

Equity Trading by Institutional Investors: To Cross or Not to Cross? The Case of the Norwegian Petroleum Fund.

Randi Næs

Norges Bank*

Bernt Arne Ødegaard

Norges Bank

and

Norwegian School of Management

13 January 2000

Abstract

The costs to institutional investors of trading equity are of obvious practical as well as academic interest. The current academic literature on this issue shows that there are large unresolved issues, both regarding the components of the costs (implicit/explicit costs, costs of non-trading, market impact), and magnitude of the costs. To date, the empirical academic literature has concentrated on data from equity trading at organized exchanges. This paper adds to the extant research by including evidence on using alternative mechanisms for facilitating equity trading, so called crossing. In this paper we use equity trades of one large institutional investor, the Norwegian Petroleum Fund, to investigate the costs of trading equity using alternative trading venues. The results show that for trades that *were* crossed, the average implicit and explicit costs were lower than found in similar cases in the academic literature. We do, however, find that the orders that did *not* get crossed were special. Using as a benchmark the date of the desire to order, we find that they have a *negative* cost. We view this difference between crosses and market orders as promising issues for further research.

Keywords: Costs of Equity Trading, Trading Mechanisms, Crossing, Institutional Equity Trading.

JEL Classification: G10, G23

*The views expressed are those of the authors and should not be interpreted as reflecting those of Norges Bank. Comments from Steinar Ekern are gratefully acknowledged.

Introduction

The costs to institutional investors of trading equity are of obvious practical as well as academic interest. The academic literature on this issue, much of which is summarized in Keim and Madhavan [1998], shows that there are large unresolved issues, both regarding the components of the costs (implicit/explicit costs, costs of non-trading, market impact), and magnitude of the costs. To date, the empirical academic literature has concentrated on data from equity trading at organized exchanges. This paper adds to the extant research by including evidence on using alternative mechanisms for facilitating equity trading, so called crossing. In this paper we use equity trades of one large institutional investor, The Norwegian Petroleum fund (hereafter “the fund”), to investigate the costs of trading equities using alternative trading venues.

The focus of the paper is on the strategy for performing the actual equity transactions. There are different ways of transacting in equity markets. One can trade at the established exchanges, or use alternative markets, such as “crossing.” A crossing network is an electronic network where participants typically submit quantities desired of a stock and the quantities are matched, either automatically or manually. The agreed price in a cross will typically be derived from the primary market, such as the closing price at NYSE.

There are several reasons to believe that a crossing network provides reductions in the costs of trading equity for an investor such as the Petroleum fund. First, crossing commissions are low compared to commissions charged by exchange brokers. Second, there is no bid-ask spread in a cross because liquidity is provided by the traders themselves and not by dealers. Finally, there is no direct price impact because prices in crossing networks are set independent of order size. There may however be an “implicit” price impact if the existence of a large crossing order is known to participants in the primary market. Also, we don’t know what a crossed stock could have been bought or sold for in the open market, and thus it is hard to say how “good” the obtained price is. The anonymity provided by crossing networks makes this trading method attractive to informed traders as well. Uninformed traders might therefore incur costs related to adverse selection. Because crossing networks do not guarantee execution, opportunity costs could also be significant (see Harris [1993] and Keim and Madhavan [1998]).

The data on the Petroleum fund’s transactions in equities are special in that we know that the fund only traded in the market when the particular stocks could not be crossed. Coupled with transaction data from the relevant exchange, our data set should therefore enable us to look more closely at the costs related to adverse selection and missed trading opportunities for investors using crossing networks.

Because of our focus on alternative trading strategies, we have chosen to use data for US equity markets. There are several reasons for this. We want to compare with extant research, most of which is done using US data. The practice of crossing is most prevalent in the US market. It is also easy to get relevant microstructure data for the NYSE.

It is far from clear how investors should submit orders so as to minimize their costs of trading. We do know, however, that order size is an important determinant of trading costs. As a first step towards an analysis of the fund’s submission strategy we therefore consider the absolute and relative size of the fund’s transactions. The Petroleum fund is not found to be a dominant player in the US equity market.

To get an idea of the fund’s relative trade performance, we compare the fund’s transaction costs to

average cost estimates reported in the literature. The average cost of trades for the fund compares favorably to cost estimates reported in the literature for similar institutions. Our results show that there are significant differences between crossed and market orders. Investigating the causes of these differences is one of several interesting topics for further research which we discuss at the end of the paper.

The paper is structured as follows. Section 1 summarizes some of the relevant academic research in the area, whereas section 2 gives an institutional overview of trading in the US equity markets. In section 3 we discuss some details about the fund's trades and the microstructure data from NYSE. In section 4 we investigate how large the fund's trades are, both in an absolute and a relative sense. We then, in section 5, compare the fund's trading costs with relevant cost estimates from the academic research. Section 6 suggests some avenues for further research, before section 7 concludes.

1 Relevant Research

1.1 Empirical evidence on trading costs

Much of the relevant research in this area is recently summarized in a paper by Keim and Madhavan [1998]. We will use their categorization of the costs of trading equity. These costs can be split into two main categories, explicit costs and implicit costs.

The explicit costs are the actual out of pocket costs of trading, such as brokers fee's. These costs are easy to quantify. There has been a significant reduction in broker commissions in the US equity markets over the last decade. In another recent study by Keim and Madhavan [1997], the average commissions reported is 0.2% of trade value. Systematic variations in commission costs are documented, depending on what broker type and market mechanism investors use. There also seem to be a positive correlation between the explicit and implicit cost components of a trade, meaning that commissions tend to be higher the larger and more complex the trade is.

