



ELSEVIER

European Economic Review 39 (1995) 1277–1301

EUROPEAN
ECONOMIC
REVIEW

A comparison of the cost of trading French shares on the Paris Bourse and on SEAQ International

Frank de Jong^{a,*}, Theo Nijman^a, Ailsa Röell^{b,c}

^a *Department of Econometrics, Tilburg University, 5000 LE Tilburg, The Netherlands*

^b *London School of Economics, London, UK*

^c *Université Libre de Bruxelles, Brussels, Belgium*

Received May 1993; final version received October 1994

Abstract

This paper analyses the cost of trading French shares on two exchanges, the Paris Bourse and London's SEAQ International. Using a large data set consisting of all quotes, limit orders and transactions for a two month period, it is shown that for small transactions the Paris Bourse has lower implicit transaction costs, measured by both the effective and quoted bid–ask spread. The market in London, however, is deeper and provides immediacy for much larger trades. Moreover, we find that the cost of trading is decreasing in trade size, rather than increasing over the range of trade sizes that we examine. This suggests that order processing costs are an important determinant of bid–ask spreads, since competing market microstructure theories (adverse selection, inventory control) predict bid–ask spreads increasing in trade size.

Keywords: Cost of trading shares; Paris Bourse; SEAQ International

JEL classification: G15

* Corresponding author. We would like to thank Henk Berkman, participants of the European Finance Association meetings in Lisbon and Copenhagen, seminar participants at Erasmus University Rotterdam and University of Limburg, two anonymous referees, and the editor for helpful comments. Financial support from SPES is gratefully acknowledged.

1. Introduction

The growing importance of London as an international stock market where shares from other European countries are traded, constitutes a major change in the structure of Europe's financial markets. In recent years, London's SEAQ International has attracted considerable trading volume from the continental exchanges. This increased competition from London has induced the domestic exchanges to modernise and adapt their trading systems. An example is the move towards fully automated trading systems in Spain and Italy. It seems natural to suppose that London has attracted large volume because trading costs are lower, particularly for large trade sizes. In this paper we investigate this conjecture empirically for French equities traded in both London and in Paris.

The stock trading systems in these two financial centres differ considerably: London is a quote-driven dealership market whereas Paris is a continuous auction. Theoretical work suggests that these differences in market architecture could have an impact on trading costs and the depth of the markets, see for example Madhavan (1992) and Pagano and Röell (1993). An investigation of the relative merits of the two trading systems is an important input for policy regarding market design and regulation. In this paper we use a large data set, a simultaneous record of all quotes, limit orders and transactions in both London and Paris, to compare the implicit cost of trading French shares on the Paris Bourse and on SEAQ International. The bid–ask spread is a major component of the total cost of trading, and we will provide several measures of the spread on both exchanges. First, the average quoted spread is estimated from the Paris limit order book and market makers' quotes in London. Second, the average effective spread is estimated using the difference between quotes and actual transactions prices. Estimates of the quoted and effective spread are presented for different transaction sizes. The dependence of the spread on trade size is of theoretical interest, because it can be used to assess the validity of market microstructure theories that predict that the bid–ask spread will be increasing in trade size.

Both the quoted and the effective spread are not directly observable in our data set. On the Paris Bourse part of a limit order can be hidden from the public information system, so that the limit order book seems less deep than it actually is. Uncorrected estimates would therefore overestimate the quoted spread in Paris. In London the problem is that there is some misreporting of transaction times, which causes a timing bias in our effective spread estimate. In order to circumvent these problems we also present model-based estimates of the average realised spread using transaction prices only. These estimators can be seen as refinements of Roll's (1984) estimator.

The setup of the paper is as follows. In Section 2 we briefly discuss the major theories that explain the existence and the size of the bid–ask spread. In Section 3, we describe the trading systems on the Paris Bourse and on SEAQ International. In Section 4 we describe our data. The spread estimates are presented in Sections

5, 6 and 7. In Section 5 we compute the average quoted spread and in Section 6 the average effective spread, both in Paris and in London. In Section 7 we take a model-based approach to estimating the realised spread that uses transactions data only. Finally, we summarise the main conclusions in Section 8.

2. Theories of the bid–ask spread

In the literature on stock market microstructure there are a number of theories that explain the bid–ask spread. Most theories view the spread as a compensation for the services of a market maker, who takes the other side of all transactions. In the literature, e.g. Stoll (1989), three cost components are distinguished: order processing cost (including dealer oligopoly profit), inventory control cost and adverse selection cost. In this section, these three components will be discussed in more detail.

The order processing cost component reflects the cost of being in the market and handling the transaction. To compensate for these costs, the market maker levies a fee on all transactions by differentiating between buy and sell prices. Much of the empirical literature, such as Madhavan and Smidt (1991) and Glosten and Harris (1988), assumes that this fee is a fixed amount per share. However, it seems more natural to suppose that order processing cost is largely fixed *per transaction*, so that expressed as cost per share it should be inversely related to trade size.

A second type of cost for the market maker is the cost of inventory management. For example, a purchase of shares will raise the market maker's inventory above a desired level. The market maker runs the risk of price fluctuations on his inventory holdings and if he is risk averse he will demand a compensation for this risk. This intuition is formalised in the model of Ho and Stoll (1981), who show that the inventory control cost is an increasing function of trade size and share price volatility.

The third type of cost for the market maker arises in the presence of asymmetric information between the market maker and his potential counterparties in trading. This theory was first proposed by 'Bagehot' (1971) and formalised in the models of Glosten and Milgrom (1985) and Kyle (1985). A trader with superior private information about the underlying value of the shares will try to buy or sell a large number of shares to reap the profits of this knowledge. The market maker, who is obliged to trade at the quoted prices, incurs a loss on transactions with better informed counterparties. To compensate for this loss he will charge a fee on every transaction, so that expected losses on trades with informed traders are compensated by expected profits on transactions with uninformed 'noise' traders. Because the informed parties would tend to trade a large quantity in order to maximise the profits from trading on superior information, the adverse selection effect is related to trade size: large transactions are more likely to be initiated by

better informed traders than small transactions, as in the model of Easley and O'Hara (1987). Therefore, the asymmetric information cost is an increasing function of trade size, and the market maker's quotes for large transactions will be less favourable than the quotes for small sizes.

These theories have been developed for markets with competitive designated market makers. In Paris, we may regard the issuers of public limit orders as market makers because they provide liquidity to the market and run the risk that their limit order will be executed against a market order placed by somebody with superior information. The inventory control theory is applicable to the extent that we can regard those who place market orders as demanders of immediacy, while those who place limit orders are making the market by absorbing inventories in return for a price concession. In practice, the distinction between the two groups is not sharp, as any trader can place both types of orders.

3. Description of the markets in French equities

In this section we describe the trading systems on the major exchanges where French equities are traded: the Bourse in Paris and SEAQ International in London. Because the trading systems are so different – Paris is a continuous auction market whereas London is a dealership market – we devote two separate sub-sections to this description.

3.1. The trading system on the Paris Bourse

The Paris Bourse uses a centralised electronic system for displaying and processing orders, the Cotation Assistée en Continu (CAC) system. This system, based on the Toronto Stock Exchange's CATS (Computer Assisted Trading System), was first implemented in Paris in 1986. Since then, trading in nearly all securities has been transferred from the floor of the exchange onto the CAC system. All the most actively traded French equities are traded on a monthly settlement basis in round lots of 5 to 100 shares set by the Société des Bourses Françaises (SBF) to reflect their unit price. The SBF itself acts as a clearing house for buyers and sellers, providing guarantees against counterparty default.

