Dealer liquidity in an auction market: evidence from the London Stock Exchange

Sylvain Friederich and Richard Payne*†

First draft: December 15, 1999
This draft: August 30, 2002

Abstract

We analyse the trade characteristics and market conditions which determine the market share of an electronic order book at the London Stock Exchange, where an "upstairs" network of dual-capacity firms is also available for trade. We hypothesise and empirically verify that execution and information risks govern the choice of execution mode. Further, we uncover strong commonality in the market share of the order book across stocks, and find that variables proxying for market-wide liquidity and informational risks also affect the choice of trading venue. These findings appear robust to possible endogeneity of the measures of order book liquidity. They suggest that competing, off-book liquidity suppliers voluntarily perform at least some of the "stabilisation" functions normally assigned to designated market-makers.

Keywords: Limit order trading. Auction and dealership markets. Commonality. London Stock Exchange.

*Department of Economics, University of Bristol and Université de Paris - I, and the Financial Markets Group and Department of Accounting and Finance, London School of Economics. Correspondence: s.friederich@bristol.ac.uk; r.payne@lse.ac.uk.

†The financial support of INQUIRE-Europe is gratefully acknowledged. Thanks to Ron Anderson, Jim Angel, Bruno Biais, Greg Connor, Vassilis Hajivassiliou, Roger Huang, Theo Nijman, Gideon Saar and Ian Tonks, as well as participants at the EFMA 2000, NASDAQ/Notre-Dame 2000 and Inquire 2001 conferences, and seminar participants at the LSE, Reading, Said Business School at Oxford, Toulouse-ESC, Warwick. Thanks also to Ian Rowell, Matthew Leighton, Jamie Lebetkin and Alan Line at financial firms and exchanges. Remaining errors are our own.

A pervasive feature of equity market structures today is that there are multiple mechanisms via which a liquid stock can be traded. More specifically, market structures appear to have converged towards hybrid systems in which a central role is given to order-driven trading. The two main equity market reforms of the late 1990's, of the Nasdaq market and the London Stock Exchange (LSE), were both moves in this direction in that limit order trading was introduced to markets previously structured as pure dealerships.

However, in all mature equity markets, order books are supplemented by the parallel supply of non-anonymous dealership services. On the NYSE the public limit order book is supplemented by the activities of the Specialist, upstairs brokers and NASD broker-dealers. The Paris Bourse is often regarded as the market closest to a pure limit order structure, but off-book trading is available on-exchange for transactions over a given size and, more importantly, London-based dealers provide off-exchange liquidity in actively traded French stocks. Other European exchanges also have separate (off-book) arrangements in place for the execution of large orders.[1]

Based on the above, understanding the operation of hybrid markets is a key issue in current empirical microstructure research. We contribute to this area, analysing the nature of the order flow routed on and off the order book in a hybrid structure, and the stock characteristics or market conditions that lead to the dominance of one mode of trading over the other. Order flow may, for example, be executed off book as the result of preferencing and internalization arrangements, or through broker-dealers free-riding on book prices and "cream-skimming" uninformed orders (Easley, Kiefer, and O'Hara, 1996). Alternatively, order flow may be routed off book because of low book liquidity or high volatility of book prices, in which case dealers are providing a valuable service.

We address these issues using two months of LSE data recorded at the end of 1999 for the FT-30 stocks. The current microstructure of the LSE comprises a central limit order book (called SETS) which the exchange describes as "the main price formation mechanism" and alongside the book there exists a network of broker-dealers who supply quotes on a bilateral basis - similar to an "upstairs market". The London market constitutes a particularly interesting venue for a number of reasons. First, the LSE's market structure exemplifies current trends in trading systems, where more freedom of choice of trading venue is left to the "buy side," and more competition

[1] This is documented in the overview of Demarchi and Foucault (1999).

is encouraged in those venues among liquidity suppliers. Specifically, there are no forced interactions between the two trading mechanisms either in price or quantity terms: no part of an order must be executed in the order book and the prices of off-book trades are not constrained by book prices. This is not usually the case for order flow executed off the book but on-exchange. On the NYSE, for example, specialist liquidity supply is constrained in a number of ways while trades negotiated with upstairs brokers must be exposed to the downstairs floor and order book. In Paris, only trades over a given size threshold can be executed off-book and only then at or within weighted average book prices. This reduces at the same time the ability of traders to route trades off the book and the incentives member firms may have to supply such off-book liquidity.

The coverage and quality of the data set are another reason why examination of the London market is interesting. The data provided by the exchange accounts for essentially all the trading activity, on and off book, in these liquid stocks and full book order information. During our sample period, the only marketplace that competed with the LSE was a system called Tradepoint, which executed a negligible fraction of order flow (Board and Wells, 2000). Data available from most European exchanges simply contains no information on trades executed off-exchange by the member firms. As the scale of this trading activity is large this is a serious deficiency.[2] Compared to data available from many other markets, for example most NYSE data sets, we are also able to distinguish trades executed against limit orders from those where an intermediary supplied liquidity.

Our empirical work focusses on characterizing the determinants of the fractions of trading activity sent on and off the central limit order book in terms of stock-specific and market-level information. We estimate a panel specification for the share of order flow sent on and off book using a 30-minute sampling frequency. The results demonstrate that the market share of the order book tends to be low if traded volume is extremely high, if mean transaction size is high and if trade imbalance (i.e. the excess of buys over sells or sells over buys) is large. These results are consistent with theory based on the influence of execution risk (Demsetz, 1968; Grossman and Miller, 1988) and information risk (Pagano and Röell, 1992) on auction versus dealer based trade. When execution risk is high due to the characteristics of trades (e.g. size) or overall activity (e.g. volume) or when information risk is high (e.g. order flow

[2] To illustrate, for January to October 2000 figures from the Paris exchange indicate book turnover of Euro 478 bn. and total (i.e. including off-book) turnover of Euro 1456 bn.

not balanced) the liquidity service supplied by intermediaries is increasingly utilized. We further show that the share of order flow sent to SETS is low if SETS spreads are high or if SETS depth is low. Hence, when the terms of trade on SETS are poor, order flow migrates to the dealer segment of the market. The preceding results are robust to various diagnostic checks and alternative estimation methods. In particular, we show that our results are robust to possible endogeneity of the liquidity-related variables appearing on the right-hand side of the regression by estimating a panel IV model.

We proceed on to examine commonality across stocks in order book utilization. Recent work (Chordia, Roll, and Subrahmanyam, 2000; Hasbrouck and Seppi, 2001) has demonstrated important cross-dependencies in trading activity and liquidity measures using US data. We first verify these findings, indeed documenting effects stronger than reported in these previous studies, possibly because our more detailed LSE data allows us to construct better measures of market tightness and depth. We further show that the proportion of order flow executed off-book is also correlated across stocks. Thus there are periods when trade tends to migrate towards or away from the central limit order book for *all* FT-30 stocks. We interpret this result as due to market-wide price or informational risk factors and evaluate this interpretation by augmenting our panel regression to incorporate market-wide trading activity and trading imbalance measures, drawn from transaction level FTSE-100 futures data. The analysis shows that when market-level trading activity is particularly great or order flow tends to be one-sided, then activity for individual stocks migrates from the electronic order book to the dealer segment of the market.

