ON- OR OFF-MARKET TRADING?

Evidence on Competition, LIQUIDITY AND EXECUTION COSTS

Trading stocks on crossing networks can have a profound effect on primary markets as well.

BY JOHANNES A. SKJELTORP & RANDI NÆS

There is currently a plethora of venues for trading equities. Some fit the needs of small retail investors, while others are more suited for large institutional investors and portfolio managers. In this article, we summarize the results from a case study of a large institutional investor in the U.S. equity market who tried to acquire a stock portfolio through crossing networks rather than the primary markets. A crossing network is a satellite-trading place: orders are matched passively (crossed) at prices taken from a primary market. Compared to regular market trading, crossing networks have low commissions, and there is no direct price-impact costs related to a crossed order. On the other hand, there may be opportunity costs associated with the fact that traders are not guaranteed execution in the network. In addition, the execution probability may be associated with adverse selection. This can happen if there are participants in the network with private information. Satellite trading does not contribute to price discovery. This has raised concerns about potentially negative effects from this form of trading on the liquidity in the primary market.

Several studies indicate that crossing networks compete for order flow in large and highly traded companies. To judge whether crossing networks represent a cost-effective alternative for investors, one needs additional information on the execution costs associated with crossing compared with the execution costs associated with the best alternative submission strategy in the market. We have access to a data set that is almost tailored to investigate these questions. The investor in our study was acquiring a U.S. stock portfolio worth around US$1.76 billion during a six-month period from the beginning of 1998. In order to avoid potentially large price-impact costs, the investor initially tried to cross all the stocks in the target portfolio. Our data set includes the dates when the decision to trade was made, the desired quantities of the stocks, and the identity, timing, and subsequent market price for the stocks that failed to be crossed.

In an earlier study of this data set, a significant difference in *ex post* abnormal performance is documented in favour of the non-crossed stocks.[1] This indicates a presence of informed investors in the crossing network. In the current study, it is investigated whether this result can be (partly) explained by a difference in primary

Johannes A. Skjeltorp is a consultant in the research department of Norges Bank, the central bank of Norway. Randi Naes is an advisor in Norges Bank's research department.

Figure 1

THE FUND'S OPPORTUNISTIC CROSSING STRATEGY

The order of submission strategy of the Norwegian Government Petroleum Fund.

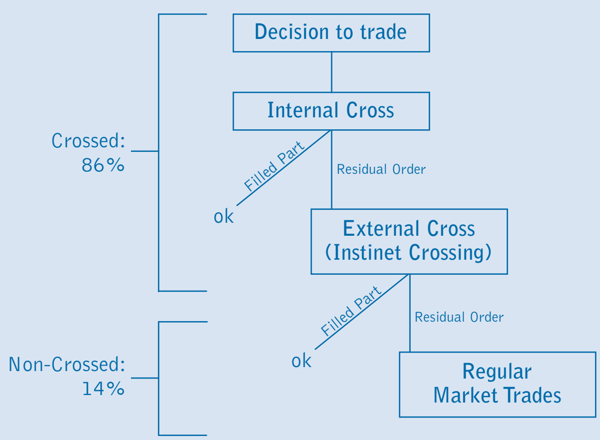

The figure illustrates the order submission strategy followed by the Fund. In each round of trading, the Fund first tried to cross as much as possible with other customers of the transition manager (internal cross). Then the remaining part of the orders were tried to be crossed externally. If the order was not filled after the external crossing attempt, the rest was bought in the regular market.

market liquidity between the two groups of stocks. Moreover, alternative trading strategies are simulated based on real historical price/volume paths of the stocks traded. The simulation exercise provides us with a rough guide to the likely execution costs associated with alternative trading strategies as well as an additional measure of liquidity.

We find support for the hypothesis (stated earlier) that stocks that are easy to cross are highly traded. What is special with our results is that this finding also applies in a sample consisting of only the most liquid stocks in U.S. markets.[2] More interestingly, we find that the actual crossing strategy was inexpensive: even though the crossed stocks were among the most liquid stocks on the NYSE, it would have been very hard to acquire them at lower execution costs by submitting limit orders to the regular market on the same dates. It is tempting to interpret our results as an indication that crossing networks are cream-skimming the order flow. We should keep in mind, however, the finding that the crossing networks were also taking informed orders away from the primary market. The net effect is therefore ambiguous.

