 offers a brief conclusion.

1.0.1 Competition between liquidity suppliers

While our above reported results suggest that the Nasdaq spread was excessively large in March 2000, they leave open the question whether the Island spread was competitive or not.

- On the one hand, as entry on the Island market is very cheap, and many traders are placing limit orders in that ECN, one could expect that the supply of liquidity would be competitive on that market. This would be in line with the theoretical results obtained by Biais, Martimort and Rochet (2000), when the number of liquidity suppliers goes to infinity, and with the hypothesis underlying the econometric approach of Sandas (2001).

- On the other hand, limit orders placed on Island before the decimalization of Nasdaq could have left rents to the liquidity suppliers. This could reflect collusion or coordination on non–competitive order placement strategies, undercutting the Nasdaq quotes enough to attract order flow, but not so much that it would drive Island liquidity suppliers profits to zero.

While the clustering of Island quotes just next to the Nasdaq quotes is suggestive that the second term of the alternative is plausible, the results reported above could be consistent with competitive liquidity supply on Island. In the next sections we offer empirical analyses which speak to these issues.

Our results also highlights a trade–off faced by the Nasdaq market organizers. 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.

Note also that rounding the Island quotes enabled NASDAQ to get around price priority constraints, and reduced the ability of the ECN to advertise good quotes and thus attract orders. The latter made it very important for Island to use another vehicle than NASDAQ screens to disseminate information. Hence its excellent website.

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 2000 they have been estimated to capture 26% of the dollar volume of Nasdaq trading (Mc Andrews and Stefanadis, 2000). The major ECN, Instinet (INCA), was estimated to represent 14% of the trading volume on Nasdaq in March 2000 (see The Economist, May 20th, 2000), while Island (ISLD) amounted to 6%, and Redi Book, B-trade (Bloomberg's Tradebook), Brut and Archipelago (ARCA) accounted for less than 2% each. Other ECNs include Attain (ATTN), NTRD (Pim Global Equities), and STRK (Strike Technologies.)

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.[1] This best execution rule implies that ECNs, as all brokers, must allocate orders according to price priority. In our understanding, 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). Quite to the contrary, Huang (2000) shows empirically that ECNs are important contributors to the price discovery process.

2.2 Island

Island is a web–based transparent limit order book. It can freely be viewed in real time through internet. When an order is transmitted to Island, if it is not immediately marketable, it is stored in the Island order book. 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

[1] Island has applied for the status of exchanges, but has not yet been granted it.

(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.

Since the best Island quotes are displayed on the Nasdaq screen, if the grid size was the same in the two market segments, the Nasdaq inside touch 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 touch, at least one 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 and the Inside Quote (IQ) file). Regarding this Nasdaq data, in the current version of the paper we work only with March 2000 data; in the next draft we plan to also use June 2001 data.

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 80417 trades for the March 2000 sample and 82838 trades for the June 2001 sample. The average trade size in the March 2000 sample is 559.4 shares, while its counterpart for the June 2001 sample is 645.3 shares. The corresponding average dollar value per trade is $ 54436.45 for the March 2000 sample, and $ 18378.08 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 limit order book on Island and Nasdaq

4.1 Before the decimalization

4.1.1 The inside spread on Nasdaq in March 2000

During that sample period, there was a ot of clustering in the Nasdaq spread on one or two ticks, i.e., one sixteenth or one eigth.[2]

[2]In the current draft, a large fraction of the results presented in this subsection were obtained for DELL and 3Coms. In the next draft we will extend the analysis to the other stocks in our sample.

For example, for Dell, the inside Nasdaq spread in March 2000 was equal to one tick, i.e., one sixteenth, most of the time (94.99 %). Other realizations were 0, only 1.76 % of the time, two ticks, 3.22 % of the time, and three ticks, only .03 % of the time. Correspondingly the time weighted–average of the Nasdaq inside spread for DELL (reflecting market makers's quotes and ECNs' rounded quotes) was $ 16.242/256, which is very close to one tick, and the mode and median were exactly equal to one sixteenth. These results are illustrated on Figure 1.A.

For 3Coms, the distribution of the inside spread was a little less concentrated. The inside spread was equal to 0 only 1.3 % of the time, one tick 80.77 % of the time, two ticks 13.94 % of the time and three ticks 2.977 % of the time. Correspondingly, the time weighted average spread was 19.5/256, while the median and mode were just one tick.