In sum, we feel that this evidence is indicative of dealers performing a stabilization and liquidity supply role similar to that assigned to dedicated market-makers. These dealers step in and are willing to intermediate an increased proportion of order flow when stock-specific or market-wide risks arise. Based on this assessment, there appears to be no clear need to re-introduce designated intermediaries in London.

The rest of the paper is structured as follows. Section 1.1 provides some theoretical motivation and empirical background and sets out our research questions in more detail. Section 2 provides more information on the current operation of the London equity market, introduces and provides basic analysis of the LSE data and discusses how we process it. Section 3 presents our empirical analysis of order flow and discusses the results derived. Finally, Section 5 summarizes our findings and indicates directions for future research.

1 Related Work

1.1 Theory

The empirical focus of the current paper is to understand when the dealer segment of a hybrid market dominates the order book segment in terms of quantities traded and vice versa. Several strands of the theoretical microstructure literature yield predictions that guide us in the choice of variables that should determine the relative success of the two venues.

The most fundamental area of relevant literature is that analysing the rationale for the existence of dealers/market-makers in continuous financial markets, something not part of the standard Walrasian framework. The classic reference for this is Demsetz (1968) who views market-makers as suppliers of continuity or immediacy. The need for their services arises from the fact that buy and sell orders in continuous markets are not synchronised, such that order flow imbalances occur naturally. Some traders are then willing to pay a higher price in order to transact now rather than wait for matching orders to arrive on the other side of the market. Several empirical implications follow from this intuition;

- Demand for dealership services is negatively related to trading activity in a stock.
- Dealers should have greater value at times of one-sided market order flow.
- Increased trade size, for given total activity, should increase dealer participation.

Grossman and Miller (1988) build on Demsetz' contribution to investigate the implications of the demand for and supply of immediacy for market structure in a world of liquidity-motivated trading. At the heart of the model is execution risk, arising from the possibility that an adverse change in the equilibrium price may occur in the time it takes for an order to execute. Dealers appear and charge for the opportunity costs of maintaining a continuous presence in the market and for bearing execution risk. The model predicts that assets for which demand for immediate execution is high will tend be to traded on dealership structures. In turn, the demand for immediacy will depend to a large extent on the degree of price uncertainty (volatility). Thus, we get the following empirical prediction;

- The demand for dealership services should be positively related to a return volatility.

A second way to motivate a relationship between volatility and off-book trading is via the "free option" analogy of limit orders (Copeland and Galai, 1983): one would expect traders to switch to bilateral negotiation of trades rather than post exposed limit orders at times of greater volatility. In the setting of Copeland and Galai, this effect is due to the fact that high volatility intervals are associated with information asymmetries between liquidity demanders and suppliers.

Off-book trading in London can be thought of as constituting an upstairs market using the economic definition of Grossman (1992): "A downstairs market refers to an organized exchange where all members agree that trades take place publicly in a central place. An upstairs market refers to a market where trading takes place privately." Thus, we draw from this second strand of literature for potential theoretical predictions.

Burdett and O'Hara (1987) model upstairs firms in their twin roles of wholesale brokers or matchmakers and suppliers of immediacy (dealers). In the paper, continuous auction systems cannot accommodate large trades easily because of the information risk inherent to them, and the services of a dual-capacity intermediary are required. The analysis focusses on the trade-off these firms face between the benefits of risk syndication (locating other counterparties to the block transaction) and the adverse price movements brought about by the leakage of information which this search causes.[3] This analysis thus generates a prediction comparable to that in Demsetz (1968) regarding the relationship between trade size and order book trading, though with an emphasis on information risk.

Grossman (1992) argues that upstairs brokers have private information regarding the trading interests of large traders ("unexpressed pools of liquidity") such that when downstairs depth is low, upstairs trading can result in better execution. Other models of the upstairs trading process focus on non-anonymity as a way to alleviate information risk, arguing that in the upstairs market, trader intentions/characteristics can be screened, thus generating superior upstairs execution for uninformed block

[3]Miller (1991) mentions that "until recently at least, the upstairs desks functioned primarily as a search market" but [...] "the amount of 'positioning' and hence of market-making liquidity provided by upstairs firms has increased substantially in recent years." Stoll (1993) similarly evokes the "dealerization" of the NYSE, - although there is no direct evidence of this type of liquidity supply at the NYSE to our knowledge.

traders (Seppi, 1990). (The theoretical part of Madhavan and Cheng (1997) contains a similar intuition). A common thread in these models of upstairs trading then is that upstairs firms supply liquidity in sizes that the book cannot accommodate, such that incentives to route an order upstairs should be inversely related to book depth. Hence;

- We expect a negative relationship between off-book activity and order book depth.

The last main strand of theoretical literature relevant in the current context compares the performance of dealership and auction markets when these markets are faced with similar trading conditions. Papers relevant here include Madhavan (1992), Pagano and Röell (1992) and Pagano and Röell (1996). It should be stressed, however, that these papers do not consider the simultaneous operation of auction and dealer markets for the same asset, rather they look at each type of market architecture in isolation.

Pagano and Röell (1992) provide an interesting discussion of the relative benefits of auction and dealership markets along a number of different dimensions. As in some of the upstairs markets literature, they argue that dealership markets offer opportunities for screening of trader intentions and thus the ability to reject the trading advances of those deemed informed. Based on this intuition, one would argue that in times of high information asymmetries dealers would still trade with those verifiably non-informed while informed agents would seek to trade in the order book. Realizing this, those supplying liquidity to the order book would likely do so on poor terms such that few trades would be consummated on the book. Based on similar intuition, Easley, Kiefer, O'Hara, and Paperman (1996) argue that less-liquid stocks are more suited for trade on dealership structures, due to greater informational asymmetries for these stocks. Thus the following prediction is generated;

- High information asymmetry should increase non-anonymous off-book trading.

Pagano and Röell (1992) go on to point out that in dealer markets traders are always forced to pay the spread, this not being the case in auction markets. Finally, they argue that in a centralized market such an as order book, information can enter prices more rapidly while in a decentralized dealer market information resides in the hands of individual dealers and thus prices impound information less rapidly. The implications of such transparency issues for liquidity and trading costs under various market architectures are discussed more fully in Pagano and Röell (1996).