The implicit costs are much harder to quantify, and there are differences of opinion about their existence and relative significance. There are three main suggested components of the implicit costs: bid ask spread, price impact and opportunity costs.

1.1.1 Bid Ask spread

Consider the example of a specialist market structure like the NYSE. The specialist will at any time quote bid and ask prices valid for a minimum quantity, typically 1000 shares. The difference between the bid and ask price can be viewed as the price that the specialist demands for his services. Early work in market microstructure focused on the quoted spread, and used half of this spread as an estimate of the cost of a one way transaction. There are however problems using the quoted spread as a cost measure. Trades may occur inside the spread, if either the specialist or one of the "trading crowd" better the price. Large blocks may be traded outside the exchange at negotiated prices.

These problems have lead several authors to propose measures of a "true" bid ask spread, often referred to as the "effective bid ask spread." These measures are ex post measures estimated from transaction prices, and try to find an "average" spread that a transaction was exposed to.

Estimators of the effective spread tend to be considerably smaller than the quoted spread. Lee [1993] uses an estimator based on quotation data and finds evidence that the effective spread is only about half of the quoted spread. Madhavan et al. [1997] use a version of the serial covariance estimator first proposed by Roll [1984], and find that the quoted spread was almost three times greater than the effective spread for a sample of 274 NYSE stocks in 1990.

1.1.2 Price Impact Costs

The bid ask spread does not account for the fact that for large orders, prices may have to move in order to be able to execute the order. Or, in other words, traders may need to “walk the demand or supply curves.” The resulting price impact may be decomposed into a temporary component reflecting the liquidity cost of the trade, and a permanent component reflecting new information. The information cost is related to the adverse selection problem studied in most of the theoretical market micro structure literature. There is always a risk that an order is informed, and this is presumably larger for large orders.

In theory, the total price impact of a trade can be easily computed if one knows what the price of the stock would have been if the trade had not occurred. In practice, this so called “unperturbed” price is of course not observable. A common empirical measure of the price impact is the deviation between the transaction price and a proxy for the unperturbed price, where the proxy is some weighted average of pre- and/or post-trade prices for the stock.¹

In studies of large-block trading it is common to focus on pre-trade benchmark prices. These studies document significant price impact for trades of 10000 shares and above. Keim and Madhavan [1996] show, however, that the choice of a benchmark price makes a large difference in the estimated price impact. Based on data on block-trades for one institutional investor, they find that the average price impact for a seller-initiated transaction vary from -4.3% to -10.2% when the unperturbed price is defined as respectively the previous day’s close and the price three weeks before the trade. This result strongly suggests that the unperturbed price for block trades should be defined as the date on which the decision to trade was made.

1.1.3 Opportunity Cost

The final source of the implicit costs of trading equity is the opportunity costs of not trading. This cost is due to the investor not being able to accurately implement the desired portfolio. Some equities are delayed, during which time the market price may move in an undesirable direction. In other cases the investor may not be able to fill the order at all. For index trackers this cost may be important, because one may be exposed to “tracking errors” whereas the total portfolio deviates from the desired one.

Treynor [1981] has proposed a theoretical measure of the total cost of trading which incorporates all the mentioned cost components including the opportunity cost of not trading. This measure, which Perold [1988] called the “implementation shortfall”, is defined as the difference in performance between the portfolio of actual trades and a matching “paper” portfolio where the stock returns are computed assuming that the trades were executed at the prices prevailing on the dates of the decision to trade. In addition to capturing all relevant cost components, the implementation

¹This measure captures one-half of the bid-ask spread plus the price impact.

shortfall overcomes the problem of measuring costs on an individual trade basis when the order consists of a package of sub-trades.

To correctly measure the opportunity cost of not trading, more detailed data are necessary than those readily available to most researchers. First, data on the date of the decision to trade are required. Second, to be able to correct for dynamic inconsistencies, researchers should have detailed information on the underlying motivations for the trade, such as investment objectives, target price, and trade horizon.² Even this may be insufficient, however, because a strictly correct measurement sometimes will require unavailable information such as the data on trades that never took place.

Fortunately, relevant data on the order submission process of institutional investors have been increasingly available. As a result, the most recent empirical studies are in fact based on the implementation shortfall approach.

1.2 Evidence on the Determinants of Trading Costs

There is an extensive empirical literature on the measurement of trading costs. However, because the different cost components of a trade are typically jointly determined, and because individual trades are often part of a larger package of trades, one cannot make inferences about the cost of a trade by adding up separate unconditional estimates of the component costs found in previous studies. Recent studies based on the implementation shortfall approach overcome these problems, but require data which are not readily available. In Keim and Madhavan [1998], the current state of knowledge on the determinants of trading costs of institutional investors is summarized as follows:

- “Implicit trading costs are economically significant when compared with explicit costs (and relative to realized portfolio returns).
- Equity trading costs vary systematically with trade difficulty and order placement strategy.
- Other important determinants of trading costs are differences in market design, investment style, trading ability, and reputation.
- Even controlling for trade complexity, there is considerable variation in trading costs.
- Accurate prediction of trading costs requires more detailed data on the entire order submission process, especially information on pre-trade decision variables such as the trading horizon.
- For institutional traders, the concept of “best execution” is difficult to measure and hence enforce.”

2 The institutional framework

In this section we provide a brief overview of the institutional framework for trading equity securities in the US capital markets.³ Trading alternatives for US stocks includes

²The cost estimate should be able to capture that a trader for example has chosen to scale back his or her order due to market movements.

³Sources for the discussion in this section include O’Hara [1995], Hasbrouck et al. [1993], Arshadi [1998], the NYSE website at <http://nyse.com>, the Nasdaq website at <http://Nasdaq-amex.com>, the Instinet website at <http://www.instinet.com>, the ITG website at <http://www.posit.com> and the AZX website at <http://www.azx.com>.