These observations suggest that, before the decimalization, the relatively coarse tick size prevailing on Nasdaq was likely to be rather constraining. Potentially, this could have resulted in excessively wide Nasdaq spreads. We investigate this point further in the remaining of the paper.

4.1.2 The limit order book on Island 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 Nasdaq. We analyse empirically if this took place.

The average spread for our 7 stocks and our sample period was equal to 50.71/256, which is approximately $ 0.198. The mode was 32/256, i.e., $ 0.125.

As illustrated in Figure 1, there is marked clustering in the data. Figure 1.A offers a comparison of the Island and Nasdaq spreads for DELL, while Figure 1.B presents the histogram for all the stocks in our sample. The most frequent spread size on Island was the Nasdaq tick, one sixteenth. But this occured much less often than for the Nasdaq inside quote. Interestingly the spread on Island was also quite often just one Island tick below these levels. In fact, 15/256 is the second most frequent value of the spread. This is likely to reflect undercutting, by a fine increment, of the Nasdaq spread by the 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 924.63 shares at the best ask quote, 952.65 at the second best ask quote, and 844.65 at the third best ask quote. The corresponding average depths on the three first best levels in the book on the bid side were 789.78, 757.56, and 692.9.

4.1.3 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. This can raise synchronicity problems. First, there are point in times when there is more than one change per second in the Nastraq file. In this case, as the Nastraq clock operates on a second by second basis, we cannot tell which was the last change, and consequently we cannot tell which state of the IQ spread prevails after this second. In this case we draw randomly one of the data points in that interval. Second, the Nastraq and Island clocks may not be perfectly aligned. The consequences of non–synchronicity problems are discussed in the appendix.

We obtain the following results for our March 2000 sample for DELL (these results are illustrated graphcally in Figure 3): The best Island bid quote was better than the best Nasdaq (IQ) bid 43.19% of the time, while the Nasdaq bid was better than the Island bid 33.43 % of the time, and the two bids were equal 23.38% of the time. The best Island ask quote was better than the best Nasdaq (IQ) ask 40.97 % of the time, while the Nasdaq ask was better than the Island ask 26.18 % of the time, and the two asks were equal 32.85% of the time.

For 3Coms, the best Island bid quote was better than the best Nasdaq (IQ) bid 31.05% of the time, while the Nasdaq bid was better than the Island bid 45.38 % of the time, and the two bids were equal 23.56% of the time. The best Island ask quote was better than the best Nasdaq (IQ) ask 17.29 % of the time, while the Nasdaq ask was better than the Island ask 47.15 % of the time, and the two asks were equal 35.57% of the time. The Island best quotes are much less frequently better than the Nasdaq Inside Quotes for 3Coms than for Dell. This is likely to reflect the fact that, as DELL is a very liquid stock, its natural spread is very tight, so that the price grid constraint imposed by Nasdaq is more often binding.

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 situatiuon where Island betters the Nasdaq market makers quotes, and we document undercutting by Island orders on a finer grid that the sixteenth grid. In appendix 1 we offer some further discussion of the impact of rounding for the data.

4.2 The limit order book on Island after the decimalization

The average spread on Island for our 7 stocks in June 2001 was equal to $ 0.64, with a mode of $ 0.04. This is definitely below the corresponding figures for March 2000. Figure 5 presents the histogram of spread sizes on Island for our June 2001 sample. As in the March 2000 case, there is a lot of clustering on the Nasdaq price grid: The Island spread is most frequently equal to .02 (more than 10% of the time), .01 and .03 (more than 8% of the time), and .04 and .05 (more than 5% of the time). There is also quite a bit of clustering 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.B) points at the stark reduction in the Island spread generated by the marker reduction in the Nasdaq tick. Interestingly, there was no such reduction in the Island tick during that period.

These results are consistent with the view that the supply of liquidity on Island in March 2000 was not competitive. Under the plausible assumption that the cost of supplying liquidity has not decreased a lot between our two sample periods, the tight spreads prevailing after the decimalization of Nasdaq suggest that Island liquidity suppliers earned rents at the wider spreads they quoted in 2000. Indeed, our results suggest that before the decimalization of Nasdaq, Island liquidity suppliers simply undercut slightly Nasdaq quotes, while leaving the Island spread at a quite proftable level.