A further paper that should be mentioned here is a recent contribution by Viswanathan and Wang (2002). The focus of this paper is the problem of a customer choosing which of three alternative market structures on which to trade. The alternatives are a pure limit-order market, a pure dealership and a hybrid of the two (defined as a system where orders beyond a certain, exogenously defined cut-off size are routed to dealers while smaller orders are sent to the order book). The authors show that, under certain conditions on order flow, risk-neutral traders prefer a pure limit order system, while risk-averse traders would choose a dealership structure (under the condition that the number of dealers is large). Finally, when it is the case that risk-averse traders prefer dealership to auction structures, it is shown that the hybrid structure dominates both. These results are important from the perspective of market design and can be used, for example, to justify the choice of a hybrid trading arrangement for London. However, the value of these results from the perspective of building an empirical model of where to trade in a hybrid market is smaller, given their assumption that all trades over a certain size threshold *must* be executed off-book in the hybrid market.

Finally, a focus of our empirical analysis is cross-dependencies on order book usage. However, above we have discussed the implications of stock-specific information risk (Pagano and Röell, 1992) and price risk (Grossman and Miller, 1988). Our intuition for commonality in book usage is driven by the possible existence of market-wide information or price risk. Such market-level shocks might drive trading onto or off of the order book for many stocks simultaneously for precisely the reasons isolated in our stock-by-stock examples. We know of no paper which explicitly models such market-wide effects but appeal to intuition and to existing related empirical work in the next subsection.

1.2 Related empirical work

One of the main stated aims of the 1997 reform introducing limit order trading in London was to lower trading costs for public investors. Naik and Yadav (1999) examine changes in the cost of trading before and after the introduction of SETS, in the spirit of the study of the Nasdaq market by Barclay, Christie, Harris, Kandel, and Schultz (1999). They arrive, by and large, at comparable conclusions: public order exposure has produced a more competitive market, with less cross-subsidisation across different categories of traders, and various spread measures indicate a decline in trading costs. Another study using recent SETS data is Board and Wells (2000), who examine the amount and characteristics of order flow executed on Tradepoint (an alternative order-driven trading system) since the launch of SETS. They report that in 1998 Tradepoint had a market share greater than 1% for only four stocks, a figure which has declined to more or less zero since then.

Turning to empirical studies of "upstairs" trading, the most well-known is the aforementioned paper by Madhavan and Cheng (1997). To this day, the authors are the only researchers to have been able to identify trades facilitated upstairs at the NYSE. They examine whether these trades are routed upstairs by block traders who want to make use of a lack of anonymity to signal that they are uninformed, and therefore obtain better price execution. Madhavan and Cheng (1997) report evidence consistent with this hypothesis, after correcting for the selectivity bias induced by the fact that traders who are able to signal they are uninformed will tend to select the upstairs venue. With UK market data comparable to ours, Ellul (2001) investigates patterns in volatility and examines trader choice using the selection model introduced in Madhavan and Cheng (1997). He reports much the same findings. However, application of the Madhavan and Cheng (1997) technique to LSE data is problematic in that dealer trades on the LSE can be reported with up to a 3 minute delay in normal conditions. This means that at the transaction level, the book and off-book trades cannot be sequenced reliably and thus the values taken by a number of trade-related variables are not measurable without error. Accordingly, in his analysis of where trades are executed, Ellul (2001) does not include any trading activity variables in his specification and thus cannot evaluate many interesting propositions. Finally he shows only that book usage tends to be low when book liquidity is low. Two other recent studies of upstairs trading, Booth, Lin, Martikainen, and Tse (2002) and Smith, Turnbull, and White (2001) report, for the Helsinki and the Toronto stock exchanges respectively, that upstairs trades have a lower permanent price impact than those executed downstairs, which they ascribe to the difference in anonymity between venues. Booth, Lin, Martikainen, and Tse (2002) find that most of the price discovery occurs on the anonymous order book, while Smith, Turnbull, and White (2001) report that upstairs brokers are likely to execute order flow they identify as uninformed against their own inventory, consistent with the "screening" hypothesis.

The analysis in the current paper is also related to earlier research on cross-traded securities in European markets, which documented that most of the block activity in liquid Continental issues is reported by London broker-dealers (de Jong, Nijman, and Röell, 1995). Röell (1992) suggested that the total depth available for trade in large size from the quotes available from London-based dealers could be much superior to that obtainable from the Paris Bourse electronic order book. However, for lack of detailed exchange data at the time, she could only use limit order data from screen captures for one stock and one day in 1991 and was thus unable to examine the role of total depth more systematically. A recent paper by Bessembinder and Venkataraman (2002) studies the execution of large trades on the Paris Bourse and reports that they tend to be granted price improvement compared to order book terms. This study, however, is subject to the caveats mentioned earlier: the Paris Bourse *only* allows trades to be executed off its order book subject to both quantity and price constraints (they must take place within a spread computed from the order book). It is arguably because of these constraints that most block trades in active French issues are routed to London dealers (where off-book liquidity supply is easier) but the authors do not include this activity in their estimations.

The analysis perhaps most closely related to ours is contained in two papers. Madhavan and Sofianos (1998) study the determinants of Specialist participation in overall NYSE trading activity, both across stocks and over time. An important contribution of this paper for interpreting the nature of the trading system in operation at the NYSE is to document large cross-sectional variation in Specialist market shares. This implies that very active stocks effectively trade in an auction market, with Specialist participation at a minimum, while illiquid stocks trade in a quasi-dealership market. They also report that the Specialist tends to step in and supply liquidity at times of high volatility, wide spreads, or in large trade sizes, and in a way related to the Specialists' inventories. Although related to ours, their aims and results of that paper are still significantly different since, as mentioned above, the Specialist has regulatory obligations to stabilize markets. We are examining whether freely competing liquidity suppliers tend to behave similarly, but without obligations or privileges.

The second related paper (emerging around the same time as our work) is a study of on and off-book activity on the Australian Stock Exchange by Fong, Madhavan, and Swan (2001). In addition to a selection model, as in Madhavan and Cheng (1997), these authors estimate a panel regression model for the ratio of off to on-book trading activity, similar to that which we propose in Section 3. However, an issue with the

Australian data is that trading frequency is very low dictating that the authors use an annual sampling frequency in their regression analysis. This renders many interesting microstructure issues untestable. However, certain of their results from these annual data, notably those on the effects of book liquidity on off-book activity, corroborate those that we find.

Unlike the stock-specific factors discussed in the previous subsection, our examination of commonalities in the use of trading systems is largely empirically motivated and follows up on recent work by Chordia, Roll, and Subrahmanyam (2000) and Hasbrouck and Seppi (2001) which reports clear evidence of common factors in order flows, liquidity measures and returns. Market-wide liquidity shocks could have many causes but the intuition behind market-wide informational asymmetries is less straightforward. One way to think about such issues is presume that our sample stocks have at least one common return driver and that a subset of traders have access to a model that provides better than average forecasts of this driver. The activity of these traders would then generate market-wide informational asymmetries and thus similar tendencies for on or off book execution on a stock-by-stock basis.