- NYSE or one of the five regional exchanges,
- the National Association of Securities Dealers Automated Quotes System (Nasdaq), or
- a so called Alternative Trading System (ATS).

One important feature of how a market system operates is how prices are determined. Prices are typically driven by orders or quotes. A market is said to be order driven if some investors, by placing limit orders, establish the prices at which trading takes place.⁴ In a quote driven market, dealers announce the prices at which they are prepared to buy and sell.

2.1 The NYSE

At the NYSE, a group of traders, known as specialists, are responsible for maintaining an orderly market in specific stocks assigned to them. Limit orders are executed against incoming market orders. Occasionally, the specialists buy and sell from personal inventories in order to improve market liquidity and to reduce the spread. In return for these services, they get full access to the limit order book and first look at the orders arriving on the computerized routing systems. The role of the specialists implies that prices are partly quote and partly order driven. Orders originating off the floor reach the exchange either electronically through the so called NYSE superDOT system or through the floor brokers. Floor brokers typically represent the orders which are largest and most difficult to execute. The specialists act both as agents and principals. Other traders are member firm brokers who trade as agents for their customers, and independent floor brokers who trade as principals for profits.

2.2 The Nasdaq

The *Nasdaq Stock Market* has traditionally been a quote driven market with market makers standing ready to buy and sell for their own accounts over the market's networks of terminals. Customers submit their orders to brokers. If the broker's own firm makes market in the security, the trade may be executed by the firm taking the counterpart position. Otherwise the broker routes the order to a market maker with the best quote and executes the order. After the implementation of new order handling rules in January 1997, the Nasdaq market is no longer fully quote driven. The "limit order display rule" requires that Market Makers display investors' limit orders in their quotes when they are priced better than the Market Makers quote. Previously, limit orders were halted until the dealer's quote reached the limit price. The "quote rule" requires Market Makers to publicly display their most competitive quotes. This means that a customer who place a market order can be sure that the quoted price reflects the most favorable price at which Market Makers are willing to trade.⁵

⁴*Limit orders* request execution at a specified price or better; i.e. they will be executed only if and when the specified price is reached. *Market orders* request immediate execution.

⁵Previously, Market Makers could place orders on private systems with better prices than their public quotes. Private system prices were only available to financial professionals.

2.3 Alternative Trading Systems (ATs)

Well known US examples of trading mechanisms that facilitate trading without brokers and exchanges are Instinet Corporation, the Portfolio Systems for Institutional Trading (POSIT), and the Arizona Stock Exchange (AZX). Currently, Instinet and POSIT are registered as broker-dealers whereas AZX is registered as an exchange. The recent success of these markets is mainly due to the anonymity they provide to traders as well as to low commissions.

Instinet Corporation was founded in 1969. Originally, Instinet was designed to allow institutions to trade NYSE stocks directly with each other through a computerized "institutional network". Later, the operations were expanded to include OTC stocks and their market makers as well. Today Instinet provides block trading, order management and crossing services in global equities. Customers can either (i) enter anonymously an order to buy or sell a given number of shares at a given price, (ii) accept a posted bid or ask in the electronic book,⁶ or (iii) try to negotiate a better price by sending a private, anonymous message to the party who posted the original order. Instinet Crossing is an after-hours equity trading product where orders from anonymous clients are matched based on benchmark prices. The benchmark prices are the last sale on the primary exchange and the midpoint of the closing inside bid and ask for OTC securities.⁷

POSIT electronic trade matching was introduced by Investment Technology Group (ITG) in 1987. Just like Instinet Crossing, orders entered into the POSIT system are matched and executed at a single price. The price is the midpoint of the bid/ask spread in the stock's primary market at the moment the match is run. In the beginning POSIT had only one daily match. Today there are six daily matches. Orders which match are automatically executed. Customers can choose to keep unmatched orders in the system for future matches. All traders are anonymous and no user knows the specifics of orders and residuals.

In 1991, another electronic single price trading system was established, known as the *AZX*. The price discovery mechanism used by *AZX* may be described as a modern version of the old call markets.⁸ Buy and sell orders are accepted throughout the day at the price range the traders want to trade. Orders are all limit-or-better orders and they go into the Open Book or the Reserve Book. The Open Book orders are immediately displayed (anonymously), whereas Reserve Book orders are not seen until the auction is over. At certain times on Mondays, Wednesdays, and Fridays, orders are matched. The auction price is set such that supply and demand, as represented by Open Book asks and bids respectively, are most nearly balanced. Open Book orders are filled first - by time of entry. Reserve Book orders are filled next - more aggressive prices before less aggressive prices. *AZX* uses a different pricing mechanism from Instinet Crossing and POSIT in that it establishes its own prices. Another important feature of the *AZX* is that it provides anonymity to their traders whereas at the same time allowing them to monitor the relative strength of buy and sell orders for the underlying stocks.

⁶The electronic book on a stock includes information about dealers' bid-ask quotes, the number of trades, the number of outstanding orders, and the number of customers monitoring the stock.

⁷The inside bid and ask is respectively the highest bid and the lowest ask being quoted among all the Market Makers in the stock

⁸In a call market, orders are aggregated so there is only one large trade between many parties at a consensus price.

2.3.1 Relative Efficiency of Different Market Systems

The extent to which differences in market structure affect trading costs is of great interest to all market participants.