Every morning at 10 a.m. the trading day opens with a batch auction where all eligible orders are filled at a common market clearing price. Nowadays the batch auction is relatively unimportant, accounting for no more than 10 to 15% of trading volume. Its role is to establish an equilibrium price before continuous trading starts. Continuous trading takes place from 10 a.m. to 5 p.m.

In the continuous trading session there are two types of orders possible, limit orders and market orders. Limit orders specify the quantity to be bought or sold, a required price and a date for automatic withdrawal if not executed by then, unless

Table 1
Simplified trading screen of CAC system

Bid			Ask			Transactions		
No. ^a	Shares	Price	Price	Shares	No.	Shares	Price	Time
1	200	763	770	800	3	400	765	10:08
1	500	762	774	100	1	50	765	10:08
1	400	761	775	200	1	50	770	10:06
4	450	760	778	1000	1	50	770	10:02
1	50	754	779	100	1	100	768	10:02

^a No. denotes the number of limit orders involved.

the limit order is good till cancelled ('à révocation'). Limit orders cannot be issued at arbitrary prices because there is a minimum 'tick' size of FF 0.1 for stock prices below FF 500, and FF 1 for higher prices. More than one limit order may be issued at the same price. To these orders, strict time priority for execution applies.

After the opening, traders linked up to the CAC system will see an on-screen display of the 'market by price' as depicted in Table 1. For both the bid side and the ask side of the market, the five best limit order prices are displayed together with the quantity of shares available at that price and the number of individual orders involved. The difference between the best bid and ask price is known as the 'fourchette'. Brokers can scroll down to further pages of the screen to view limit orders available beyond the five best prices. In addition, some information concerning the recent history of trading is given: time, price, quantity and buyer and seller identification codes for the five last transactions, the cumulative quantity and value of all transactions since the opening, and the price change from the previous day's close to the latest transaction.

In practice, the underlying limit order book tends to be somewhat deeper than suggested by the visible display of limit orders. This is because traders who are afraid that they might move the market by displaying a very large order may choose to display only part of their limit order on-screen. The remaining part, known as the 'quantité cachée' or undisclosed quantity, remains invisible on-screen but may be called upon to fill incoming orders as the visible limit orders become exhausted. Strict price priority applies also to the hidden orders, but not time priority. Röell (1992) suggests that due to the *quantité cachée* the visible depth of the market is about two thirds of the actual depth when hidden quantities are included.

Market orders only specify the quantity to be traded and are executed immediately 'au prix du marché', i.e. at the best price available. If the total quantity of the limit orders at this best price do not suffice to fill the whole market order, the remaining part of the market order is transformed into a limit order at the transaction price (for a detailed description of this system see Biais et al. (1992)).

Hence, market orders do not automatically walk up the limit order book, and do not always provide immediate execution of the whole order ¹.

The member firms of the Bourse (the ‘Sociétés de Bourse’) key orders directly into the CAC system via a local terminal. All market participants can contribute to liquidity by putting limit orders on display. In particular, the Sociétés de Bourse may act in dual capacity: as agency brokers, acting on behalf of clients, and as principals, trading on own account. Their capital adequacy is regulated and monitored by the Bourse.

There is some scope for negotiated deals if the limit order book is insufficiently deep. A financial intermediary can negotiate a deal directly with a client at a price lying within the current fourchette, provided that the deal is reported to the CAC system as a ‘cross order’. For trades at prices outside the fourchette, the member firm acting as a principal is obliged to fill all central market limit orders displaying a better price than the negotiated price within five minutes.

3.2. *SEAQ International*

SEAQ International is the price collection and display system for foreign equity securities operated by London’s Stock Exchange. For each foreign equity included in SEAQ International, the system provides an electronic display of bid and ask prices quoted by the market makers registered for that equity.

The French equities in our sample are designated as firm quote securities, which means that during the relevant mandatory quote period (9:30 to 16:00 London time, i.e. 10:30 to 17:00 Paris time in our sample) the registered market makers are obliged to display firm bid and ask prices for no less than the ‘minimum marketable quantity’, also referred to as the Normal Market Size (NMS), a dealing size set by the exchange’s Council at about the median transaction size. Market makers are obliged to buy and sell up to that quantity at no worse than their quoted prices. In addition, when a market maker displays a larger quantity of shares than the minimum marketable quantity, his prices must be firm for that quantity. Outside the mandatory quote period, market makers may continue to display prices and quantities under the same rules regarding firmness of prices.

SEAQ International market makers are not allowed to display prices on competing display systems which are better than those displayed on SEAQ International. Market making in French shares is fairly competitive, see Röell (1992): during our sample period, most French equities were covered by at least ten market makers, and usually many more.

¹ A trader who wants to trade a certain quantity immediately can circumvent this mechanism by placing a *limit order* at a very unfavourable price. This limit order will then be executed against existing orders on the other side of the market that show a more favourable price.

4. Description of the data

The data consist of a comprehensive record of quote changes and transactions in the French equities of our sample, collected over a two month period in the summer of 1991 by the Paris and London stock exchanges.

The Paris data set is a transcription of all changes in the trading screen information for all shares on the CAC system for 44 trading days in the summer of 1991, starting May 25 and ending July 25. We have available a complete record of the total limit order quantity at the five best prices on both the bid side and the ask side of the market and all transactions. This enables us to reconstruct at every point in time the visible limit order book for each security in the sample, up to the cumulative volume of the observed best limit orders. However, we do not observe the ‘quantité cachée’, so the actual limit order book might be deeper than the observed quantities suggest. Due to the automated trading system, the data are relatively clean. The time stamps indicate exactly the time of the transaction or quote change. Also, quote and trade information is in correct sequence, so that it is possible to infer exactly whether a trade is buyer- or seller-initiated.

For each transaction an indicator records whether the transaction is a ‘cross’ negotiated outside the CAC system. Cross transactions need not be reported immediately to the exchange, so that their timing may not be totally accurate. We also have available broker identification codes of the buying and selling parties, which allow us to identify series of small transactions that were initiated by the same person as part of one large transaction. The transaction price per share for such transactions is defined as the quantity weighted average of the prices of the small transactions that together make up the larger one.

In this paper we concentrate on ten major French stocks, listed and described in Table 2. Panel A concerns the Paris data. For most series there are between five and ten thousand transactions in the data set. Excluding cross transactions, the median transaction value is between FF 50,000 and FF 150,000 (\$5,000–\$15,000 at the time). The distribution of transaction size is very skewed: the mean is about twice the median, indicating that a few large transactions account for a large share of total turnover. The cross transactions are relatively large: their median value is about 2 to 5 times as large as the median value of regular transactions, and the mean value is up to 10 times the mean value of regular transactions. Although there are relatively few crosses (between 2 and 5% of the total number of transactions) they account for a large share of total trading volume.