2 The market and a first look at the data

2.1 The trading venues

The London equity market operated for about ten years as a pure dealership system. Those trading arrangements came under criticism for their opacity and the high trading costs incurred by smaller investors and were thus reformed in October 1997, giving a central role to order-driven trading. Since then, an individual or institution wishing to trade equity on the London market has a choice of two fundamental trading methods. First, the investor could contact his dealer and instruct her to submit a limit or market order to the SETS order book.[4] The second trading method is to call one's dealer and request an immediate trade with the dealer as counterparty. Note, however, that the supply of dealer services is entirely voluntary and unconstrained. Also, an investor is not forced to choose solely between these two extremes and,

[4]It should be noted that SETS is only available for trade in around 200 of the most liquid stocks and our discussion should be interpreted accordingly. However, these 200 stocks account for around 95% of overall LSE volume.

typically, investors simultaneously use both methods. In particular, an investor is not forced to execute any portion of an order on the SETS book.

The two trading methods differ along a number of dimensions. SETS is among the most transparent of limit order books in major equity markets as every outstanding limit order is displayed to member firms. There is also immediate publication of details of all book trades to participants. Finally, settlement is integrated into the SETS system. Shares traded on the book must be held in dematerialised form and have a $t + 5$ day settlement cycle. All settlement details are handled by an outside firm called CRESTco.

Off-book trading (i.e. direct trade between dealers, direct matches of client orders by dealers or clients trading against dealer inventory) is very different in nature. There is no pre-trade transparency as quotes are provided on a purely bilateral basis. Post-trade transparency is high in that dealers are forced to report any trade consummated within 3 minutes of its occurrence. Finally, there are no settlement restrictions in that paper shares may be traded and non-standard settlement cycles (i.e. not $t + 5$) may be requested.

A couple of other features of the London market relevant to our study are as follows. First, there are non-standard reporting rules for a class of block and portfolio trades known as "Worked Principal Agreements" (WPAs). These arrangements are a legacy of the dealership system in place pre-October 1997 and in our analysis we filter them from the data (see next subsection). A second feature of the market which impacts upon this study is the existence of three broker-dealer firms who supply retail order execution systems and are known as "Retail Service Providers" (RSPs). These systems guarantee prices that are at or within the book spread and, as such, essentially remove the incentives of retail investors to use the SETS order book.

2.2 The data set, filtering and basic analysis

The data we employ cover the months of November and December 1999. To keep our analysis manageable we focus only on data for constituent stocks of the FT-30 index. The FT-30 stocks are a convenient subsample of the entire universe of equities as they are drawn from various size/liquidity deciles. All stocks in this index are actively traded and, in aggregate, the raw dataset comprises 1,060,000 trades.

Basic information on our sample stocks is given in Table 1. The table presents firm names and identifiers, market values as of 1/12/1999 (the middle of our sample) and the final column indicates each security's "normal market size" (NMS) at the same date. The NMS is a measure of the average institutional trade size in a stock as computed (and regularly reviewed) by the Exchange.[5] We will use this measure repeatedly in our descriptive statistics as well as in the computation of our regressors to express trade size or activity using a stock-specific standard.

The goal of this analysis is characterising how the proportions of order flow executed on and off the order book are determined. Clearly, then, we need to ensure that the transactions included in our subsequent panel regression analysis are those for which the two trading mechanisms might actually compete. The first category of trades we exclude from the data are those which occur outside the hours of order book operation. These trades represent about 1.3% of the entire sample by frequency and are generally very large. It is interesting to note that a large proportion of these trades occur right after the book closes. Around 25% are reported in the minute after the book closes and 75% in the 30 minutes following book closure. This is consistent with these trades being prearranged with London dealers to take place at the closing price, as found in US data by Barclay and Hendershott (2000). Second, for reasons of sheer size and due to delayed trade publication arrangements, we remove all trades eligible for WPA classification (0.06% of trades by frequency; these may be single-stock as well as portfolio WPAs, defined as a trade effected in at least 20 stocks simultaneously, and where the total quantity traded must be at least 8 NMS). Finally, we remove all trades with non-standard settlement cycles as they would not have been eligible for book execution. These represent 26.5% of all trades, most very small.

In aggregate, this filtration of the data removes 27.6% of the trades by frequency and 21.8% by value, leaving about 770,000 trades in the data. Descriptive statistics on trading activity by market segment are given in Table 2. First note that almost exactly the same number of (filtered) trades occur in the two market segments. In traded quantity terms the share of the book in activity is lower at around 45%. We also see that book trades tend to be smaller than non-book trades (on average) and that the dispersion in book trade size is only around one third that of non-book trades.

[5] Roughly, a stock's NMS is calculated as follows. First, the exchange computes 2.5% of average (recent) daily traded volume. This average is then rounded to one of a set of pre-specified NMS values. Note, therefore, that a 1 NMS trade is very large.

The empirical quantiles of the book and off-book trade size distributions are given in the final rows of Table 2. Average trade size is much larger than median trade size for both book and off-book trades. The key thing to note from these quantiles, however, is that for all up to the 95th, the quantile of the book distribution is an order of magnitude larger than its non-book counterpart. On the other hand, the 95th percentile of off-book trade size is over 50% larger than that of book trades. The implication of these results is that off-book trades are generally very small but also contain a very large share of very big transactions. We would expect many of the very small trades to be executed off-book due to the RSP arrangements discussed in Section 2.1. Also, we would expect a high proportion of very large transactions to be handled off the book due to the likelihood that the terms on which they are executed would need to be pre-negotiated.

Finally for this section, in Table 3 we supply stock-by-stock evidence on the extent of order book usage. For the vast majority of the sample stocks, the share of the book in quantity traded (NMS) is between 40 and 55%. When looking at the number of trades routed to the book the comparable range is 50–70%. This clearly indicates that the book is an important source of liquidity for all of our stocks. The final columns of Table 3 give the quantiles of book and off-book trade size by stock. This evidence simply confirms that drawn from Table 2 on the *type* of trades routed to the dealers. Dealers tend to execute a very large share of both very small *and* very large trades.

3 Order flow interactions: panel analysis

3.1 Specification

We now turn to examination of the determinants of the market share of the limit order system. To detect whether order flow migrates from the book to the dealers in response to changing market conditions, we use a panel regression framework. The dependent variable in our analysis is the market share of the limit order book by number of shares traded, which lies in the $[0, 1]$ interval. We compute this market share over half-hour intervals of all trading days. This gives 17 intervals for each trading day (as SETS is available from 8:00 to 16:30 in our sample period). In cross-sectional terms, the final sample is made up of 28 stocks after removing two of the original sample.[6] The sample period covers 40 trading days, after removing the 24^{th} and 30^{th} of December during which the market was formally open but when trading activity was exceptionally low. This leaves 19,040 observations in the panel data set used for estimation.