Several studies find that the quoted and effective spreads on Nasdaq stocks are wider than comparable NYSE listed stocks. This suggests that small traders, whose price impact costs and opportunity costs are of little relevance, face lower trading costs on the NYSE than on Nasdaq. As discussed in the section on relevant research, a simple comparison of spread costs would be misleading for large traders. Large traders often execute their trades at prices inside or outside the quotes. Besides, they often face non-negligible costs from price impact and missed trading opportunities. For these investors, there are mixed empirical evidence as to which market system is preferable.

There is an extensive literature on the relative efficiency of the NYSE and Nasdaq. To our knowledge, however, there are few results on the relative efficiency of alternative trading systems.

3 Data

3.1 The Norwegian Government Petroleum Fund

In 1990 the Norwegian Government established a fund as a vehicle for investing the Government's income from petroleum related activities in international capital markets. The Ministry of Finance is responsible for the management of the fund, whereas Norges Bank (the Central Bank) performs the operational portfolio management. Norges Bank is to optimize the returns on the investments subject to the investment restrictions set by the Ministry of Finance. Funds were allocated for the first time in 1996 and invested according to the guidelines used for the foreign exchange reserves. By the end of 1997 the fund value was about USD 15.4 billion, invested mainly in foreign government securities.

In October 1997 the Ministry of Finance issued new regulations for the fund, to apply from 1 January 1998. The new guidelines state that between 30% and 50% of the fund is to be invested in equity securities.⁹ The composition of the fund portfolio was changed into partly equity during the first half of 1998, by buying equity in (mainly developed) foreign markets. The new investment strategy was implemented through a specified benchmark portfolio.¹⁰ The fund employed four index managers to establish the equity portfolio. Most of the trades were crossed, mainly through an informal crossing network established among these index managers. The remaining trades were executed in the open market. By the end of June 1998, the market value of the total Petroleum Fund portfolio was USD 17.7 billion, and the market value of the stock portfolio was USD 7.2 billion. US stocks represent 28.5% of the benchmark portfolio for the equity securities.¹¹¹²

⁹The portfolio is rebalanced each quarter to give an asset distribution of 40% equities and 60% fixed income securities.

¹⁰The benchmark portfolio forms a basis for actual investments and serves as a reference for evaluation of investment performance.

¹¹The benchmark portfolio for the equity investments is based on the Financial Times/ Standard & Poor's Actuaries World Index (FT/S&P). FT/S&P aims to cover around 80 percent of the value of all listed stocks in 29 countries worldwide. By the end of 1998, the index consisted of 2400 stocks. The Petroleum Fund benchmark stock portfolio consists of 2040 of these stocks from 21 countries.

¹²As the US stock markets are among the biggest and most liquid markets in the world, the conclusions drawn here on the performance of the Petroleum fund's transactions do not necessarily carry over to the rest of the Petroleum

Table 1 summarizes the implementation of the US part of the fund’s equity portfolio. The portfolio was established in the period from January 1998 to June 1998. The total portfolio investment was USD 1751 mill. Of this amount USD 1501 mill or nearly 86 percent was crossed.¹³ The fund traded at NYSE at 3 of the 16 trading dates. The last column shows each trading date’s percentage of the total dollar portfolio investment. The highest trading volume on one date amounted to USD 300 mill, or 17.1 percent of the total portfolio investment.

Table 1 Establishing the US stock portfolio.

Transaction volume in million USD, for each date on which the fund traded. January to June 1998. For anonymity reasons we do not show the actual dates, but the table is in chronological order.

date	Crosses		Market	All
	Internal	External		
1	174			174
2	184			184
3		115		115
4	58		3	61
5	19			19
6		30		30
7	163		73	236
8	300			300
9	14			14
10	14			14
11	231			231
12	70			70
13	8			8
14	23			23
15	97			97
16			174	174
	1356	145	250	1751
Percent	77	8	14	

3.2 Stock market data

As a source for actual market data from the NYSE, we use the NYSE Trades and Quotes (TAQ) database for the period from January 1998 through June 1998. The TAQ data base includes historical trade prices and quantities, with their associated market conditions, transaction by transaction. The data are time-stamped to the nearest second. TAQ contains all equity transactions reported on the so called *Consolidated Tape*, which includes all transactions on NYSE, AMEX, Nasdaq and the regional exchanges.

For each of the stocks traded by the fund we search for data on this stock on the TAQ tape. In some cases we are not able to match the trade with TAQ data. Except for one date, the matching percentage is in the 82 to 94 percentage range. We also remove stocks that split around the fund’s

Fund stock portfolio.

¹³For the entire stock portfolio, the corresponding percentage is 83.

trades.

4 The size of the fund's transactions

Market microstructure theory and empirics tell us that an important determinant of the cost of trading is an order's size. Unless it is known that a trader is uninformed, larger orders will, all else equal, have a bigger price impact. We therefore give some statistics for the size of the fund's orders. We consider two measures of order size, one absolute and one relative.

4.1 Number of shares traded

A common definition of a *large* order is an order of ten thousand shares or more. Tables 2 and 3 give some summary statistics for the fund's trades, both in number of shares and in dollar values. Additionally, figure 1 shows the distribution of order sizes for the fund. As we observe, the majority of orders for the fund is below ten thousand shares, and most of them below twenty thousand.

Table 2 Number of shares traded per transaction.

Averages of number of shares traded per transaction.

No Shares	mean	std	min	q1	median	q3	max	n
All orders	6898	9654	11.0	2000	3800	7700	115200	4200
Crosses only	7013	9661	19.0	2000	3800	8000	109400	3494
Market only	6329	9598	11.0	2000	3550	6900	115200	706

Table 3 Dollar values per transaction

Averages of dollar values per transaction. Values in thousands of USD.