The data from the London exchange cover the months May to July 1991. First, there is a chronological record of all the market maker quotes as displayed on the SEAQ International system: the name of the market maker, his bid and ask quotes and the sizes for which they hold good. Typically, there are about 10 to 15 market makers in each security; many of them are international security houses, see Röell (1992). Their quotes are firm for sizes that can range from NMS up to about 10 times NMS. Market makers do not update their quotes very frequently: on a

Table 2
Descriptive statistics of transactions data ^a

A. Paris		CAC				Crosses		
Firm	Full name	Average price	Median value	Mean value	nobs	Median value	Mean value	nobs
AC	Accor	771	114	197	5255	384	2531	148
AQ	Elf-Aquitaine	358	179	303	9855	183	1607	598
BN	BSN	889	62	182	10728	266	1039	378
CA	Carrefour	1919	90	164	9943	366	1268	307
CS	Axa-Midi	989	62	120	6482	89	2266	221
EX	Generale des Eaux	2518	129	247	9585	366	2070	475
OR	l'Oreal	584	64	145	6813	116	694	271
RI	Pernod-Ricard	1162	84	131	3626	327	838	123
SE	Schneider	685	68	134	4329	388	876	183
UAP	Un. Ass. de Paris	538	134	222	5206	54	728	402

B. London					
Firm	Median value	Mean value	nobs	NMS ^b	
AC	1094	2049	393	2000	
AQ	1473	2966	1168	5000	
BN	862	1487	853	2500	
CA	950	2293	771	500	
CS	758	1858	291	1000	
EX	1106	2545	905	500	
OR	732	1691	449	2500	
RI	630	1479	210	1000	
SE	1100	1970	204	2000	
UAP	1532	2460	518	2000	

C. Percentiles of transaction size distribution ^c								
%	Paris				London			
	90	95	99	99.5	90	95	99	99.5
AC	0.25	0.50	1.0	1.4	2.5	4.4	12.4	15.0
AQ	0.40	0.60	2.0	3.0	3.0	4.6	15.7	22.0
BN	0.21	0.36	0.8	1.2	1.6	2.4	5.0	6.0
CA	0.40	0.60	1.6	2.1	5.0	8.0	30.0	40.0
CS	0.29	0.48	1.1	3.0	4.0	6.0	13.3	21.6
EX	0.46	0.80	2.0	4.0	4.2	7.5	18.0	26.7
OR	0.22	0.39	1.0	1.5	2.2	3.4	9.8	18.8
RI	0.25	0.43	1.2	1.8	3.3	5.1	8.8	8.8
SE	0.25	0.48	1.0	1.5	3.8	5.2	9.7	10.0
UAP	0.20	0.75	2.5	4.6	5.0	7.5	12.5	15.5

^a Price is average transaction price in FF; value of transactions in FF1000; nobs is number of observations.

^b See Panel A. NMS is Normal Market Size in number of shares.

^c Percentiles expressed in NMS; crosses included in Paris sample.

typical day their opening quotes are not changed more than once or twice, though occasionally there are eventful days where quote changes are much more frequent.

Second, there is a record of transactions: date, time, price and size, as reported to the stock exchange. The data set does not tell us who initiated the transaction, or which side is taken by a market maker.²

Table 2, panel B shows some statistics for the London data. There are fewer transactions in London than in Paris, but the median size of the transactions is much larger. The NMS is generally valued at about FF 1 million (\$100,000), a rather large transaction by Paris standards. The average value of transactions in London is about 10 times the average value of regular transactions in Paris, and still somewhat larger than the mean value of crosses in Paris.

Table 2, panel C shows some numbers concerning the distribution of the trade size. There are many more large transactions in London than in Paris. For example, the 90th percentile in London is about as large as the 99.5th percentile in Paris, where the latter includes the cross transactions. We also computed patterns of the number of trades and the distribution of volume by time of day. These show a clear U-shaped pattern, as in McNish and Wood (1990). For more details we refer to the working paper version of this paper, De Jong et al. (1993).

5. The quoted spread

In this section we provide an analysis of the cost of immediacy on the Paris Bourse and SEAQ International. The worst price that can be obtained in an urgent transaction is determined by the limit order book in Paris and the market makers' quotes in London. Thus we measure the cost of immediacy by the quoted spread. For Paris, the average quoted spread is determined as the average difference between bid and ask prices in the limit order book for a certain size. In London, the quoted spread is the difference between the best bid and ask quotes of the market makers. Although prices are negotiable in London, one cannot always count on 'within-the-touch' prices for an immediate transaction.

² In our estimates we use the common classification rule of attributing the initiation to the side of the bargain which gets a price worse than the reigning mid-quote; and we attempt to correct for potential biases induced by mis-classification. Indeed, transactions can be both customer-initiated and inter-dealer trades. It is usual for a large deal to be taken on initially by a large market maker, who subsequently passes on parts of it to final holders or even other market makers in the stock. Thus, a series of transactions is recorded; and indeed, some of the unwinding may take place on the Paris Bourse. Lack of data on the identities of traders precludes us from identifying such follow-on transactions. The reader should be aware that this may inflate the number of transactions recorded for London relative to those in Paris, where trades are more likely to involve final customers (though even there intermediaries take on large negotiated positions which they may want to unwind subsequently via the limit order book). And our measures of transaction cost will necessarily include the cost to the first market maker of unwinding his position in the subsequent transactions.

In order to compute the quoted spread in Paris it is necessary to construct the limit order book. We observe all new limit orders, as well as all transactions that fill limit orders and orders that are withdrawn, so that we can recursively build up the order book over the day. There are two problems in constructing the order book, however. First, there is the unobserved ‘quantité cachée’, which makes the book deeper than observed. Second, we observe only the limit orders at the five best prices, so that we do not have prices for larger order sizes. In constructing the book we impute the fifth best limit order price for all sizes beyond the range for which the bid and ask price are observed.³ A first way to measure the average quoted spread would be by a simple calendar time average of the observed spreads between bid and ask prices. An obvious drawback of that spread measure is that periods in which there is hardly any trading are given the same weight as periods of equal length in which trading is heavy. An estimator that conditions on the actually observed trade pattern is the transaction time average of the difference between bid and ask prices, see De Jong et al. (1993). A further refinement of that estimator is obtained if we condition not only on the pattern of trades over the day, but also on the size of transactions. The results of Biais et al. (1992) suggest that indeed large transactions tend to take place at times when it is relatively cheap to trade large quantities. This is formalised in our preferred estimator, which averages the quoted spread over times that transactions in a particular size class occurred:

$$S_Q(z, \bar{z}) = \frac{\sum_{i=1}^N I(z < z_i \leq \bar{z})(A[t_i, z_i] - B[t_i, z_i])}{\sum_{i=1}^N I(z < z_i \leq \bar{z})}, \quad (1)$$

where $A[t_i, z_i]$ denotes the ask price of a transaction at time t_i of size z_i , $B[t_i, z_i]$ the corresponding bid price and $I(\cdot)$ is an indicator function that takes the value one if the trade size exceeds the lower bound z and is smaller than or equal to the upper bound, \bar{z} , and takes the value zero otherwise. Table 3 reports the quoted spread S_Q for several size classes. The quoted spread is clearly increasing in trade size, nearly doubling from the smallest to the largest size class. For London, the ‘touch’ was averaged by transaction size class, and shows no clear pattern. In London, therefore, trade size does not seem to depend on the ‘touch’.