We employ a 30 minute sampling frequency, rather than anything higher, for a couple of reasons. First of all, we are interested in when trading activity, in general, tends to be routed to one venue or another (rather than when a single trade might be routed to a given venue). Second, since off-book trades are potentially reported with a delay of up to three minutes, transaction level analysis of the hypotheses in which we are interested is not possible as the relevant variables cannot be accurately measured. For example, due to the reporting delay of off-book trades, at a given point in time the total trading activity over the last k minutes is not available to market participants as some of the trades occurring in this interval may not have been reported. For these two reasons, we conduct our analysis at a higher level of aggregation than is done in the switching regression studies based on Madhavan and Cheng (1997).

3.1.1 Choice of explanatory variables

Based on theory, evidence from descriptive statistics, and our own discussions with broker-dealers, we consider the following variables as potential determinants of switching of order flow between the two market segments:

1. Order book depth; the time-weighted average of the total quantity available on the bid and ask sides of the book (computed using data on all limit and market orders submitted in each half-hourly interval) expressed in NMS terms. Expected coefficient sign positive.

2. Cost of book trading; (time-weighted) average percentage inside spread from book quotes. Expected coefficient sign negative.

These variables are suggested by models of upstairs trading. Based on Demsetz (1968) and information-based models for off-book trades (Seppi, 1990; Pagano and Röell, 1992; Madhavan and Cheng, 1997) we also include the following regressors;

[6] We removed data on Tate & Lyle because it was a recent addition to SETS as of 1 November 1999 and is an outlier in term of traded volumes. We also removed Marconi because the company underwent significant organisational changes over these two months.

3. Order flow imbalance; the absolute value of the quantity of buyer-initiated trades less the quantity of seller initiated, in NMS terms.[7] Expected coefficient sign negative.

4. Average trade size; for each stock-interval, computed in NMS terms. Expected coefficient negative.

5. Total value traded: for each stock-interval, in NMS terms. Expected coefficient positive. To investigate possible non-linearities we include two extra regressors. Dummies are created that pick out the intervals in the lowest and highest quartiles of the traded value distribution taken across all stocks, respectively, and these dummies are then interacted with value itself.

The predictions in Grossman and Miller (1988) and the "free option" analogy of limit orders (Copeland and Galai, 1983) suggest inclusion of the following:

6. Realized Volatility; first lag of the square root of the sum of the squared 5-minute transaction price returns within a 30-minute interval (Andersen, Bollerslev, Diebold, and Labys, 1999). We lag volatility to avoid picking up a mechanical correlation between book activity and book midquote return volatility. We also add an interaction term to account for the overnight breaks in the sample.[8] Expected coefficient sign negative.

Finally, we include several dummies to take account of trades routed to RSPs and to account for low book usage at the open;

7. Two "small trade" dummies; for all stocks and time bins together, we examine the distribution of trade sizes for trades smaller than or equal to one-tenth of

[7]This simple measure is similar to that used in Blume, MacKinlay, and Terker (1989). Note that construction of this measure requires us to sign trades. The direction of order book trades is already properly assigned in our data (identified by the direction of the order hitting the limit outstanding), but the direction of dealer trades is less clear. We sign dealer trades in two different ways. First, we make use of the "dealing capacity" flags in the data which indicate whether the broker-dealers who were party to a trade were acting as agents or for their own portfolio. We adopt the convention that the direction of the trade will be that of the party acting as principal, since they are the ones who have a mandate to trade and are therefore likely to be the "aggressor". Alternatively, we sign trades based on the distance from the quotes. The two methods yield very similar results.

[8]To be clear, we create an interaction of a dummy picking out intervals that are the first of the trading day with the lagged volatility regressor. In this way we account for a different relationship between book share and lagged volatility from the same trading day and book share in the first interval of a day and volatility from the final interval of the preceding day.

an NMS. We use the median value (0.005 NMS) and the 95th percentile (0.05 NMS) as cutoff points to define "tiny" and "very small" trades, respectively. If, within a half-hourly bin, the proportion of the value traded accounted for by each of these two categories of trade exceeds a "usual" level (the mean, computed across days and time bins, within stocks) by one standard deviation, the dummy equals one for that bin. Expected coefficient signs negative, and the coefficient on the "very small" dummy should be greater than that on the "tiny" dummy.

8. "Opening" dummy variable; unity in the first 30 minutes of each trading day to account for low level of book usage at the open. Expected coefficient sign negative.

3.1.2 Choice of estimator

The selection of a panel estimator is done in very standard fashion. The first step is to determine whether persistent individual effects are present in the data or if the time series can simply be pooled. The null hypothesis that a single constant term is adequate is rejected by an LM test. Having detected individual effects, we test whether they can be modelled as random or whether they are correlated with the regressors. A Hausman test rejects the hypothesis that the difference in estimated coefficients between random and fixed effects is not systematic, and we therefore fit a fixed-effects model as our baseline specification. We extend the baseline model in a number of directions subsequently.

3.2 Results

Our panel regression results are presented in table 4. A first general comment is that all explanatory variables appear to be significant, most of them strongly so, since the significance levels given in the table are for two-tailed tests, whereas we have clear expectations regarding the signs of most of these coefficients. The only coefficient that is not significantly different from zero is that on the volatility interaction term (and hence in all further estimations we drop this variable). A second point to note is that all coefficients for which we had clear priors from theory and intuition have the expected sign.

It appears from our results that the share of the order book in transaction activity drops in times of market turbulence, as measured by volatility or imbalances in the value of market order flows. The first of these results is consistent with the notion that dealership services are of greater value (and are more heavily employed) when execution risk is high, although the estimated coefficient on this variable is only marginally significant at the 5 % level. The latter result is consistent with a view of dealers as providers of immediacy when trading conditions are unusual and, as in Demsetz (1968), when an order book is unlikely to perform well due to lack of two-sided trading interest.

As conjectured, we uncover a non-linear relationship between trading volume and the book's share of activity. When volume is at low or medium levels, increased volume is associated with increased book usage. This result is consistent with the prediction implicit again in Demsetz (1968) that continuous auction markets should work better when activity is high. However, when trading volume is extremely high (i.e. within the top quartile) the relationship between volume and order book activity is negative. We believe that this result is due to the fact that very large trades are essentially always done off the order book (as the book is unlikely to contain enough natural liquidity to accommodate them) and extremely high volume intervals are those containing such big trades.

Order book trading activity is strongly linked to book liquidity as liquidity or information-based models of upstairs trading such as Grossman (1992) or Seppi (1990) suggest. Specifically, book share is low when book liquidity is low, as measured either by high spreads or low depth. Hence, when the book terms of trade are poor those demanding liquidity look elsewhere for more favourable execution. This result strongly suggests that off-book activity is not simply driven by preferencing and internalisation arrangements as these would imply that the order flow sent to both market segments should behave as separate processes, and off-book share of activity was insensitive to the state of book liquidity.