Trade value (in 1000\$)	mean	std	min	q1	median	q3	max	n
All orders	386	688	0.5	87	174	373	9050	4200
Crosses only	396	683	0.5	89	177	389	8962	3494
Market only	339	710	12.1	83	157	300	9050	706

We want to investigate whether there are any obvious differences in size between the orders that were crossed, and those that went to the market. Tables 2 and 3 give some summary statistics across the two groups, and figure 2 compares the distribution of order size across the two types of orders.

There seems to be little differences in sizes here. One thing to note is that some of the largest orders were done in the market. This may be a sign that it is problematic to "cross" very large volumes for one stock.

4.2 Relative volume

It is also interesting to look at how large the fund's trades were relative to the daily trading in the market for the same stock. We therefore calculate what fraction the fund's trades were of the total

Figure 1 The frequency distribution of order size(number of shares) for the fund.

The histogram shows the frequency distribution for the size (in number of shares) for the fund's equity trades. Data for all the trades, Jan to June 1998.





Figure 2 The distribution order size for the fund.

The histogram shows the frequency distribution for the size of each of the fund orders. On the left: Crosses, on the right: Market Orders.



quantities traded on the NYSE during a day. We look at both NYSE volumes on the day of the actual trade, and average daily NYSE volumes over the month of trade. Table 4 summarizes the relative trading volume against NYSE.

Table 4 Relative trading volume for the fund’s transactions.

Relative trading volume is defined as the percentage fraction of the fund’s daily trade relative to the daily total volume traded at the NYSE. We use two measures: The total volume traded at the NYSE that day for that particular stock, and the average daily NYSE trading volume during that month for that stock. Numbers in percent. “Mean” is an equally weighted average and “std” is its standard deviation. “Vw” is a value weighted average, using the value of the fund’s trades as weight.

	Relative to	mean	std	vw	min	median	max	n
All orders	That day	1.4	2.67	1.7	0.00	0.9	100	3971
	That month	1.2	2.14	1.8	0.0	0.8	73.6	3972
Crosses	That day	1.3	2.56	1.6	0.00	0.9	100	3308
	That month	1.1	1.29	1.4	0.0	0.8	25.2	3308
Market	That day	1.8	3.11	2.0	0.01	1.0	36	663
	That month	1.8	4.31	4.0	0.0	0.9	73.6	664

The median relative volume is 0.8% of the total trading during the day. This is probably a better measure than the mean, because the distribution is skewed to the right. This indicates that the fund’s average trade was relatively modest compared to the daily NYSE trading.

The average trade is higher compared to trade “that day” than compared to average trading during “that month,” suggesting that the fund did not in general trade on dates when trading activity in the market was peaking.

We also calculated the equally and value weighted averages for each trading date. Varying from one to two percent of daily transaction volume, these confirm the impression that the magnitude of the fund’s transactions have been modest.


The column max shows the highest relative volume over all trading dates. As we see there is at least one stock in which the fund was the whole market. It is truly a special case, this was an order of 700 shares in a thinly traded equity. The second highest relative volume was 40% of the day’s trades. There is in fact only 10 cases where the fund’s trades represented above 20% of the market trading that day.

Because averages can hide a lot of information, it is also interesting to see the *distribution* of the relative volumes. Figure 3 shows a histogram of relative volume against NYSE compared to transactions done “that day”. In the histograms the large transactions discussed above are left out.

Clearly, by both metrics used above, the Petroleum fund is not a dominant player in the US equity markets. But the fund is neither an unimportant player, the median trade of \$177 thousand is for example in the same range as the trades used in Keim and Madhavan [1997], which used a sample of large institutional investors.

Figure 3 Histogram of relative volume. Comparing to NYSE trade “that day.”

Relative volume (in percent). The comparison is done relative to the volume traded on the NYSE that day. Data for all dates on which the fund traded.



5 Comparing the trading costs for the bank with empirical evidence

As a measure of implicit trading cost, Keim and Madhavan [1998] use

$$\text{implicit cost} = \frac{P^a}{P_d} - 1$$

where P^a is the average price of all the executed trades in the order and P_d is the benchmark price, where they use the closing price for the stock on the day before the decision to trade the stock.

Explicit trading costs are defined as

$$\text{explicit cost} = \frac{\text{Commission per share}}{P_d}.$$

Total trading costs is then defined as the sum of these two.

Table 5 is taken from their paper and gives results for applying these measures to a large sample of institutional orders.

Table 5 Measured trading costs for a sample of institutions, taken from table 4 of Keim and Madhavan [1998].

Average trading cost by trade size quartile for common stock trades for 21 insitutions for the period January 1991 to March 1993.

Implicit trading costs are defined as $P^a/P_d - 1$, where P^a is the average price of all the executed trades in the order and P_d is the closing price for the stock on the day before the decision to trade the stock. Explicit trading costs are defined as (Commision per share/ P_d). The sample is partitioned by trade size quartile defined as number of shares traded divided by total outstanding shares, with quartile cutoffs determined separately for buy and sell transactions. Costs are reported in percent. Standard errors are in parenthesis.

Trade Size	Exchange Listed Stocks				Nasdaq stocks		
	Quartile	Total	Implicit	Explicit	n	Total	n
Buyer-Initiated Trades							
Smallest	0.31	0.18	0.13	7,392	0.76	1,755	
	(0.02)	(0.02)	(0.00)		(0.06)		
2	0.36	0.19	0.17	6,577	1.01	2,571	
	(0.03)	(0.03)	(0.00)		(0.07)		
3	0.53	0.32	0.21	6,503	1.08	2,645	
	(0.04)	(0.04)	(0.00)		(0.09)		
Largest	0.90	0.65	0.25	5,570	1.80	3,577	
	(0.05)	(0.05)	(0.00)		(0.10)		
Seller-Initiated trades							
Smallest	0.33	0.15	0.18	5,736	0.29	696	
2	0.31	0.11	0.20	5,291	0.50	1,142	
3	0.38	0.17	0.21	4,766	0.71	1,666	
Largest	1.42	1.13	0.29	3,830	2.63	2,602	

We calculate the same measures of implicit and explicit trading costs for the fund's trades. The results are summarized in table 6.