³ An alternative procedure is to exclude those observations for which we do not observe the quoted bid and ask price up to the required size. That procedure introduces a selection bias in the spread measure because the five best limit orders add up to a large size only when the market is deep. Hence, that procedure underestimates the spread. Comparison of this alternative procedure with the procedure described in the main text showed that the selection bias is more serious than the bias caused by imputing the fifth best price for unobserved limit orders. See also Anderson and Tychon (1993) who report large selection biases for Belgian stocks. Clearly this biases the average quoted spread downwards. On the other hand, ignoring the quantité cachée biases the average upwards. The net effect of these data imperfections on the estimates of the quoted spread is indeterminate.

Table 3

Transaction time average of percentage quoted spread by size class

A. Paris ^a							
Size:	≤ 0.1	0.1-0.5	0.5-1	> 1.0	All		
AC	0.237	0.271	0.472	0.523	0.252		
AQ	0.179	0.218	0.324	0.363	0.201		
BN	0.182	0.228	0.422	0.555	0.197		
CA	0.209	0.234	0.336	0.375	0.225		
CS	0.356	0.434	0.697	0.720	0.389		
EX	0.134	0.154	0.225	0.263	0.148		
OR	0.309	0.396	0.706	0.782	0.342		
RI	0.359	0.404	0.662	0.778	0.378		
SE	0.356	0.449	0.831	0.993	0.386		
UAP	0.421	0.465	0.716	0.925	0.453		
B. London ^a							
Size:	≤ 0.1	0.1-0.5	0.5-1	1-2	2-5	> 5	All
AC	1.325	1.346	1.336	1.241	1.301	1.345	1.315
AQ	0.995	0.953	0.961	0.960	0.897	0.972	0.954
BN	0.880	0.852	0.845	0.830	0.812	0.849	0.852
CA	1.260	1.245	1.276	1.219	1.160	1.183	1.228
CS	2.418	2.206	2.285	2.132	2.054	2.166	2.208
EX	0.957	0.974	1.061	1.008	0.996	1.075	1.006
OR	1.505	1.652	1.625	1.656	1.680	1.716	1.624
RI	2.208	2.130	2.253	2.201	2.095	1.909	2.159
SE	2.003	2.072	1.954	1.972	2.158	1.922	2.025
UAP	1.841	1.655	1.740	1.681	1.706	1.488	1.685
C. Percentage imputed values in Paris limit order book ^b							
Size:	≤ 0.1	0.1-0.5	0.5-1	> 1.0	All		
AC	0.0	0.0	6	10	0.1		
AQ	0.0	0.0	14	28	0.7		
BN	0.0	0.0	9	40	0.1		
CA	0.0	0.3	8	19	0.5		
CS	0.0	0.6	11	20	0.4		
EX	0.0	0.1	4	10	0.3		
OR	0.0	0.1	13	40	0.2		
RI	0.0	0.4	15	38	0.4		
SE	0.0	0.5	32	75	0.6		
UAP	0.0	0.0	6	11	0.3		

^a Quoted spread by $S_D(\underline{z}, \bar{z})$ definition as a percentage of transaction prices.

^b This table reports the percentage of transactions for which either in the bid or the ask price was constructed by imputing limit order prices if the limit order book contained too few orders. For details see Section 5.

A comparison of both markets shows that the quoted spread in Paris is much smaller than the quoted spread in London for all transaction sizes below NMS. However, for larger transactions the quoted spread in Paris rises quickly as the

limit order book runs out.⁴ Some care has to be taken with these results because the estimates of the quoted spread in Paris ignore the hidden quantities and are marred by the problem that we only have data on the five best limit orders. The direction of the overall bias caused by these problems is not clear.

6. Effective spread

In this section we compute spread estimates that are based on the difference between quotes and actual transaction prices and will therefore be referred to as measures of the effective spread. The estimator of the effective spread that we propose is twice the average absolute difference between the quoted mid-price and the transaction price:

$$S_E(z, \bar{z}) = \frac{\sum_{i=1}^N I(z < z_i \leq \bar{z}) \cdot |p[i] - m[i]|}{\sum_{i=1}^N I(z < z_i \leq \bar{z})} \quad (2)$$

where as before $I(\cdot)$ is the indicator function, $p[i]$ is the actual transaction price (average price paid per share) and $m[i]$ is the mid-price at the time of the i th transaction, defined as the average of the best bid and ask quote (or best buy and sell limit orders) for the smallest possible order size.

In practice, the market mid-quote may temporarily deviate from the security's 'true' equilibrium value in response to market makers' and other speculators' inventories. Other agents who are aware of this can obtain lower (or even negative) transaction costs, because they can place market orders to buy (sell) when quotes are low (high) relative to the true value. Our spread measure does not take account of this. Thus, it does not try to measure trading costs for the actual population of market order placers, some of whom may well be market making in this way. Rather, our spread measures the trading cost for an agent whose only source of information regarding the security's value is the display of price quotes on the exchange.

For Paris, there are at least two important differences between the effective spread measure and the quoted spread measures of the previous section. The first is that the limit order book data are used only to construct the mid-price. This means that the effective spread estimate in Paris is not affected by the *quantité cachée* and the availability of only the five best limit order prices. The second important difference is that the implicit assumption that the market is equally deep on both sides is dropped. One would expect that trades are more likely to take place on the deeper side of the market. If so, the effective spread measure should be lower than the quoted spread measure for larger trade sizes. See also Biais et al. (1992) on this point. In London transactions are routinely priced within the touch,

⁴ In Table 3, Panel C the percentage of imputed limit order prices is reported.

and therefore the quoted spread will be an overestimate of the realised cost of trading. Surprisingly, calculations show that about 50% of the transaction prices are outside the touch. This is also true for small transactions below NMS, for which the quotes are binding. This is confirmed by estimates made by the London Stock Exchange (1992b). A likely explanation for this phenomenon are errors in the reported time of transaction.

Is the effective spread measure a good indicator of the cost of immediacy? No, for two reasons. First of all, an impatient trader cannot choose the deeper side of the market. Secondly, in London not all traders are able to negotiate within-the-touch prices, depending on how urgent their need to transact is and how sure the market maker can be that their trade is not information driven.

Estimates of the average effective spread are reported in Table 4. In calculating the estimates we excluded all transactions outside the continuous trading period in Paris or the mandatory quote period in London because outside normal trading hours the mid-quote is not a reliable proxy for the market consensus valuation of the stock.

Table 4, panel A shows the average effective spread in Paris. All transactions within the continuous trading period were used, including ‘crosses’. The table clearly shows that the effective spread in Paris does not increase with trade size. In contrast, in the previous section we have seen that the quoted spread increases with size. The dependence of the quoted and effective spread in Paris on trade size is illustrated in Fig. 1, where the quoted spread and the effective spread estimate for the Accor series are graphed.

Estimates of the average effective spread in London are reported in panel B. The most striking result here is that the effective spread in London seems to be *declining* in trade size. This effect was also observed by Breedon (1993), Tonks and Snell (1992) and Röell (1992). A comparison of Table 4 and Table 3 shows that in London the average effective spread for transactions smaller than NMS is sometimes larger than the quoted spread. This seems impossible: the rules of SEAQ International oblige market makers to stand firm at the best quoted price for transactions smaller than NMS. A likely explanation for this anomaly is a *timing bias* due to inaccurately reported times of transactions.⁵

⁵ Another possible explanation is inaccurate or late updating of quotes. We checked this possibility by using in Eq. (2) the most recent mid-price quoted in Paris rather than in London. This should alleviate the problem because quote (limit order) updates in Paris are much more frequent than quote changes in London. The estimated effective spreads using Paris mid-prices were very similar to the estimates using London quotes. Therefore, it is errors in the reported transaction time that appear responsible for the problem. In the appendix we show that this biases our effective spread measure upwards, because the market mid-price may have moved between the actual transaction time and the reported time. In Table 4, panel C we report bias-adjusted effective spreads for London. These are for some series (four out of ten) substantially smaller than the uncorrected spreads.