The data set

We have access to a detailed data set from the Norwegian Government Petroleum Fund (hereafter the Fund). The Fund is a vehicle for investing the Norwegian Government's income from petroleum-related activities, and a large institutional investor in many countries' capital markets. We have transactions data for the Fund's build-up of a U.S. stock portfolio worth US$1.75 billion during the first half of 1998. The benchmark portfolio was the U.S. part of the FTSE All-World index which currently consists of about 480 of the largest companies in the U.S. To construct liquidity measures and simulate alternative submission strategies, we use additional transaction data from the NYSE Trades and Quotes database (TAQ). The TAQ database contains all the intraday trades and quotes for stocks listed on the New York Stock Exchange, American Stock Exchange (AMEX) and NASDAQ's National Market System and SmallCap issues.

To conduct the transition, the Fund employed four index managers. One of the managers was chosen as a transition manager. The transition manager first tried to find sellers among its own customers (internal crossing). If this was not possible, the manager searched for counterparties among the customers of the other three index managers or sent the order to an electronic crossing network (external crossing). Orders that could only be partly crossed or not crossed at all were sent to the primary market. The Fund had no impact on the outcome of what stocks they were able to get in the crossing network. The strategy is illustrated in Figure 1. According to the discussion in Ruyter (1999), this is the typical order submission strategy followed by large index managers. Nearly 86% of the total portfolio ended up being crossed. For the first two months, internal and external crossing prices were set as the primary market closing prices that day. For the remainder of the period, all crossing prices were set as the volume weighted average price (VWAP) of trades in the primary market during the day.

Execution probability and primary market liquidity

The most interesting piece of information in the data set is the date and identity of stocks that were initially tried crossed, and then subsequently bought in the open market. In Næs/Ødegaard (2000) this information is used to document a significant difference in *ex post* abnormal performance between the stocks that were crossed and stocks that could not be crossed. This result is interpreted as evidence of informed investors in the crossing network.[3] An alternative explanation for the result could be that investors require a higher return on the non-crossed stocks due to differences in liquidity. To investigate this possibility, and to get some information about the nature of competition between crossing networks and the primary markets, we analyze in detail the relation between the probability of getting a stock crossed and the liquidity and trading activity in the primary market.

Liquidity measures

Market liquidity is a comprehensive concept. We try to capture as many dimensions of liquidity as possible by calculating several spread, volume, and volatility measures. Three measures of the spread are considered: the average quoted dollar spread, the average quoted percentage spread, and the effective spread. The effective spread is often considered the most appropriate measure of trading costs because it takes into account that trades often are executed inside (price improvement) or outside the spread ("walking the book"). To capture market depth, we calculate the average quoted number of shares on the inner quotes, and the daily and intraday Amivest liquidity ratio. A high liquidity ratio indicates better ability for the market to absorb large trades without affecting the price. To get a broader picture of the volume and trading activity in the primary market, we also calculate total shares traded, the dollar value of shares traded, and the average trade size. Finally, we calculate two measures of volatility: the standard deviation of daily returns over the 10 days prior to the initial crossing attempt, and a 15-minute return standard deviation for capturing intraday volatility.

Results

In Table 1 we report the different liquidity measures over the relevant dates for three order categories: "Crossed stocks" are the orders that were fully crossed, "Cross/Market" are the orders that were partly crossed and partly bought in the market, and "Market stocks" are the orders that were fully filled in the market. To confirm our sample with the stocks in the S&P500 index, we also calculate the average liquidity measures for the S&P500 index over the same dates as well as for the entire period when the Fund was trading.

The numbers in the table strongly indicate that stocks that were easy to cross had higher liquidity than stocks that were hard to cross. The relative spread difference between the "Crossed stocks" and "Market stocks" (non-crossed stocks) group is 22%, which is both economically and statistically significant.

Table 1

AVERAGE LIQUIDITY AND ACTIVITY MEASURES

For each group tests for differences in means between the crossed stocks and the two other groups are performed where * indicates significant difference in the measure at the 5% level and ** indicates a significant difference at the 1% level. We also test whether there is a difference between the sample stocks and the S&P500 stocks where an *a* indicates a significant difference at the 1% level.