Finally, we see that the frequency of appearance of very small and tiny trades at the market tends to reduce the order book share. This is likely due to the RSP effect discussed in Section 2.1. Also, after controlling for such effects, order book share is negatively related to average trade size within an interval, consistent with the predictions drawn from Demsetz (1968) as well as from Burdett and O'Hara (1987).

In sum, we see that in "unusual" conditions (imbalanced order flow for example), dealers become relatively more active. We also observe increased dealer activity when book liquidity is low (as measured by either spreads or depth). Taken together our results seem to imply that the dealers voluntarily perform a liquidity provision and a stabilisation role. When book liquidity is deficient and trading conditions are unusual they step in and trade alongside the book.

3.3 Endogeneity and fragmentation

We have interpreted the previously reported results using a number of priors on directions of causality between order book share and other variables, and associated priors on signs of these effects. Our results broadly conform with these priors. An important issue remains to be addressed, however.

When multiple trading venues are available for a given stock, the liquidity of each venue will depend upon the level of trading activity it attracts. If, in our example, investors route trades off-book when they deem book liquidity to be insufficient this may in turn discourage agents from supplying liquidity to the book. Hence, the order flow sent to a venue is determined simultaneously with that venue's liquidity. The econometric worry that arises from this intuition is that the liquidity variables on the right-hand side of our regression specification are likely to be endogenous.

To address the endogeneity issue, we employ an instrumental variable estimation technique in panel format, using two lags of the spread and depth regressors respectively as instruments for their own contemporaneous values. The choice of the number of lags to use as instruments was determined by examining the time-series dependence in these variables on a stock-by-stock basis. However, using a different number of lags between 1 and 10 does not change the results in any important way.

Estimation results are reported in table 5. The majority of them are consistent with those we found in previous estimations. The only major difference is that the coefficient on the volatility regressor loses significance. Hence, across the estimations we have performed and given the other conditioning variables we have included, it seems that the effect of volatility on the share of the order book in trade is not clear. Again, we suspect that this is due to those reasons mentioned above: the effect of volatility on book share operates via an effect on the book liquidity variables that also appear in the specification. Consistent with this view, if one re-estimates any of our specifications with depth and spreads omitted from the right-hand side, the volatility regressor is always negative and strongly significant.

More importantly, we see that the liquidity (spread and depth) regressors retain both their signs and significance (as do all the remaining regressors). This would suggest that our previous results are not driven by the effects of fragmentation on book liquidity but are really due to poor order book liquidity leading to those demanding immediacy looking to other trading venues.

3.4 Other robustness checks

In order to further investigate the robustness of our results we performed a number of diagnostic checks on the regression residuals. A first examination indicated cross-sectional heteroskedasticity and residual correlation. We model this more carefully in the next subsection but, for now, note that a GLS estimator robust to this and to group-wise heteroskedasticity yielded estimated coefficients and standard errors which were extremely comparable to the ones produced by the simple fixed-effects specification, implying that while statistically present, these problems must be mild.[9]

A second diagnostic check was to examine the time-series dependence in residuals. We found some evidence of residual autocorrelation (mainly at very low displacements) on a stock-by-stock basis. Computation of the panel autocorrelation test statistic proposed by Baltagi and Wu (1999) indicated that the presence of autocorrelation could not be rejected at the 5% level of significance. Thus we investigated the impact of this autocorrelation on our results in two ways. First we estimated a fixed effects specification which explicitly allowed for AR(1) errors. Second we estimated a fixed-effects specification including a lagged dependent variable on the right-hand side (using the estimation technique suggested by Anderson and Hsiao (1981)). The results from both of these estimations were qualitatively identical to those reported in Table 4. (In particular, the estimated coefficient on the lagged dependent variable as a regressor did not appear significant).

Our final concern was over the nature of our dependent variable. As a market share, it is bounded at zero and one, and there are inevitably a small number of observations at both of the censoring points (with observations at zero more frequent than at unity), mainly for the less liquid stocks in our sample.[10] To account for this, we fitted a

[9]The GLS results and those mentioned but not reported later in this section are available upon request from the authors.

[10]Economically speaking, and largely for the least liquid stocks in the sample, there are half-hour

fixed-effects panel Tobit estimator with double censoring (at 0 and 1). The results from this estimation are economically and statistically very similar to those in Tables 4 and 5 and thus are not presented to conserve space.[11] Finally, as an alternative way to account for the bounded nature of the dependent variable, we applied a logistic transform such that the book share was mapped from the unit interval on to the entire real line. Again, when estimating a fixed effects panel specification, this did not significantly affect the sign or statistical significance of the parameter estimates.

4 Commonality in book liquidity and order book usage

Our panel regression analysis has thus far demonstrated that order book utilisation in trade is strongly linked to a number of stock-specific factors which, on a theoretical level, can themselves be linked to execution and information risks. None of the explanatory factors that we have considered thus far impart market-wide effects to the stock-by-stock determination of order book usage. Recent work (Hasbrouck and Seppi, 2001; Chordia, Roll, and Subrahmanyam, 2000) points to the existence of strong covariation in liquidity and trading activity across stocks, patterns which are likely to generate cross-stock commonality in the attractiveness of order book versus off-book trading. Moreover, the existence of market-level informational shocks or execution risk (via market trading activity or volatility for example) will, by similar arguments to those presented in Section 1.1, lead to a likely trader preference for off-book trading on an individual stock basis.

Thus, in this section we test for the existence of common factors in order book liquidity measures, checking the evidence presented in Hasbrouck and Seppi (2001) and Chordia, Roll, and Subrahmanyam (2000) for our London data. More pertinently in the current context, though, we test for commonality in order book usage across stocks. This analysis gives us a pre-test for the possible dependence of book utilisation on market-wide risk factors. We employ two methodologies to test for commonality. First, as in Hasbrouck and Seppi (2001), we use a standard principal components analysis (PCA). Second, as in Chordia, Roll, and Subrahmanyam (2000), we also

intervals in which trading activity is concentrated entirely amongst the dealers or on the order book. Across all stocks, less than 5% of observations are at the upper or lower censoring point.

[11]These results are available upon request from the authors.

carry out "market model" regressions for our variables of interest. These are stock-by-stock time-series regressions of our variable of interest on a concurrent market-wide value for that variable, constructed as the average across all stocks (aside from the stock currently under analysis).

Panel (a) of Table 6 presents the results from our PCA. Note that in this analysis we have, where appropriate, removed intra-day patterns from each variable and standardised variance to unity. What the Table presents for each series is the proportion of total variation in liquidity and order book usage across stocks that can be explained by each of the first three principal components. Our two book liquidity variables are the book depth and spread variables defined in Section 3.1.1. We also perform the computation for the order book share in trade.