Table 6 Average Trading Costs for the Norwegian Petroleum fund's transactions.

Average Trading Costs for the Norwegian Petroleum fund's transactions, January to June 1998. Each date at which the Petroleum fund transacted. Costs are measured following Keim and Madhavan [1998]: Implicit trading costs are defined as $P^a/P_d - 1$, where P^a is the average price of all the executed trades in the order and P_d is the closing price for the stock on the day before the trade. Explicit trading costs are defined as (Commission per share/ P_d).


		All	Implicit	Explicit	n
All trades	mean	0.12	0.09	0.03	3909
	stdev	(1.95)	(1.95)	(0.14)	
	vw avg	0.30	0.29	0.01	
Crosses only	mean	0.09	0.06	0.03	3252
	stdev	(2.01)	(2.01)	(0.15)	
	vw avg	0.27	0.27	0.01	
Market Orders only	mean	0.30	0.25	0.05	657
	stdev	(1.60)	(1.60)	(0.04)	
	vw avg	0.46	0.43	0.03	

As we observe from the table, the average cost of transactions for the Petroleum fund of 0.12% compares favorably to the sample of trades used by Keim and Madhavan [1998], where the average cost for the smallest stock quartile was 0.31%, and the cost for larger quartiles even higher.

We also look at the frequency distribution of the implicit costs, both for all trades and split into crosses and market orders. Figures 4, 5 and 6 show the distributions.

Figure 4 The distribution of implicit costs. All orders.

The histogram shows the frequency distribution for the implicit cost. All trades.



It should be mentioned here that the time periods are not the same, average costs of trading may have declined since the trades used in Keim and Madhavan [1998]. Note also the clear differences

Figure 5 The distribution of implicit costs. Market orders only.

The histogram shows the frequency distribution for the implicit cost. Only trades that were traded in the market.





Figure 6 The distribution of implicit costs. Crosses only

The histogram shows the frequency distribution for the implicit cost. Only trades that were crossed.



between orders that were crossed and orders that were executed in the market. One problem is that we do not properly measure the “implementation shortfall,” because we do not know the date at which the initial desire for each stock was. We will return to this issue.

To make further comparisons with Keim and Madhavan [1998], table 7 splits the trades into quartiles based on trade size. The split is not completely comparable to their table, because we use the dollar size of the order, not the fraction of company's outstanding equity that the order constituted, but because the fund is tracking a value weighted index, these should be highly correlated.

Table 7 Average Trading Costs for the Norwegian Petroleum fund's transactions, sorted by size.

Average Trading Costs for the Norwegian Petroleum fund's transactions, January to June 1998. Costs are measured following Keim and Madhavan [1998]: Implicit trading costs are defined as $P^a/P_d - 1$, where P^a is the average price of all the executed trades in the order and P_d is the closing price for the stock on the day before the trade. Explicit trading costs are defined as (Commission per share/ P_d). Size sorted quartiles.

Size quartile			All	Implicit	Explicit	n
Largest	All trades	mean	0.23	0.23	0.01	978
		stdev	(1.96)	(1.96)	(0.01)	
	Crosses only	mean	0.23	0.22	0.00	858
		stdev	(2.00)	(2.00)	(0.00)	
	Market Orders only	avg	0.28	0.25	0.03	120
		stdev	(1.68)	(1.68)	(0.02)	
2	All trades	mean	0.27	0.25	0.02	978
		stdev	(1.88)	(1.88)	(0.01)	
	Crosses only	mean	0.27	0.26	0.01	801
		stdev	(1.96)	(1.96)	(0.00)	
	Market Orders only	avg	0.24	0.20	0.04	177
		stdev	(1.49)	(1.49)	(0.02)	
3	All trades	mean	0.05	0.03	0.03	978
		stdev	(1.85)	(1.85)	(0.02)	
	Crosses only	mean	0.00	-0.02	0.02	800
		stdev	(1.91)	(1.91)	(0.00)	
	Market Orders only	avg	0.30	0.25	0.05	178
		stdev	(1.57)	(1.57)	(0.03)	
Smallest	All trades	mean	-0.06	-0.14	0.08	978
		stdev	(2.06)	(2.08)	(0.27)	
	Crosses only	mean	-0.16	-0.25	0.09	796
		stdev	(2.13)	(2.14)	(0.30)	
	Market Orders only	avg	0.38	0.32	0.06	182
		stdev	(1.67)	(1.67)	(0.05)	

To see if there are any obvious differences across the dates, table 8 reports estimates using the same cost definitions for each date at which the Petroleum fund traded. Note the large differences across dates, implicit costs range from -1.84% to 1.46% .

Table 8 Average Trading costs for the fund on each day with transactions.

Average Trading Costs for the Norwegian Petroleum fund's transactions, January to June 1998. Each date at which the Petroleum fund transacted. Costs are measured following Keim and Madhavan [1998]: Implicit trading costs are defined as $P^a/P_d - 1$, where P^a is the average price of all the executed trades in the order and P_d is the closing price for the stock on the day before the trade. Explicit trading costs are defined as $(\text{Commission per share}/P_d)$.