Table 4
Average percentage effective spread

A. Paris ^a								
Size:	≤ 0.1	0.1-0.5	0.5-1	> 1.0	All			
AC	0.245	0.236	0.251	0.230	0.242			
AQ	0.193	0.202	0.188	0.160	0.196			
BN	0.187	0.188	0.201	0.181	0.187			
CA	0.227	0.212	0.221	0.200	0.221			
CS	0.372	0.378	0.429	0.327	0.375			
EX	0.151	0.154	0.158	0.145	0.153			
OR	0.325	0.315	0.305	0.171	0.322			
RI	0.368	0.352	0.384	0.401	0.364			
SE	0.362	0.359	0.311	0.178	0.361			
UAP	0.458	0.416	0.434	0.381	0.438			
B. London ^b								
Size:	≤ 0.1	0.1-0.5	0.5-1	1-2	2-5	> 5	All	
AC	1.214	1.133	1.051	1.334	1.487	0.703	1.181	
AQ	1.823	1.196	1.354	1.016	1.049	1.553	1.257	
BN	1.520	1.009	1.087	0.992	1.355	1.538	1.151	
CA	1.588	1.541	1.153	1.181	1.124	1.195	1.294	
CS	3.085	1.494	1.377	1.263	1.911	2.200	1.679	
EX	1.448	1.101	1.045	1.063	1.159	1.446	1.150	
OR	2.126	1.163	1.540	1.420	1.956	1.688	1.469	
RI	1.869	1.227	1.075	1.648	1.694	1.288	1.398	
SE	1.374	1.322	1.316	1.237	0.930	1.008	1.250	
UAP	1.664	1.258	1.140	1.452	1.327	1.221	1.297	
C. London, bias corrected								
Size:	$S_Q(0, 1)$	$S_E(0, 1)$	0-1	1-2	2-5	> 5	All	σ
AC	1.364	1.128	1.128	1.334	1.487	0.703	1.181	0.00
AQ	0.833	1.508	0.833	0.573	x	0.378	0.163	1.69
BN	0.852	1.153	0.852	0.489	1.176	1.422	0.849	1.14
CA	1.272	1.403	1.272	0.949	0.852	0.972	1.121	1.09
CS	2.260	1.722	1.722	1.263	1.911	2.200	1.679	0.00
EX	0.977	1.146	1.977	0.849	0.966	1.371	0.984	0.99
OR	1.607	1.420	1.420	1.420	1.956	1.688	1.469	0.00
RI	2.194	1.331	1.331	1.648	1.694	1.288	1.398	0.00
SE	2.015	1.328	1.328	1.237	0.930	1.008	1.250	0.00
UAP	1.709	1.237	1.237	1.452	1.327	1.221	1.297	0.00

^a Average based on all transactions in continuous trading period (10-17).

^b Transactions only in mandatory quote period (9.30-16).

^c See Panel B. Bias correction described in appendix. x = true spread could not be computed. σ is the estimated standard deviation of the pricing error.

The observation that effective spreads do not increase with trade size is important because it is not in line with the inventory control and adverse selection models of the spread discussed in Section 2, or with the assumption that order




Fig. 1. Quoted and effective spread Accor in Paris by size.

processing cost is fixed per share. Constant processing costs per transaction, and therefore declining per share, could be an explanation for the empirical result that the cost *per share* is smaller for large trade sizes. We return to this issue in Section 7 where we estimate a parametric model for the dependence of the effective spread on trade size.

Comparing the effective spread in London with the effective spread in Paris, it appears that the London effective spread is considerably higher than the Paris one, so Paris seems to be cheaper. There are a number of caveats. First, there are few transactions larger than NMS in Paris (indeed most of those are cross transactions) while in London roughly half the transactions exceed NMS. Second, the full cost of trading also includes taxes and other explicit transaction costs. We return to this point in the concluding section.

7. Model-based estimates of the realised bid–ask spread

The spread measures of the previous section relied on data from the limit order book or quotes to construct an estimate of the unobserved consensus value of the stock. The estimators proposed in this section do not require such a proxy, and are therefore less sensitive to the problems encountered in Section 6. In particular, timing bias is not a problem. The price paid for this improvement is the need to make some parametric assumptions about the process that generates prices. We build some simple models to estimate the realised spread and to estimate the dependence of the spread on trade size.

The simplest model that we consider is based on the work of Stoll (1989) and George et al. (1991). In their models it is assumed that the transaction price, p_t , is

equal to the mid-price prior to the trade, y_t , plus or minus one-half times the total spread, S . We allow for an error term, u_t , in the price equation, that picks up various effects on the transactions price that are not captured by the mid-price and the transactions type, such as price discreteness and trade size. Thus, the price equation is

$$p_t = y_t + (S/2)Q_t + u_t. \quad (3)$$

where Q_t indicates whether the transaction is initiated by the buyer (+1) or the seller (-1), and will be referred to as the 'sign' of the trade. In order to obtain an equation in transaction prices only, we make some assumptions on the dynamics of the mid-price.

If there is asymmetric information between market makers and other traders, the market maker will revise his mid-price after a trade has occurred. Moreover, for inventory control reasons he will also change his quotes because the trade changes his inventory. Let $(1 - \pi)$ be the fraction of the spread that persists in subsequent transaction prices. In addition, between two trades public information on the stock's value may come in, captured by a constant, β_0 , and an error term, e_{t+1} , so that the evolution of the mid-quote is given by

$$y_{t+1} = y_t + \beta_0 + (1 - \pi)(S/2)Q_t + e_{t+1} \quad (4)$$

Substituting the transaction price for the mid-quote using (3), the following equation holds:

$$\Delta p_t = \beta_0 + (S/2)Q_t - \pi(S/2)Q_{t-1} + \xi_t, \quad \xi_t = e_t + \Delta u_t. \quad (5)$$

This is a valid regression model under the additional assumption that Q_t and ξ_t are uncorrelated. If the pricing error u_t is due to rounding, Q_t and ξ_t might be correlated. Also, Q_t and e_t (and hence ξ_t) might be correlated if good (bad) public news tends to induce a rush of buyers (sellers). But one would expect speculators to realign their quotes in response to public information in such a way that the expected trading interest on the two sides of the market remains roughly balanced. We think this assumption is reasonable and therefore we treat Q_t and the error as uncorrelated and estimate all regressions by ordinary least squares.

If (e_t, u_t) is a joint white noise process, the regression has a first-order moving average error structure. Empirically, the MA structure of the errors was clearly significant. Moreover, the innovations in the true price are probably heteroskedastic, as suggested by the results of Hausman et al. (1992). One of the reasons for the heteroskedasticity is the difference in the calendar time span between transactions. However, there may be other factors that cause a time-varying conditional variance. Instead of specifying the form of heteroskedasticity, we estimate by OLS, which under the stated assumptions gives consistent point estimates, and compute heteroskedasticity and autocorrelation consistent (HAC) standard errors using the method proposed by Newey and West (1987).