Our results indicate very strong commonality across stocks in both book liquidity and order book utilisation. For the depth measure, for example, the first principal component explains almost 50% of all variation. This is far higher than the comparable number in Hasbrouck and Seppi (2001), who state in the conclusion of their paper that "It is perplexing that common-factor variables do not on balance play a more prominent and unambiguous role given the importance traders attach to the proximity of other traders." We can only speculate as to why we find stronger effects, but one reason may be that we are able to construct a less noisy depth proxy because we use data on all quantities to buy and sell submitted in the order book in the construction of our depth measure, whereas Hasbrouck and Seppi only use the best bid and ask. Spreads have smaller indications of commonality in our data, but even here the first principal component explains close to 15% of variation. Our most interesting and novel result, however, is that there is also a very strong indication of commonality in order book use. The order book share variable has a first principal component that explains 12% of total variation. Thus, there is strong evidence that trading tends to migrate from the dealer segment of the market to the order book (and vice-versa) at similar times across our sample stocks.[12] The key question that arises then, is what drives this commonality? Can we identify market level influences on stock-by-stock order book utilisation?

[12]Using aggregate figures from an NYSE Annual Report, Hasbrouck and Seppi (2001) argue that common components in orders are not likely to be driven by large program trades. We can rule this out more directly, since large program trades, eligible for WPA status, are filtered out in our data before the computation of the dependent variable (see Section 2.2).

4.1 Market-wide factors in the panel framework

The presence of common factors in measures of book liquidity on the right-hand side of our previous panel equation and in book use on the left-hand side across securities strongly suggests that we extend our previous analysis to allow for market-wide activity and liquidity influences. We construct two market-wide regressors to account for the commonalities. First, we construct a market-wide regressor based on the Demsetz (1968) execution risk intuition — when at the market level trading is heavy one would expect individual stocks to trade more on the book due to increased likelihood of buying and selling interest matching naturally. Second, we use a market-wide order flow imbalance regressor to capture market-wide information risk and the intuition that when market buying and selling pressure in one-sided, this is likely to result in one-sided stock-level order flow and thus reduced usage of the SETS order book. We compute these market-level regressors from FTSE-100 futures contract data, a very actively traded asset. Our proxy for market activity is a simple count of the number of futures contracts traded in every half-hour interval. To this, we add a measure of market-wide order flow imbalance, computed in identical fashion to our stock-by-stock variable i.e. taking the absolute value of buy minus sell trades, in number of contracts traded.

Table 7 reports the estimation results. Both market regressors come out clearly significant at the 1% level, while their signs indicate that SETS market share is higher when overall market activity is high, and tends to be lower when market order flow is one-sided. These results are very stable across all three previously used estimation methods (although only the baseline fixed-effects specification is reported here for brevity). It also appears that including the market activity and imbalance regressors somewhat reduces the significance of the stock-specific ones in the estimation. All other right-hand side variables remain comparable in magnitude and significance to what they were previously, indicating that market-wide factors have separate power in explaining the choice of trading venue.

For interpretation of these results, we evoked reasons related to portfolio-wide liquidity shocks or market-wide information risk at the end of section 1.2 which could drive order flow off the order book across securities at certain times. Current research is developing the idea of "systematic liquidity" as a priced factor, an idea originating in Amihud and Mendelson (1986). Easley, Hvidkjaer, and O'Hara (2002) argue on the other hand that measures of liquidity which appear to play a role in pricing are in fact only proxying for information risk. The current results suggest that both market-wide liquidity and information risks have implications for trading system use (and, beyond this, design).

5 Summary and conclusion

We study a hybrid trading system for very liquid stocks, where both limit order trading and broker-dealer services are available throughout the trading day, and liquidity suppliers may compete on and off the order book, without particular privileges. At an intra-day frequency, a panel regression analysis demonstrates that the choice of where trades are executed is sensitive to market conditions and the state of the order book. Our results imply that the share of the order book in trading activity is increasing in its liquidity (as measured by low spreads and large depth) and decreasing in average trade size and imbalances in liquidity demand. There is some evidence that increased volatility of returns tends to reduce the market share of the order book. Finally, for low and medium levels of trading volume, volume and book usage are positively related. For very high volume levels there is a negative relationship between these variables.

This analysis sheds some light on what underlies the demand for dealership services in an order-driven environment, and why real-world market structures are hybrid. Our results suggest a natural degree of interaction between a liquid order book and a network of broker-dealers, with no obvious need for rules designed to force consolidation. Moreover, our results do not support the view that there is a need for designated intermediaries who must be granted a variety of privileges (in particular informational ones) against obligations to "do good", consistent with the conjecture in Stoll (1998). Instead, they indicate that dual-capacity firms voluntarily take on a stabilizing role, absorbing order imbalances or meeting very high demands for immediacy in the presence of an order book without formal regulatory commitment to do so.

Finally, we uncover evidence that the utilization of on and off book trading opportunities on an individual stock level is correlated across stocks and is moreover linked to the evolution of market-wide variables, specifically trading activity and order flow imbalances. If market-level trade is heavy then individual stocks tend to trade on the order book more frequently. If, however, market-level order flow tends to be one-sided then this reduces the utilization of the order book. We interpret these results as pointing to the importance of market-wide liquidity and information shocks in the determination of stock-level activity and order flow distribution.

It is generally accepted that less-liquid stocks are less suitable for order-driven trading, either for liquidity or informational reasons (Easley, Kiefer, O'Hara, and Paperman, 1996), but our findings illustrate the need to consider liquidity dynamically: there are times when execution or information risks are high for various reasons - markets could be slow or the composition of order flow between informed and uninformed participants could change. At those times, even stocks that are the most actively traded on average may behave like less-liquid issues, requiring intermediaries to smooth the trading process.

References

Amihud, Y., and H. Mendelson, 1986, "Asset Pricing and the Bid-ask Spread," *Journal of Financial Economics*, 17, 223–249.

Andersen, T., T. Bollerslev, F. Diebold, and P. Labys, 1999, "The Distribution of Exchange Rate Volatility," Working Paper 99-08, Financial Institutions Center, The Wharton School, University of Pennsylvania.

Anderson, T., and C. Hsiao, 1981, "Estimation of Dynamic Models with Error Components," *Journal of the American Statistical Association*, 76, 598–606.

Baltagi, B. H., and P. X. Wu, 1999, "Unequally Spaced Panel Data Regressions with AR(1) Disturbances," *Econometric Theory*, 15(6), 814–23.

Barclay, M. J., W. G. Christie, J. H. Harris, E. Kandel, and P. H. Schultz, 1999, "The Effects of Market Reform on the Trading Costs and Depths of Nasdaq Stocks," *Journal of Finance*, 54(1), 1–34.

Barclay, M. J., and T. Hendershott, 2000, "Price Discovery and Trading Costs After Hours," Working Paper.

Bessembinder, H., and K. Venkataraman, 2002, "Does and Electronic Stock Exchange Need an Upstairs Market?," Working Paper.