To gain more insights in the order placement process on Island, we computed the frequency with which limit orders or trades occurred at different points on the pricing grid. We differentiate four categories:

- prices which are multiples of one dollar, such as, e.g., $ 65,
- prices which are multiples of 10 cents, such as, e.g., $ 65 and 20 cents,
- prices which are multiples of 1 cent, such as, e.g., $ 65 and 21 cent,
- and prices which are expressed using finer resolutions on the grid, such as, e.g., $ 65 and 456 cents.

Our results are presented in Figure 6. Orders are most frequently placed at price levels which are multiples of one cent. But finer levels of the grid (one thousandth) as well as coarser levels (one tenth) are also frequently used. Interestingly, fine ticks are more often used at the best quotes, where undercutting is an important stragey, than further away from the quotes. Also, the finer tick is more frequent for transaction prices than for orders. This is consistent with undercutting orders, placed at finer levels on the grid, attracting trades relatively more frequently than matching orders.

While the spread has decreased after the 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 book. Our results, presented in Figure 2, show that the depth at the quotes is not lower in the June 2001 than in March 2000.

5 The profits earned by Island liquidity suppliers

5.1 Econometric approach

Generalizing (slightly) Biais, Martimort and Rochet (2000) to allow for constant marginal order handling costs (c), the best ask price in the book at time t ($A_{1,t}$) is:

$$A_{1,t} = E(v|Q_t > 0, H_t) + c + \pi,$$

where Q is the market order received at time t, where a positive sign corresponds to a buy order, H_t is the information set of the liquidity suppliers just before receiving the order, and π is their oligopolistic mark–up. Similarly, the best bid price ($B_{1,t}$) is:

$$B_{1,t} = E(v|Q_t < 0, H_t) - c - \pi.$$

Substracting the bid from the ask, the spread is:

$$A_{1,t} - B_{1,t} = \alpha + 2(c + \pi),$$

where:

$$\alpha = [E(v|Q_t > 0, H_t) - E(v|Q_t < 0, H_t)]$$

denotes the informational component of the small trade spread.

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:

$$m_{t+\Delta t} - A_{1,t} = [E(v|H_{t+\Delta t}) + \epsilon_{t+\Delta t}] - [E(v|Q_t > 0, H_t) + c + \pi].$$

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

$$\begin{aligned} E(m_{t+\Delta t} - A_{1,t}|Q_t &> 0, H_t) = \\ E([E(v|H_{t+\Delta t}) + \epsilon_{t+\Delta t}] - [E(v|Q_t &> 0, H_t) + c + \pi]|Q_t > 0, H_t). \end{aligned}$$

Applying the law of iterated expectations:

$$E([E(v|H_{t+\Delta t}) + \epsilon_{t+\Delta t}]|Q_t > 0, H_t) = E(v|Q_t > 0, H_t).$$

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

$$E(m_{t+\Delta t} - A_{1,t} | Q_t > 0, H_t) = c + \pi.$$

A similar equality holds for the bid side:

$$E(m_{t+\Delta t} - B_{1,t} | Q_t < 0, H_t) = -(c + \pi).$$

The intuition is that, on average, the informational component of the spread differences out, so that the difference between the transaction price and the subsquent midquote, i.e., the gross trading profit of the liquidity supplier, is equal to his non–informational cost (c) plus his oligopoly rent (π).

Building on the above analysis, we obtain two moment conditions, which can be used to estimate the deep parameters α, c and π, and also test the model:

$$E(A_{1,t} - B_{1,t} - [\alpha + 2(c + \pi)] | H_t) = 0,$$

and:

$$\begin{aligned} E(m_{t+\Delta t} - (B_{1,t} - c - \pi))1(Q_t &< 0) \\ -(A_{1,t} + c + \pi)1(Q_t &> 0) | H_t) = 0, \end{aligned}$$

where H_t includes all the variables present in the information set of the liquidity supplier at the time of the trade.

To conclude this subsection, we discuss identification issues.