All Trades	S&P500 Jan-July	S&P500 stocks	Fund stocks	Crossed stocks	Crossed/Market	Market stocks
Spread measures						
Effective spread	0.1151	0.1143	0.1106	0.0982	0.1034	0.1201**
Quoted $ spread	0.1360	0.1359	0.1374	0.1222	0.1312*	0.1478**
Quoted % spread (midp.)	0.3414	0.3586	0.3023[a]	0.2629	0.3068**	0.3197**
Volume measures						
Trades	758	772	758	890	911	621*
Shares traded (1000)	1073	1097	1079	1404	1181	871**
Volume (US$ mill.)	58	57	56	76	68	42**
Ordersize (US$ 1000)	77	74	74	83	68**	67**
Liquidity ratios and depth						
Daily LR (US$ mill.)	117	108	121	180	144	82**
Intraday LR (US$ mill.)	13	12	12	18	14	8**
Depth at quotes (shares)	1851	1385	1652	1805	1596	1601
Volatility and return						
Daily volatility (10 day std.dev)	0.0238	0.0250	0.0261[a]	0.0241	0.0266**	0.0269**
Intraday volatility (% std.dev)	0.2785	0.3048	0.2847[a]	0.2960	0.3128	0.2662**

Table 2

IMPLICIT COST ESTIMATES FOR DIFFERENT TRADING STRATEGIES

The table shows the implicit cost estimates for the different trading strategies. The estimates are based on the implementation shortfall methodology. The implementation shortfall is the difference in performance between the portfolio of actual trades and a matching paper portfolio where the trades are assumed executed at the price prevailing on the date of the decision to trade. Numbers in bold indicate that the implicit cost is significantly different from zero at the 1% level. For each strategy and across groups of stocks, tests for differences in means between the original strategy and the simulated strategies are performed where ** indicates a significant difference in implicit costs at the 1% level. The difference between the equally and volume-weighted results indicate that the largest orders were the most expensive across all order types.

Implicit costs	Opportunity Cross	Pure Cross	LO1	LO2	LO3
Equally weighted					
All orders	**0.0879**	**0.1443**	**0.0603**	**0.1281**	**0.2470****
Crossed orders	0.0553	0.0553	**-0.0147****	**0.0520**	**0.1729****
Non-crossed orders	**0.2536**	**0.5867**	**0.4317****	**0.5048****	**0.6143****
Volume weighted					
All orders	0.2028	0.2534	0.0770	0.2706	0.3762
Crossed orders	0.1837	0.1837	0.0141	0.2007	0.3025
Non-crossed orders	0.3101	0.5867	0.4298	0.6615	0.7892

Measured by the number of trades, the trading volume, and the number of shares traded, the trading activity was also significantly lower in the non-crossed stocks over the entire sample. We confirm these results more formally by estimating a probit regression model of the probability of getting a stock crossed as a function of various liquidity measures.

Limit order simulation

Since the Fund was trading in the 500 largest and most liquid companies in the U.S. market, it could well be that a strategy of buying them directly in the market would have been less expensive than the crossing strategy followed by the Fund. To investigate this, we simulate alternative submission strategies in the open market. Note that we do not set up and estimate a model that maximizes the execution probability given the stock and market characteristics at the time of submission. Rather, we simulate "naive" order submission strategies to get a "lower bound" on the implicit execution costs of alternative submission strategies in the primary market. In addition we get an indication of how fast a trade for a given number of shares can be executed. This dimension of liquidity is important and not directly captured by the measures reported in the previous section.