Date		All	Implicit	(std)	Explicit	n
1	Crosses only	1.46	1.43	(2.40)	0.03	389
2	Crosses only	0.79	0.76	(1.97)	0.02	392
3	Crosses only	-0.27	-0.29	(2.29)	0.02	299
4	All trades	-0.06	-0.08	(2.06)	0.02	163
	Crosses only	-0.18	-0.19	(2.17)	0.01	138
	Market Orders only	0.58	0.49	(1.13)	0.09	25
5	Crosses only	-0.54	-0.57	(1.51)	0.03	58
6	Crosses only	-0.79	-0.82	(1.71)	0.03	92
7	All trades	-0.03	-0.06	(1.59)	0.03	375
	Crosses only	0.01	0.01	(1.78)	0.01	107
	Market Orders only	-0.05	-0.09	(1.51)	0.04	268
8	Crosses only	0.65	0.63	(1.19)	0.02	559
9	Crosses only	0.04	0.00	(0.03)	0.04	80
10	Crosses only	0.22	0.19	(0.69)	0.02	5
11	Crosses only	-0.11	-0.13	(1.67)	0.02	406
12	Crosses only	-0.39	-0.43	(1.62)	0.04	344
13	Crosses only	-0.08	-0.11	(1.36)	0.03	34
14	Crosses only	-0.14	-0.16	(0.99)	0.02	42
15	Crosses only	-1.84	-1.94	(1.76)	0.10	307
16	Market Orders only	0.54	0.49	(1.64)	0.05	364

As we discussed before, we have problems in directly comparing costs of market orders and crosses, because market orders are the “left over” orders that could not be crossed. Hence, the relevant benchmark for comparison is not the closing price the day before the trade, but the closing price the day before the decision to trade was made, which in this case is 2 or 3 days before.

To perform this analysis we calculate the costs using three different benchmarks $P_d = P_{t-1}$, $P_d = P_{t-2}$ and $P_d = P_{t-3}$, where P_{t-1} , P_{t-2} and P_{t-3} are the closing price respectively 1, 2 and 3 days before the trade date. Table 9 shows the results. The interesting observation from this table is the sudden sign reversal of the implicit cost of market orders if we use the closing price 2 days before as the benchmark.

To further investigate this, table 10 breaks down these calculations by date.

Table 9 Average trading costs for the fund's transactions, using alternative benchmarks.

Average Trading Costs for the Norwegian Petroleum fund's transactions, January to June 1998. Each date at which the Petroleum fund transacted. Costs are measured following Keim and Madhavan [1998]: Implicit trading costs are defined as $P^a/P_d - 1$, where P^a is the average price of all the executed trades in the order and P_d is the benchmark for comparison. We use three different benchmarks P_{t-1} , P_{t-2} and P_{t-3} , the closing prices respectively 1, 2 and 3 days before the trade date. Explicit trading costs are defined as (Commission per share/ P_d).

		Explicit	P_{t-1}		P_{t-2}		P_{t-3}		n
			Implicit	All	Implicit	All	Implicit	All	
All trades	mean	0.03	0.09	0.12	0.20	0.24	0.29	0.32	3909
	stdev		(1.95)		(2.86)		(3.57)		
Crosses only	mean	0.03	0.06	0.09	0.33	0.36	0.44	0.48	3252
	stdev		(2.01)		(2.93)		(3.66)		
Market orders only	mean	0.05	0.25	0.30	-0.40	-0.36	-0.49	-0.44	657
	stdev		(1.60)		(2.41)		(2.92)		

Table 10 Average trading costs for the fund's transactions, using alternative benchmarks. Day by day.

Average Trading Costs for the Norwegian Petroleum fund's transactions, January to June 1998. Each date at which the Petroleum fund transacted. Costs are measured following Keim and Madhavan [1998]: Implicit trading costs are defined as $P^a/P_d - 1$, where P^a is the average price of all the executed trades in the order and P_d is the benchmark for comparison. We use three different benchmarks P_{t-1} , P_{t-2} and P_{t-3} , the closing prices respectively 1, 2 and 3 days before the trade date. Explicit trading costs are defined as (Commission per share/ P_d).

Date			Explicit	P_{t-1}		P_{t-2}		P_{t-3}		n
				Implicit	All	Implicit	All	Implicit	All	
1	Crosses only	mean	0.03	1.43	1.46	2.64	2.67	2.11	2.13	389
		stdev		(2.40)		(3.02)		(3.39)		
2	Crosses only	mean	0.02	0.76	0.79	1.94	1.96	2.67	2.69	392
		stdev		(1.97)		(2.97)		(3.52)		
3	Crosses only	mean	0.02	-0.29	-0.27	0.44	0.46	1.52	1.54	299
		stdev		(2.29)		(3.10)		(3.79)		
4	All trades	mean	0.02	-0.08	-0.06	-0.22	-0.20	0.26	0.29	163
		stdev		(2.06)		(2.51)		(2.91)		
4	Crosses only	mean	0.01	-0.19	-0.18	-0.37	-0.36	0.11	0.12	138
		stdev		(2.17)		(2.55)		(2.89)		
4	Market orders only	mean	0.09	0.49	0.58	0.59	0.68	1.13	1.22	25
		stdev		(1.13)		(2.10)		(2.81)		
5	Crosses only	mean	0.03	-0.57	-0.54	-0.00	0.03			58
		stdev		(1.51)		(1.75)				
6	Crosses only	mean	0.03	-0.82	-0.79	-1.25	-1.22	-0.75	-0.72	92
		stdev		(1.71)		(2.14)		(2.78)		
7	All trades	mean	0.03	-0.06	-0.03	0.77	0.80	0.78	0.81	375
		stdev		(1.59)		(2.22)		(2.76)		
7	Crosses only	mean	0.01	0.01	0.01	0.88	0.88	0.88	0.89	107
		stdev		(1.78)		(2.30)		(2.89)		
7	Market orders only	mean	0.04	-0.09	-0.05	0.73	0.77	0.74	0.78	268
		stdev		(1.51)		(2.19)		(2.70)		
8	Crosses only	mean	0.02	0.63	0.65	0.53	0.55	0.73	0.76	559
		stdev		(1.19)		(2.32)		(3.42)		
9	Crosses only	mean	0.04	0.00	0.04	0.47	0.51	0.26	0.31	80
		stdev		(0.03)		(1.36)		(2.03)		
10	Crosses only	mean	0.02	0.19	0.22	1.33	1.35	2.71	2.74	5
		stdev		(0.69)		(0.73)		(1.73)		
11	Crosses only	mean	0.02	-0.13	-0.11	0.07	0.09	0.66	0.68	406
		stdev		(1.67)		(2.26)		(2.71)		
12	Crosses only	mean	0.04	-0.43	-0.39	-1.15	-1.11	-1.28	-1.23	344
		stdev		(1.62)		(2.34)		(2.92)		
13	Crosses only	mean	0.03	-0.11	-0.08	-0.63	-0.60	-0.47	-0.43	34
		stdev		(1.36)		(2.51)		(2.94)		
14	Crosses only	mean	0.02	-0.16	-0.14	-0.81	-0.80	-1.24	-1.23	42
		stdev		(0.99)		(1.76)		(2.30)		
15	Crosses only	mean	0.10	-1.94	-1.84	-2.29	-2.19	-3.73	-3.63	307
		stdev		(1.76)		(2.57)		(3.17)		
16	Market orders only	mean	0.05	0.49	0.54	-1.30	-1.26	-1.50	-1.45	364
		stdev		(1.64)		(2.19)		(2.68)		