- The second moment condition enables one to identify $c + \pi$. Denote θ the sum of these two parameters. Unfortunately we cannot identify separately the two components of θ. Yet, under the plausible assumption that the order handling cost c is constant across our two subsamples, we can say something about the oligopoly mark–up: π. Estimate $\theta_1 = c + \pi_1$ with the first period data, and $\theta_2 = c + \pi_2$ with the second period data. Note that: $\theta_1 - \theta_2 = \pi_1 - \pi_2$. Thus, if we find that $\theta_1 - \theta_2 > 0$ and since π_2 must be non–negative, we know that $\pi_1 > 0$. Note that changes in the oligopoly mark–up π are not unlikely, since the rules of the game are different in the two samples. Note also that by the same token the information content of trades might also have changed, but this is allowed in ou analysis, and does not alter our estimation of c and π, since α does not enter in our second moment condition.

- Equipped with the estimate of $c+\pi$ obtained from the second moment condition, we can estimate α from the first moment condition. Hence that parameter is identifiable.

To conclude this subsection, we discuss the relation between our approach and the literature.

As regards the theoretical underpinnings of our approach, note that we focus on competition in limit order schedules (in the line of the theoretical analyses of Biais et al, 2000, Glosten, 1994, and Bernhardt and Hughson, 1997), which differs from the signalling 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 size and sign trade. This differs from the conditional expectation condition we use, where the expectation is conditional only on the sign of the trade.

Our structural econometric analysis of competition in limit order schedules is directly in the line of Sandas (2001). Yet we take a somewhat different angle than him. While we focus only on the best bid and ask quotes (and thus analyze only the components of the small trade spread), Sandas (2001) takes a broader perspective, by analyzing orders placed at different levels in the book, and using the size of trades in his econometric approach. On the other hand, while our approach is directly inspired by the theoretical model, and does not request parametric assumptions, Sandas (2000) relies on some parametric assumptions (to specify the link between the size of trades and their information content) and on some short–cuts to the theoretical model (by postulating an exogenous joint distribution of orders and values).

Our 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. Indeed our estimate of $c + \pi$ is essentially based on the average profitability of executed limit orders. Our approach differs from theirs because i) we use the empirical approach to estimate and test a structural model, while they take a more descriptive approach, ii) we focus on the best limit orders only while they study a broader set of orders, iii) we use the fact that we have two different samples over which the order handling cost is likely to remain constant, in order to collect information about the oligopoly mark–up.

5.2 Empirical results

In our 2000 Island dataset, there are 60749 trades. The corresponding average dollar profit earned by the liquidity suppliers posting the first ask or bid quote was .0389. The average dollar bid–ask spread prevailing when these trades occured was .1387.[3] The corresponding figures for our 2001 Island dataset are the following: 62444 trades took place, the corresponding average dollar profit earned by the liquidity suppliers was .00284, and the average spread was .004. To obtain further insights about the economic effects at play behind these summary statistics we implemented the structural econometric approach described in the previous

[3]This is slightly lower than the average spread computed using all observed quotes, reflecting the fact that market orders are more likely to take place when the spread is relatively tight (as was also found, in the case of the Paris Bourse, by Biais, Hillion and Spatt, 1995).

subsection.

To carry the GMM estimation, we used the following instruments (in addition to the constant): whether the ask or bid or spread are on the Nasdaq grid, the current size of the spread, and an indicator of the morning. The results are presented in Table 1.

	$c + \pi$	α
2000	0.0389 (2.00)	0.053 (4.00)
2001	0.00003753 (1.95)	0.036 (3.9)

These estimates show that the two components of the spread ($c + \pi$ and α) are significantly positive. They also show that, while the adverse selection component of the spread is of the same order of magnitude in the two samples, $c + \pi$ is estimated to be much lower after the decimalization. Under the plausible assumption that c did not vary between the two periods, this suggest that liquidity supply on Island was imperfectly competitive before the decimalization.

Note that, our structural econometric approach does not rely on parametric assumptions on the distribution of the underlying variables. Yet, implicit to our approach is the assumption that c, α and π are constant. While the assumption that c is constant through time is reasonable, we might expect that α and π vary with market conditions. A simple, non parametric way to take this into account is to split the sample in subsamples corresponding to similar market conditions. In the next draft of the paper we will do this by differentiating across time of the day and between high and low volatility periods.