The details of the estimation procedure are as follows. In line with the previous

sections, we exclude all transactions outside the mandatory quote period (London) and the period of continuous trading (Paris). We include all other transactions, including the crosses in Paris. We take logarithms of the transaction prices and multiply those by 100 to obtain estimates of the percentage spread. The estimation equation is thus specified in returns, but only within-day returns were used because overnight returns are unlikely to follow the same process as intra-day returns, see Hausman et al. (1992). The classification of the trade as buyer initiated or seller initiated is done by comparing the transaction price with the mid-price. If the transaction price exceeds the mid-price, the trade is classified as buyer initiated ($Q_t = 1$), and if the transaction price is lower than the mid-price the trade is classified as seller initiated ($Q_t = -1$). If the transaction price is exactly at the mid-price, the trade is not classified and the value 0 is assigned to Q_t . This procedure is exact for the Paris transactions that were executed through the CAC system, but for the crosses and the London data there might be some incorrect classifications due to reporting lags.

In the literature several spread estimators have been developed for cases where no data on sign or size of the transactions are available. Generally, these estimators are unbiased under much stricter assumptions than necessary for the regression based estimator. Moreover, they are less efficient. For reasons of comparison we report two well-known alternative estimators of the realised spread. Roll (1984) proposes an estimator of the spread based on the first-order autocovariance of the returns, $\gamma_{\Delta p} = E(\Delta p_t \Delta p_{t-1})$. In the simple model (5), Roll's estimator is consistent only under some very restrictive assumptions: no serial correlation in expected returns; no error term in the price equation ($\sigma_u^2 = 0$); no serial correlation in the transaction type ($E[Q_t Q_{t-1}] = 0$); and no asymmetric information or inventory control effects ($\pi = 1$). Under these assumptions, the first-order autocovariance of the returns is equal to $-(S/2)^2$, and Roll's estimator of the spread is given by

$$S_{\text{Roll}} = 2\sqrt{\gamma_{\Delta p}}. \quad (6)$$

Roll's estimator is biased downward if there is positive serial correlation in the transaction sign Q_t (i.e. if transactions at the bid tend to be followed by further transaction at the bid and similarly for the ask). Choi et al. (1988) adjust to Roll's estimator for serial correlation in Q_t , retaining the assumptions that there are no pricing errors ($\sigma_u^2 = 0$), no serial correlation in mid-price returns and no asymmetric information or inventory control effects ($\pi = 1$). Choi et al. (1988) assume also that Q_t follows a first-order Markov process. Under these assumptions, the CSS estimator is

$$S_{\text{CSS}} = 2\sqrt{\gamma_{\Delta p}} / (1 - \gamma) = S_{\text{Roll}} / (1 - \gamma), \quad (7)$$

where γ is an estimate of the first-order autocovariance of the transaction sign.

Table 5

Model based estimates of realised spread. Model: $\Delta p_t = \beta_0 + (S/2)Q_t - \pi(S/2)Q_{t-1} + \epsilon_t$.^a

	Paris			London		
	Roll	CSS	S	Roll	CSS	S
AC	0.178	0.259	0.214 (47.855)	1.075	1.802	0.890 (10.214)
AQ	0.143	0.196	0.167 (65.196)	1.136	2.040	1.290 (13.740)
BN	0.147	0.182	0.169 (86.701)	0.679	1.354	0.781 (12.666)
CA	0.154	0.241	0.179 (56.733)	1.003	1.991	0.809 (11.010)
CS	0.274	0.359	0.330 (48.947)	2.152	3.997	1.131 (6.961)
EX	0.109	0.157	0.123 (58.959)	0.954	1.717	0.771 (12.077)
OR	0.246	0.328	0.285 (59.805)	0.849	1.401	0.992 (11.418)
RI	0.248	0.336	0.305 (35.139)	0.748	1.284	0.819 (6.765)
SE	0.253	0.371	0.316 (41.965)	2.071	4.186	1.901 (4.396)
UAP	0.349	0.521	0.404 (49.420)	0.902	1.444	0.842 (11.575)

^a Estimates of the realised spread. For definitions of estimators see Section 7. All transactions within continuous trading period or mandatory quote period were used, but overnight returns were excluded. Estimates are percentages of the transaction price. Newey–West *t*-values of the regression based estimates in parentheses.

The model-based estimates of the realised spread in London and Paris are given in Table 5. Comparing the regression-based realised spread estimator with Roll's estimators and the CSS estimator, it appears that the upward bias due to noise in the pricing equation is about the same as the downward bias due to the positive serial correlation in trade sign.⁶ The estimates suggest that the realised spread in London substantially exceeds the realised spread in Paris. Comparing the average effective spread in Table 4 with the regression-based realised spread estimate, the latter is uniformly smaller for all stocks, both in Paris and in London, suggesting that the average of best bid and ask quotes is not a good approximation of the unobserved true mid-price at the time of the transaction. This discrepancy is particularly striking in the case of the London data.

⁶ The first order serial correlation coefficient of the sign is about 0.3 for all series. The Roll estimates are therefore about the same as the regression based estimates, whereas the CSS estimates are much bigger.

Table 6

Model based estimates of realised spread, with size effects. Model: $\Delta p_t = \beta_0 + \delta Q_t + \alpha z_t + \gamma/z_t + \text{lags} + \epsilon_t$.^a

	Paris				London			
	2 δ	2 α	2 γ	Wald	2 δ	2 α	2 γ	Wald
AC	0.181 (23.338)	0.059 (2.870)	2.399 (4.279)	18.513	0.976 (8.018)	-0.066 (-1.311)	-7.459 (-1.817)	4.123
AQ	0.145 (36.234)	0.014 (1.745)	3.979 (6.944)	52.172	1.283 (11.391)	0.002 (0.048)	3.167 (0.486)	0.241
BN	0.158 (58.359)	0.038 (3.665)	0.184 (4.020)	21.349	0.591 (6.777)	0.187 (1.569)	13.980 (5.450)	30.103
CA	0.172 (44.577)	0.012 (0.527)	0.107 (2.878)	8.364	0.819 (9.655)	-0.095 (-1.501)	2.325 (1.370)	4.263
CS	0.302 (30.145)	0.073 (1.985)	0.298 (3.385)	11.630	0.793 (3.184)	0.188 (1.794)	6.996 (5.240)	27.599
EX	0.115 (41.387)	0.017 (1.163)	0.168 (3.931)	15.473	0.737 (10.204)	0.037 (0.572)	1.724 (1.770)	3.261
OR	0.257 (42.213)	0.047 (2.446)	0.854 (6.137)	37.772	0.735 (7.021)	0.113 (2.189)	70.061 (2.677)	9.312
RI	0.283 (20.015)	0.008 (0.226)	0.965 (2.077)	4.947	0.680 (3.696)	0.036 (0.408)	21.844 (1.381)	1.943
SE	0.313 (29.231)	-0.056 (-2.053)	0.561 (1.270)	9.337	2.163 (4.362)	-0.166 (-1.081)	-13.606 (-1.971)	3.896
UAP	0.348 (30.332)	-0.000 (-0.021)	3.764 (7.981)	77.760	0.814 (8.692)	0.012 (0.346)	0.606 (0.719)	0.534

^a All transactions within continuous trading period or mandatory quote period were used, but overnight returns were excluded. The size variable was censored at 2 NMS for Paris and 5 NMS for London. Newey-West *t*-values of the regression based estimates in parentheses. Wald is a χ^2 (2) test of joint significance of α and γ .