Simulation design

We simulate three limit order strategies on the same stocks and on the same dates when they were tried crossed by the Fund. In the first limit order simulation (LO1) we do not take into account the actual order sizes traded by the Fund. In other words, we assume that only one share is traded in each stock.[4] At the beginning of each crossing date a limit order is submitted with a limit price equal to the opening bid-quote ("at the quote" limit order strategy) for each stock that the Fund tried to cross. If a trade with a price lower than the limit order price is observed during the day, the order is assumed to be filled. If an order is not filled, we assume that it is executed at the opening price the next day. Because we are ignoring order size, the first limit order simulation constitutes a lower bound on transaction costs. In the second simulation (LO2) we split the actual order size into suborders. The number and size of the suborders are determined by the average order size traded in the stock the day before the Fund traded. All the suborders are assumed to be submitted sequentially. A suborder is assumed filled if the observed execution price is less than the limit price. When a suborder is filled, the next suborder is submitted at the bid quote following the fill ("chasing the market"). Unfilled orders are assumed to be executed at the opening price the next day. The third limit order simulation (LO3) is similar to the LO2 strategy except that we also take the size of the suborder into account. A suborder is only assumed filled if the observed execution price is less than the limit price and the size is equal to or larger than the size of our order. This is the most realistic of our simulated limit order strategies. In addition to the three limit order strategies, we simulate a pure crossing strategy where we assume that the Fund only traded in the crossing network. Orders that the Fund could not cross are assumed to be crossed at the closing prices in equal amounts over the 10 days after the first crossing attempt failed.

Results

Table 2 summarizes the results from our simulation exercise. In order to compare the performance of different submission strategies, we apply an empirical version of the so called "implementation shortfall approach" proposed by Perold (1988). This measure is designed to capture both explicit cost components such as broker's fees, and implicit components such as spread costs, price impact costs, and costs related to delayed or uncompleted trading (opportunity costs). We only focus on the implicit cost components. The cost estimates are given as a percentage of the closing price on the day before the decision to trade.

Neither the pure crossing strategy nor the two first limit order strategies (LO1 and LO2) have significantly different execution costs from the opportunistic crossing strategy. Thus, not even the most simplistic and unrealistic limit order simulation (LO1), which constitutes our "lower bound" on primary market execution costs, is able to significantly beat the actual crossing strategy. The most realistic limit order strategy (LO3) is shown to be significantly more expensive. In addition, we have not taken into account that the explicit costs in crossing networks are lower than in the primary market. Hence, the total execution costs would overwhelmingly favour the opportunistic crossing strategy, or potentially the pure crossing strategy. Some studies indicate that market orders are cost-effective alternatives to limit orders because market orders (worked orders) often achieve significant price improvements. However, it is hard to model price improvements and price impacts, and we have therefore not attempted to simulate market order strategies.

The non-crossed stocks have the highest execution costs regardless of submission strategy. This result supports the finding that these stocks are less liquid than the crossed stocks. Note also that the opportunity costs constitute a large part of the costs for orders in these stocks. Since unfilled limit orders generally are for stocks that rise in value, these orders are penalized by the execution at the opening price the next day. This result, together with the high costs found for the pure crossing strategy, is just another way of stating the finding that the stocks bought in the market had a high *ex post* return. Our results indicate that the *ex post* abnormal returns documented in the earlier study might be partially accounted for by a liquidity premium for the non-crossed stocks in the primary market.

Conclusions

Overall, our results indicate that crossing is a cost-effective trading method for large uninformed investors. The stocks with the lowest execution costs in the U.S. primary markets seem to be cheaper to obtain through crossing networks, even without taking into account that crossing commissions are significantly lower than broker commissions.

However, one should keep in mind the result in Næs and Ødegaard (2000) that the crossing network also takes informed orders away from the market. We do find a highly significant difference in liquidity between the two groups of stocks which is statistically significant. But, a liquidity premium between the two groups is probably not large enough to explain the entire difference in abnormal performance of 1% per month documented in their study. ▮

References

P. Handa and R. A. Schwartz, "Limit order trading." *Journal of Finance*, 51, pp.1835-1861, 1996

R. Næs and B. A. Ødegaard, "Equity trading by institutional investors. To cross or not to cross?" Discussion Paper 3/2000, Norwegian School of Management BI, 2000

A. Perold, "The implementation shortfall: Paper versus reality." *Journal of Portfolio Management*. 14(3), pp.4-9, 1988

T. Ruyter, "The dinosaur, the rock, and the hard place." *Plan Sponsor*, March 1999

Endnotes

* The views are those of the authors and should not be interpreted as reflecting those of Norges Bank.
1. See Næs/Ødegaard (1999).
2. The U.S. part of the FTSE All World Index
3. The difference in risk-adjusted returns is found to be as much as 1% over the 20 days following the trades. This result indicates that the "best" stocks are hard to get in the crossing network, a result which is inconsistent with pure liquidity trading.
4. See Handa/Schwartz (1996).