6 Conclusion

This papers analyzes the competition to supply liquidity on Nasdaq and Island. Our empirical results suggest the following:

- Before the 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 Nasdasq 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). While both the limit orders undercutting the Nasdaq quotes and those matching them were profitable, the latter were more profitable than the former.

- After the Nasdaq decimalization, the Island spread is much tighter (while the depth at the quotes is not reduced). There is still clustering at the Nasdaq quotes or slight undercutting, but the Nasdaq spread is now set on a tighter, and less constraining, grid.

Our empirical results suggest that the wide dissemination of information and the reduction in the costs of accessing markets brought about by the internet technology have not been sufficient to eliminate market power in financial markets, in particular in the supply of liquidity.

They also underscore that, in addition to the competition between liquidity suppliers within a marketplace, competition between trading mechanisms plays an important role. Our empirical findings point at the competitive pressure exerted by ECNs such as Island on the Nasdaq system. In light of thess results, decimalization on Nasdaq can be interpreted as reflecting (at least in part) a reaction of that market to regain market share. Our empirical findings also point at the reduction in Island spreads, brought about by the reduction of the Nasdaq spread, generated by the decimalization which took place on Nasdaq.

References

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

Bernhardt, D., and E. Hughson, 1997, Splitting orders, *Review of Financial Studies*, 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*.

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

Chen, H., and J. Ritter, The seven percent solution, forthcoming, *Journal of Finance*.

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

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

Harris, L., and J. Hasbrouck, 1996, Market versus limit orders: The SuperDOT evidence on order submission strategy, *Journal of Financial and Quantitative Analysis*, 31, 2, 213–231.

Huang, R., 2000, The quality of ECN and Nasdaq market maker quotes, *Journal of Finance,* forthcoming.

Kyle, P., 1985, Continuous auction and insider trading, *Econometrica*.

Mc Andrews, 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.

Simaan, Y., D. Weaver, and D. Whitcomb, 1998, The Quotation Behaviour of ECNs and Nasdaq Market Makers, Working paper, Baruch College.

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 for DELL. It was equal to $ 31.49/256. Similarly the average Island spread for 3Coms computed using the Nasdaq DQ data (67/256) was greater than its Island data counterpart. 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 quite 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 than in 15 % of the cases synchronicity problems induce mistakes in our best quotes comparisons.