In addition to the effect of the *sign* of the trade (buyer or seller initiated) the *size* of the trade may also be an important determinant of the price. The microstructure theories discussed in Section 2 predict that due to asymmetric information and inventory control the spread will be an increasing function of trade size. To estimate the effect of size we extend model (3) in the spirit of Glosten and Harris (1988) and Madhavan and Smidt (1991). The price equation is extended with a linear term in the size of the transaction. In Section 6 we found some evidence for a decreasing spread for large trade sizes. This effect is captured by adding the inverse of trade size to the price equation, which picks up possible non-linear dependence of the spread on size. The extended pricing equation is

$$p_t = y_t + (S/2)Q_t + \alpha z_t + \gamma z_t^{-1} + u_t, \quad (8)$$

where z_t is the signed trade size. First-differencing (8) we obtain the equivalent of

Table 7

Implied realised bid-ask spreads

Paris ^a					
Size:	0.1	0.5	1.0	2.0	
AC	0.199	0.213	0.241	0.299	
AQ	0.155	0.154	0.160	0.173	
BN	0.162	0.177	0.196	0.234	
CA	0.173	0.178	0.184	0.196	
CS	0.313	0.340	0.376	0.449	
EX	0.117	0.123	0.132	0.148	
OR	0.265	0.281	0.305	0.352	
RI	0.293	0.289	0.292	0.299	
SE	0.311	0.286	0.257	0.201	
UAP	0.367	0.351	0.349	0.348	
London ^a					
Size:	0.1	0.5	1.0	2.0	5.0
AC	0.932	0.935	0.906	0.841	0.644
AQ	1.289	1.285	1.286	1.288	1.294
BN	0.666	0.696	0.783	0.967	1.525
CA	0.818	0.774	0.726	0.631	0.347
CS	0.882	0.901	0.988	1.173	1.734
EX	0.747	0.757	0.775	0.811	0.921
OR	1.026	0.848	0.867	0.975	1.306
RI	0.902	0.741	0.738	0.763	0.865
SE	2.079	2.067	1.991	1.829	1.334
UAP	0.819	0.821	0.827	0.839	0.876

^a Realised spread computed from the estimates in Table 6, Panel A.^b Realised spread computed from the estimates in Table 6, Panel B.

regression equation (5) but now including current and lagged trade size and the inverse of size as regressors: ⁷

$$\Delta p_t = \beta_0 + (S/2)Q_t + \alpha z_t + \gamma z_t^{-1} + \text{lags} + e_t + \Delta u_t. \quad (9)$$

In order to reduce the influence of very large transactions (outliers) on the estimates, we 'censor' large trade sizes. For Paris, we pick the threshold at 2 NMS, which is about the 99.5% quantile. ⁸ In London many more trades would be censored at 2 NMS, between 10 and 25 percent. Therefore, we use a threshold of 5 NMS for the London data, which corresponds to the 95% quantile, see Table 2C.

Estimates of the trade size augmented model are given in Table 6. For Paris,

⁷ We do not impose restrictions on the coefficients of the lagged regressors. We do not want to run the risk of imposing invalid restrictions and thus misspecifying the model. Not imposing such restrictions does not affect the consistency of the estimators of the parameters of interest (S , α and γ).

⁸ Hausman et al. (1992) also censor trade size at the 99.5% quantile.

the coefficients of the size and the inverted size are small but significant for most cases. The Wald test of joint significance of the size and inverted size parameters is larger than its 5% critical value (5.99) for all series except one. On the other hand, for London less evidence for a trade size effect on the realised spread is found. The size effect is jointly significant only for BSN (BN) and Axa-Midi (CS). Partly this may reflect the smaller sample size of the London series.

Table 7 shows the implied estimates of the realised bid–ask spread. The estimates in Table 6 were used to construct these spreads. The spread in Paris is slowly increasing for large sizes, but in London there is no clear pattern, some spreads increase and some decrease with size. We confirm the previous conclusions that for all sizes up to 2 NMS the realised spread in Paris is uniformly smaller than the realised spread in London for all stocks.

Our results are robust to an extension along the lines of Hasbrouck (1991), who advocates a more extensive model to assess the dynamic effects of transactions. More specifically, in his model the price effect of a transaction can last for more periods than the one period assumed implicitly in Eq. (4). Our regression based spread estimator can be extended easily to include more complex dynamics by adding lagged regressors. The spread estimates obtained using four (rather than one) lags of trade sign show only minor differences with the reported estimates and the conclusions do not change.

A decomposition of the realised spread in cost components is beyond the scope of this paper. In De Jong et al. (1994) we calculate that the price impact of a transaction is between 25% and 40% of the total spread. In Stoll's (1989) model, this gives the sum of asymmetric information and inventory control components.

8. Summary and conclusions

In this paper we compare the cost of trading French shares in Paris and in London. The estimates of the average quoted spread, which reflect the cost of immediate trading, suggest that the quoted spread on the Paris Bourse is lower than London's SEAQ International for small transactions, roughly up to the normal market size. Röell (1992) however shows that for very large sizes the Paris limit order book often does not contain enough limit orders and the average quoted spread rises steeply, hence the Paris market is not very deep. The London market with its competing market makers provides more immediacy for large sizes. The quoted spread in London for small sizes is relatively large.

The estimates of the effective and realised spread show a slightly different picture. It appears that the few large transactions that are executed in Paris (often 'crosses') have a fairly low spread, lower than the spread in London. Our regression-based estimates suggest that at trade sizes of twice the NMS the realised spread in Paris is still considerably lower than in London. On the whole, we conclude that if the trader is patient and prepared to wait for counterparties,

transaction costs for large sizes can be fairly low in Paris compared with SEAQ International.

The full cost of trading on either exchange includes taxes and other levies as well. Information on such explicit transaction costs are presented in London Stock Exchange (1992a). The commissions and fees in London are on average 0.14% of the transaction value and in Paris about 0.5% (these percentages are for a large transaction of 1 million ECU, roughly FF 7 million). Thus explicit transaction costs are higher in Paris for large transactions. One reason is that in London many large deals are done on a 'net' basis, i.e. commissions are included in the price.

A theoretically interesting result is that the effective spread in Paris is virtually flat in trade size, whereas the effective spread in London declines with size. Hence, we do not confirm the predictions of the pure inventory control or adverse selection microstructure theory that the spread should be an increasing function of trade size. Our estimates of a simple model for transaction prices confirm this result and indicate mild support for the hypothesis that part of the order processing cost is fixed per transaction rather than per share.

All in all our results do not explain the overwhelming success of SEAQ International in capturing wholesale trade in non-British equities. Perhaps, factors such as immediacy and execution risk play a crucial role. These are not captured in our trading cost estimates.

Appendix: Adjustment for bias due to misreported transaction times

As explained in the main text, the estimates of the average realised spread in London for transaction sizes smaller than NMS are sometimes larger than the quoted spread. This seems impossible, because the true effective spread has to be smaller than the quoted spread since market makers are obliged to provide the best quoted price for transactions smaller than NMS. This anomaly is probably explained by a timing bias due to misreported transaction times in London. In this appendix we propose a model for the impact of timing bias on estimates of the effective spread that can also be used to correct the effective spread estimates for this bias.