6 Further Research

Given the results we have found, there are a number of unresolved issues that we want to look into using our data material.

One of the most promising seem to be looking at the difference between trades that were crossed and those that could not be crossed and was sent to the market. One possibility for doing that is to simply regress the fund's costs on various determinants of those costs, such as the market cap, the stocks past and future return, the state of the market, and the order size, as done in Keim and Madhavan [1997] or Jones and Lipson [1999]. One can then look across the types of orders to investigate differences. This can also be done formally using choice theoretic methods, analysing this as a binary choice regression, whether the stock was crossed or not as function of determinants, like the ones above.

To measure the "opportunity loss" more carefully, we could gather data about the exact dates when the stock was first tried crossed. The price at this date should be used as a benchmark.

It is hard to measure the exact performance of a trading strategy without an alternative strategy to measure it against. We have described the performance of the fund's *actual* trades. However, without explicitly trying to evaluate the performance of *alternative strategies*, we can not make any statements that this is a good or bad method for submitting orders. To get a better idea of how well the crossing strategy has worked, we could try to model the price impact of the fund's transactions. Even without doing a full modeling of the price impact, we can try to perform a simple limit order check that does not demand construction of the limit order book. Suppose that we at the beginning of the day put a limit order to buy at the current bidprice. If, during the day, we see the same or larger quantity traded at a lower price, our limit order would have been "hit" (assuming no price impact). If we didn't get "hit", accept the ask price at the end of the day. The costs from this strategy may be compared to the achieved prices.

7 Conclusion

This paper has investigated the Norwegian Petroleum fund's transactions in the US equity market, with the emphasis on measuring trading costs using alternative trading venues.

We looked at absolute and relative volumes, and found that the Petroleum fund is most of the time a relatively small player in this particular market, which after all is the world's largest.

For trades that *were* crossed, the average implicit and explicit costs were lower than found in similar trades for results in the literature. We do however find that the orders that did *not* cross were special. Using as a benchmark the date of the desire to order, we find that they have an on average *negative* cost. This warrants further investigation.

References

- Nasser Arshadi. NYSE, NASDAQ, and alternative trading systems: An evaluation of the SEC's proposal toward a national market system. *Financial Markets, Institutions and Instruments*, 7 (5), November 1998.
- Lawrence Harris. The winners and losers of the zero-sum game: The origins of trading profits, price efficiency and market liquidity. Working Paper, School of Business Administration, University of Southern California, May 1993.
- Joel Hasbrouck, George Sofianos, and Deborah Sosebee. New York Stock Exchange systems and trading procedures. Working Paper 93-01, New York Stock Exchange, April 1993.
- Charles M Jones and Marc L Lipson. Execution costs of institutional equity orders. Working Paper PW-99-01, Columbia Business School, March 1999.
- Donald B Keim and Ananth Madhavan. The upstairs market for large-block trades: Analysis and measurement of price effects. *Review of Financial Studies*, 9:1–36, 1996.
- Donald B Keim and Ananth Madhavan. Transaction costs and investment performance: An analysis of institutional equity trades. *Journal of Financial Economics*, 46:265–292, 1997.
- Donald B Keim and Ananth Madhavan. The cost of institutional equity trades. *Financial Analyst Journal*, pages 50–69, July/August 1998.
- Charles M Lee. Market integration and price execution for NYSE-listed securities. *Journal of Finance*, 48(3):1009–10038, 1993.
- Ananth Madhavan, Matthew Richardson, and Mark Roomans. Why do security prices change? A transaction-level analysis of NYSE stocks. *Review of Financial Studies*, 10(4):1035–1064, 1997.
- Maureen O'Hara. *Market Microstructure Theory*. Blackwell, 1995.
- A Perold. The implementation shortfall: Paper versus reality. *Journal of Portfolio Management*, 14:4–9, 1988.
- Richard Roll. A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, 39:1127–1139, September 1984.
- J Treynor. What does it take to win the trading game? *Financial Analyst Journal*, 47(1), January/February 1981.