Figure 1.A: Inside spread for DELL on Island and Nasdaq, March 2000

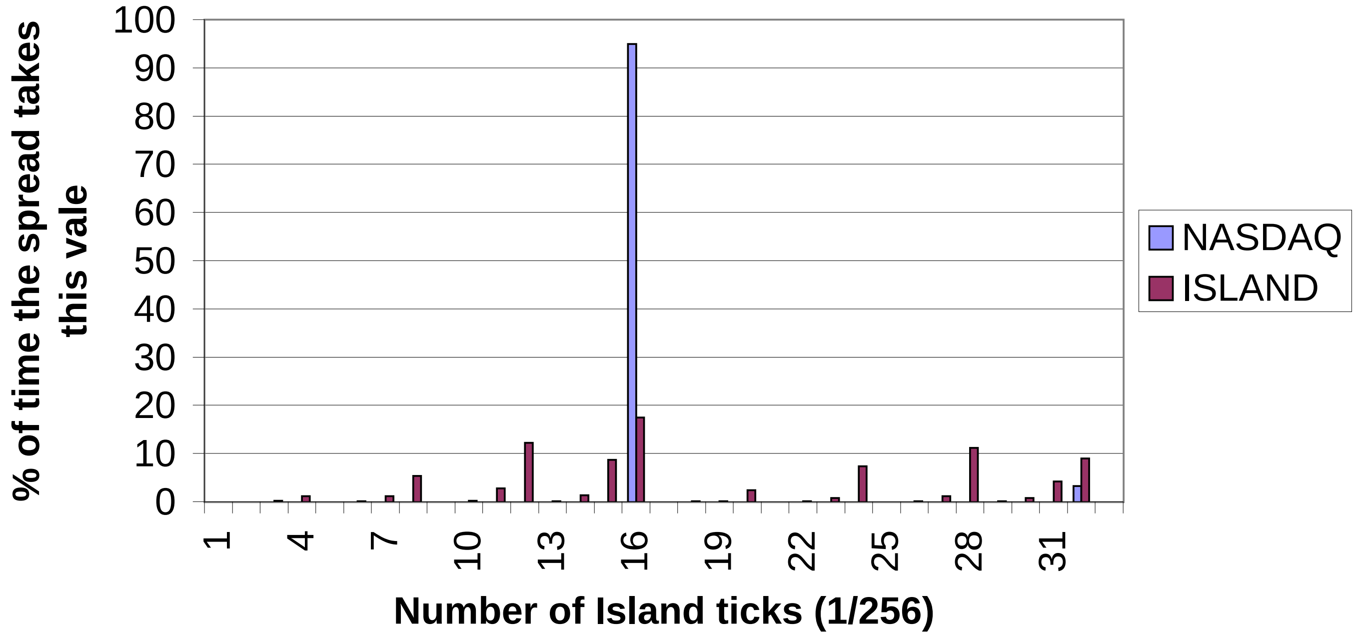

Figure 1.B: Histogram of spread size on Island, March 2000

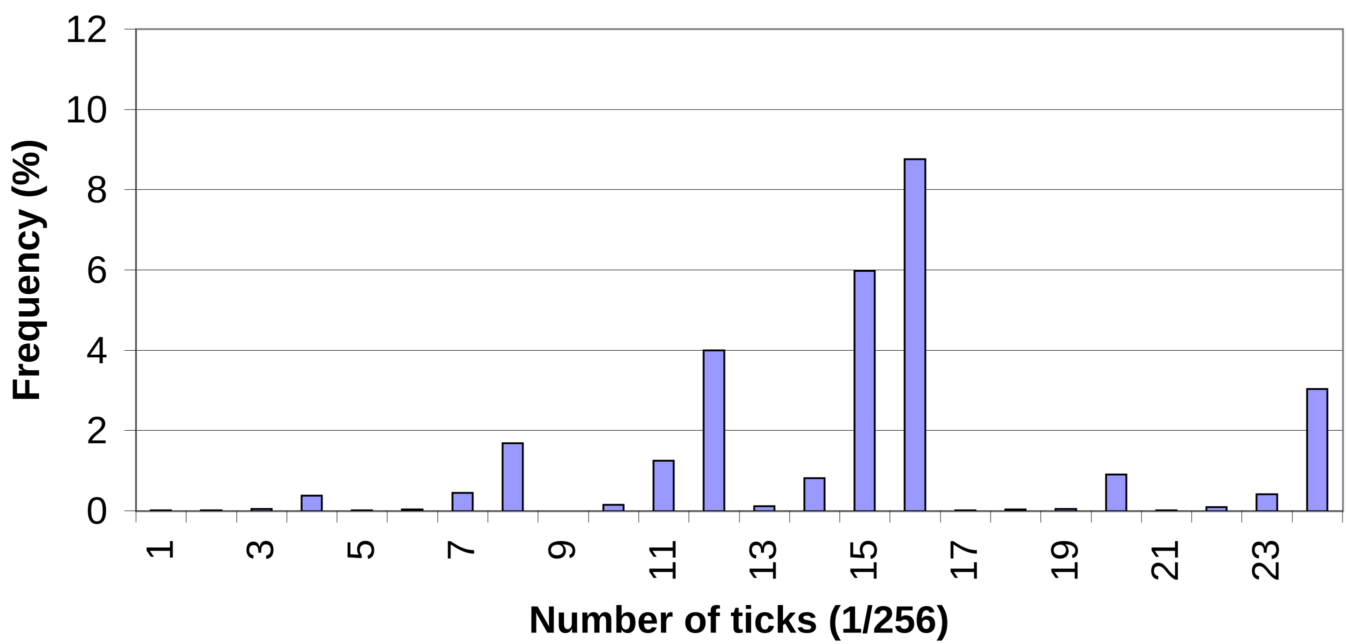