Let $S(z)$ be the average effective spread (as a function of size) that we would want to estimate. Suppose that the transaction is reported inaccurately. In general the mid-price recorded at the reporting time is different from the mid-price at the correct time. Denoting the change in mid-price by x , we effectively estimate

$$\hat{S}(z) = E | S(z) + x |. \quad (\text{A.1})$$

Suppose that x is normally distributed with mean 0 and variance σ^2 . Then we can apply the expressions in Amemiya (1985, p. 367), who shows that for a




Fig. 2. Timing bias by standard deviation of reporting error.

normally distributed variable $y \sim N(\mu, \sigma^2)$, the conditional expectation of y , given $y > 0$ is

$$E(y | y > 0) = \mu + \sigma \frac{\phi(\mu/\sigma)}{\Phi(\mu/\sigma)}, \tag{A.2}$$

where ϕ and Φ are the standard normal density and the cumulative standard normal distribution, respectively. Using this result the expectation of the absolute value in (A.1) can be written as

$$\hat{S}(z) = S(z)(2\Phi(a)-1) + 2\sigma\phi(a), \quad a \equiv S(z)/\sigma. \tag{A.3}$$

Fig. 2 shows that $\hat{S}(z)$ is always larger than the quantity that we want to estimate, i.e. $S(z)$. The estimates reported in Table 4B will therefore in general overstate the true effective spread if $\sigma > 0$.

We now turn to a method to correct for timing bias. Fundamental to the correction is the assumption that the timing error is independent of the transaction size.⁹ There is one problem with the procedure outlined just before. If we take the smallest size class to be the class from 0 to 0.1 NMS, we estimate quite a large σ . In fact, the estimated σ is often so large that Eq. (A.3) does not have a solution

⁹ In Section 6 we found that the effective spread in London was decreasing in size. A referee pointed out that this could be due to late reporting of transactions and a price impact that increases with trade size. Moreover, it is known that most small transactions are at the touch (London Stock Exchange (1992b)). Thus, for small transactions the quoted spread and the true effective spread should be the same: $S(z) = S_Q(z)$. In Table 3, panel B the average quoted spread by size class can be found. The first step in the correction procedure is to solve (A.3) for σ , given $S(z) = S_Q(z)$ in the smallest size class. The second step is to compute $S(z)$ for all other size classes from (A.3), given the estimate of σ obtained in the first step and the uncorrected effective spread estimate from Table 4B.

for $S(z)$. Therefore, we choose to base the estimate of σ on the average quoted spread for all transactions up to 1 NMS. Because quotes are firm up to 1 NMS, the quoted spread is an upper bound for the effective spread for this size class. Therefore, the estimated σ from solving (A.3) given $S(z) = S_Q(0, 1)$ gives a lower bound for the true timing error. This estimate of σ will therefore yield a conservative correction of the effective spread for other size classes.

References

- Amemiya, Takeshi, 1985, *Advanced econometrics* (Academic Press, New York).
- Anderson, Ronald W. and Pierre Tychon, 1993, Competition among European financial markets: The case of Belgian cross-listed stocks, Discussion paper 9314 (Department of Economics, Université Catholique de Louvain, Louvain-la-Neuve).
- Bagchot, Walter, pseud., 1971, The only game in town, *Financial Analysts Journal* 22, 12–14.
- Biais, Bruno, Pierre Hillion and Chester Spatt, 1992, An empirical analysis of the limit order book and the order flow in the Paris Bourse, Paper presented at CEPR Workshop on Market Microstructure, 3–4 April 1992, Konstanz.
- Bredon, Francis, 1993, Intraday price formation on the London Stock Exchange, Discussion paper no. 153 (Financial Markets Group, London School of Economics, London).
- Choi, J.Y., Dan Salandro and Kuldeep Shastri, 1988, On the estimation of bid–ask spreads: Theory and evidence, *Journal of Financial and Quantitative Analysis* 23, 219–230.
- De Jong, Frank, Theo Nijman and Ailsa Röell, 1993, A comparison of the cost of trading French shares on the Paris Bourse and on SEAQ International, Discussion paper no. 9329 (CentER, Tilburg) and Discussion paper no. 169 (Financial Markets Group, London School of Economics, London).
- De Jong, Frank, Theo Nijman and Ailsa Röell, 1994, Price effects of trading and components of the bid–ask spread on the Paris Bourse, Discussion paper no. 9454 (CentER, Tilburg).
- Easley, David and Maureen O’Hara, 1987, Price, trade size and information in securities markets, *Journal of Financial Economics* 18, 235–248.
- George, Thomas, Gautam Kaul and M. Nimalendran, 1991, Estimation of the bid–ask spread and its components: A new approach, *Review of Financial Studies* 4, 623–656.
- Glosten, Lawrence and Paul Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71–100.
- Glosten, Lawrence and Lawrence Harris, 1988, Estimating the components of the bid/ask spread, *Journal of Financial Economics* 21, 123–142.
- Hasbrouck, Joel, 1991, Measuring the information content of stock trades, *Journal of Finance* XLVI, 179–207.
- Hausman, Jerry, Andrew Lo and A. Craig MacKinlay, 1992, An ordered probit analysis of transaction stock prices, *Journal of Financial Economics* 31, 319–379.
- Ho, Thomas and Hans Stoll, 1981, Optimal dealer pricing under transactions and return uncertainty, *Journal of Financial Economics* 9, 47–73.
- Kyle, Pete, 1985, Continuous auctions and insider trading, *Econometrica* 53, 13–32.
- London Stock Exchange, 1992a, Institutional transaction costs in major European markets, *Stock Exchange Quarterly*, summer edition, 19–26.
- London Stock Exchange, 1992b, Price improvement and best execution, *Stock Exchange Quarterly*, spring edition, 25–31.
- Madhavan, Ananth, 1992, Trading mechanisms in securities markets, *Journal of Finance* 47, 607–641.
- Madhavan, Ananth and Seymour Smidt, 1991, A Bayesian model of intraday specialist pricing, *Journal of Financial Economics* 30, 99–134.

- McNish, Thomas and Robert Wood, 1990, An analysis of transactions data for the Toronto stock exchange, *Journal of Banking and Finance* 14, 441–458.
- Newey, Whitney and Kenneth West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, no. 3, 703–708.
- Pagano, Marco and Ailsa Röell, 1993, Transparency and liquidity: A comparison of auction and dealer markets with informed trading, Working paper no. 54 (IGIER, Bocconi University, Milan).
- Röell, Ailsa, 1992, Comparing the performance of stock exchange trading systems, in: Fingleton and Schoenmaker, eds., *The internationalisation of capital markets and the regulatory response* (Graham and Trotman, London).
- Roll, Richard, 1984, A simple implicit measure of the effective bid–ask spread in an efficient market, *Journal of Finance* 39, 1127–1139.
- Stoll, Hans, 1989, Inferring the components of the bid–ask spread: Theory and empirical tests, *Journal of Finance* 44, 115–134.
- Tonks, Ian and Andy Snell, 1992, Trading volumes and stock market prices, Discussion paper no. 130 (Financial Markets Group, London School of Economics, London).