Figure 2: Average depth in the Island book in March 2000 and June 2001

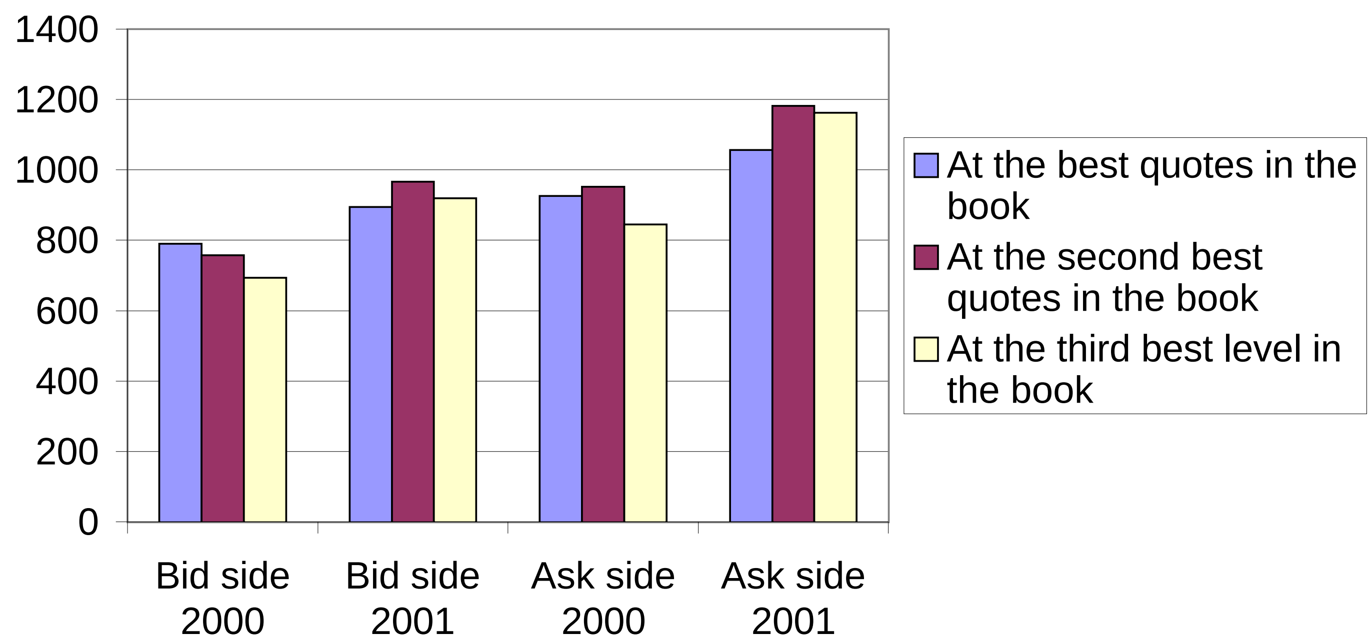

Figure 3: Are the best Island quotes better than the Nasdaq Inside Quotes? (March 2000, DELL)

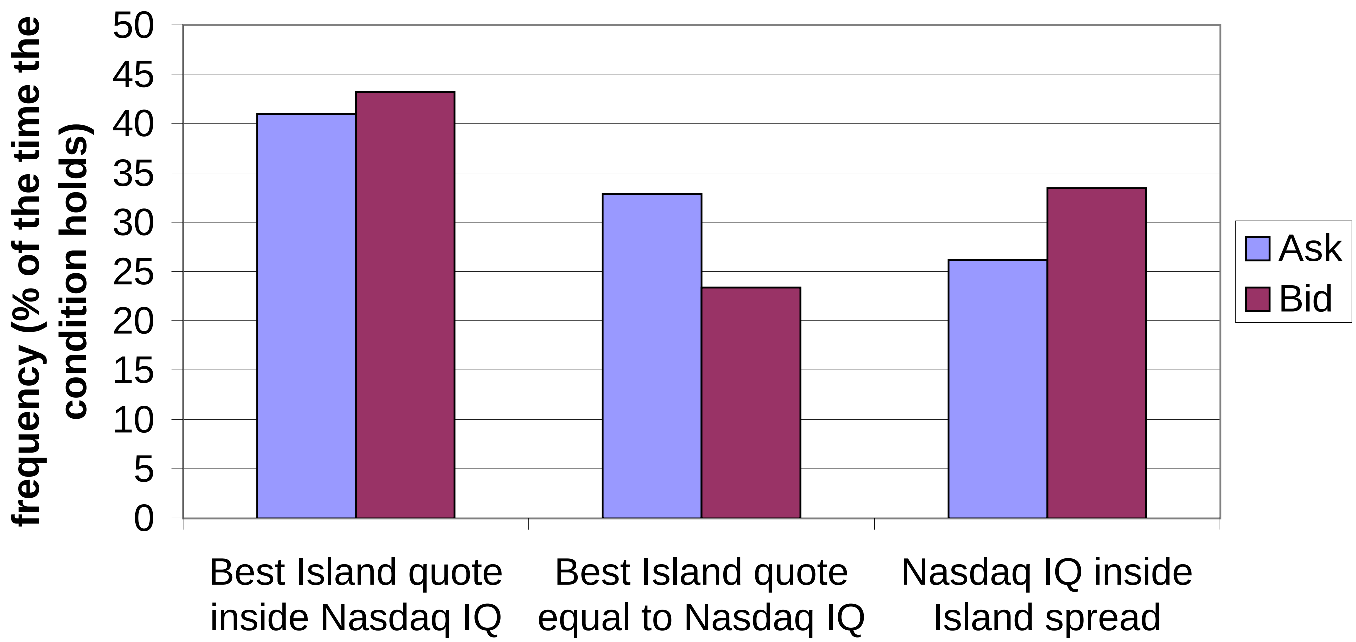

Figure 4: Size weighted mean rate of return of executed Island limit orders, March 2000, all stocks

Figure 5: Histogram of spread size on Island, June 2001

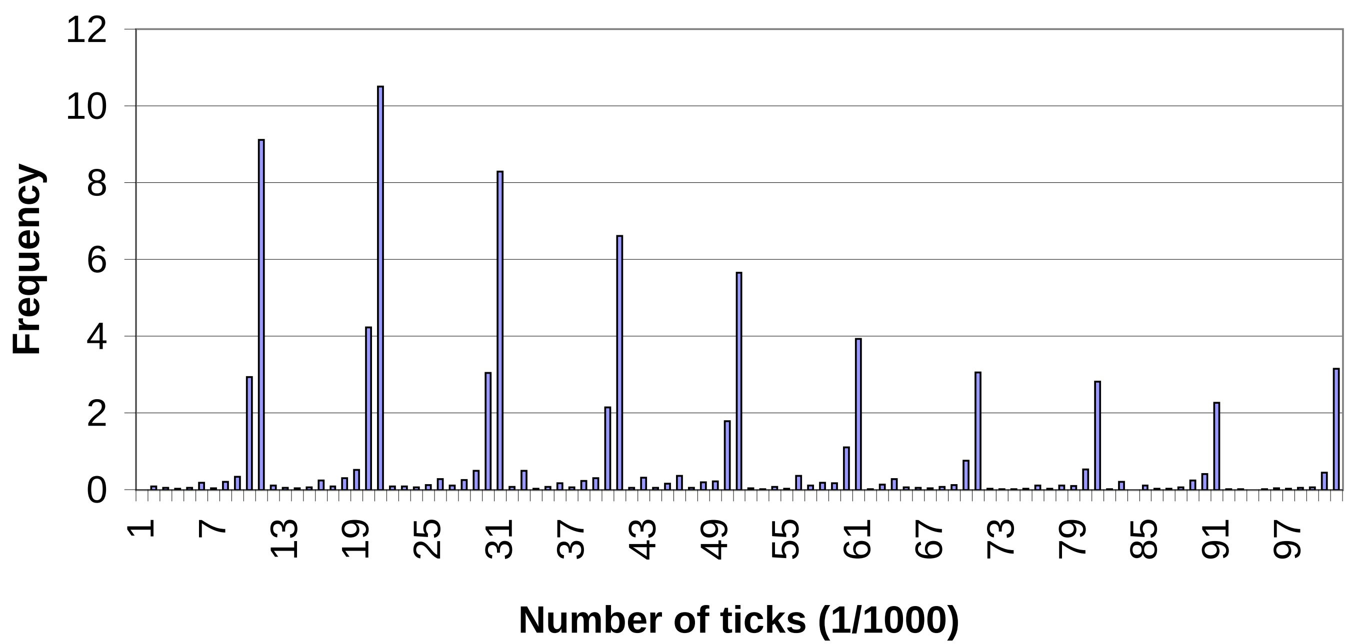

Figure 6: Where on the grid are transaction prices and limit orders on Island? (June 2001)