Blume, M. E., A. MacKinlay, and B. Terker, 1989, "Order Imbalances and Stock Price Movements on October 19 and 20, 1987," *Journal of Finance*, 44(4), 827–848.

Board, J., and S. Wells, 2000, "Liquidity and Best Execution in the UK: A Comparison of SETS and Tradepoint," Paper presented at June 2000 FMG conference "The Future of Exchanges".

Booth, G., J.-C. Lin, T. Martikainen, and Y. Tse, 2002, "Trading and Pricing in Upstairs and Downstairs Stock Markets," *Review of Financial Studies*, 15(4), 1111–1135.

Burdett, K., and M. O'Hara, 1987, "Building Blocks," *Journal of Banking and Finance*, 11(3), 193–212.

Chordia, T., R. Roll, and A. Subrahmanyam, 2000, "Commonality in Liquidity," *Journal of Financial Economics*, 56, 3–28.

Copeland, T. E., and D. Galai, 1983, "Information Effects on the Bid-Ask Spread," *Journal of Finance*, 38(5), 1457–1469.

de Jong, F., T. Nijman, and A. Röell, 1995, "A Comparison of the Cost of Trading French Shares on the Paris Bourse and on SEAQ International," *European Economic Review*, 39(7), 1277–1301.

Demarchi, M., and T. Foucault, 1999, "Equity Trading Systems in Europe (A Survey of Recent Changes)," HEC School of Management Mimeo.

Demsetz, H., 1968, "The Cost of Transacting," *Quarterly Journal of Economics*, 82(1), 33–53.

Easley, D., S. Hvidkjaer, and M. O'Hara, 2002, "Is Information Risk a Determinant of Asset Returns?," *Journal of Finance*, 57(4).

Easley, D., N. Kiefer, and M. O'Hara, 1996, "Cream-Skimming or Profit-Sharing? The Curious Role of Purchased Order Flow," *Journal of Finance*, 51(3), 811–833.

Easley, D., N. Kiefer, M. O'Hara, and J. B. Paperman, 1996, "Liquidity, Information, and Infrequently Traded Stocks," *Journal of Finance*, 51(4), 1405–1436.

Ellul, A., 2001, "As you Like it: an Investigation of Trading Behaviour and Price Volatility on Auction and Dealership Market Architectures," Working paper.

Fong, K., A. Madhavan, and P. Swan, 2001, "Why do Markets Fragment? A Panel-Data Analysis of Off-Exchange Trading," University of Sydney Working Paper.

Grossman, S. J., 1992, "The Informational Role of Upstairs and Downstairs Trading," *Journal of Business*, 65(4), 509–528.

Grossman, S. J., and M. H. Miller, 1988, "Liquidity and Market Structure," *Journal of Finance*, 43(3), 617–633.

Hasbrouck, J., and D. Seppi, 2001, "Common Factors in Prices, Order Flows, and Liquidity," *Journal of Financial Economics*, 59(3), 383–411.

Madhavan, A., and M. Cheng, 1997, "In Search of Liquidity: Block Trades in the Upstairs and Downstairs Markets," *Review of Financial Studies*, 10(1), 175–203.

Madhavan, A., and G. Sofianos, 1998, "An Empirical Analysis of NYSE Specialist Trading," *Journal of Financial Economics*, 48(2), 189–210.

Madhavan, A. N., 1992, "Trading Mechanisms in Securities Markets," *Journal of Finance*, 47(2), 607–642.

Miller, M. H., 1991, "Liquidity and Market Structure," in *Financial Innovations and Market Volatility*, ed. by M. H. Miller. Basil Blackwell, Oxford, chap. 2, pp. 23–31.

Naik, N., and P. Yadav, 1999, "The Effects of Market Reform on Trading Costs of Public Investors: Evidence from the London Stock Exchange," London Business School IFA Working Paper 296.

Pagano, M., and A. Röell, 1992, "Auction and Dealership Markets. What is the difference?," *European Economic Review*, 36, 613–623.

——— , 1996, "Transparency and Liquidity: A Comparison of Auction and Dealer Markets with Informed Trading," *Journal of Finance*, 51, 579–611.

Röell, A., 1992, "Comparing the Performance of Stock Exchange Trading Systems," in *The Internationalisation of Capital Markets and The Regulatory Response*, ed. by J. Fingleton, and D. Schoenmaker. Graham and Trotman, London, chap. 8, pp. 167–177.

Seppi, D., 1990, "Equilibrium Block Trading and Asymmetric Information," *Journal of Finance*, 45(1), 73–94.

Smith, B. F., A. S. Turnbull, and R. W. White, 2001, "Upstairs Market for Principal and Agency Trades: Analysis of Adverse Information and Price Effects," *Journal of Finance*, 56(5), 1723–1746.

Stoll, H. R., 1993, "Organization of the Stock Market: Competition or Fragmentation?," *Journal of Applied Corporate Finance*, 5(4), 89–93.

——— , 1998, "Reconsidering the Affirmative Obligation of Market Makers," *Financial Analysts Journal*, Sept.-Oct., 72–82.

Viswanathan, S., and J. Wang, 2002, "Market Architecture: Limit Order Books Versus Dealership Markets," *Journal of Financial Markets*, 5, 127–167.

Table 1: FT-30 stock information

Company name	MNEM	MV	NMS
Allied Domecq	ALLD	3537.75	75000
British Aerospace	BA.	10753.04	150000
British Airways	BAY	3955.55	100000
Blue Circle Industries	BCI	2855.25	50000
BG	BG.	13222.55	150000
BOC Group	BOC	6402.94	25000
Boots Co	BOOT	5409.26	75000
BP Amoco	BPA	123418.2	200000
British Telecommunications	BT.A	86236.94	200000
Cadbury Schweppes	CBRY	7889.84	75000
Diageo	DGE	19319.13	150000
Emi Group	EMI	4164.57	50000
Granada Group	GAA	9516.28	75000
GKN	GKN	7089.22	50000
Glaxo Wellcome	GLXO	67835.63	100000
Imperial Chemical Industries	ICI	4553.44	75000
Invensys	ISYS	11106.82	200000
Lloyds TSB Group	LLOY	44861.1	200000
Marks and Spencer	MKS	7079.25	200000
Marconi	MNI	24497.75	150000
National Westminster Bank	NWB	24183.61	75000
Peninsular and Orient Steam Nav. Co.	PO.	6116.83	50000
Prudential Corp	PRU	20079.36	75000
Royal and Sun Alliance Insurance Group	RSA	5668.35	100000
Reuters Group	RTR	9664.71	75000
Smithkline Beecham	SB.	47222.37	200000
Scottish Power	SPW	10224.98	100000
Tate and Lyle	TATE	1923.37	25000
Tesco	TSCO	11884.57	200000
Vodafone Airtouch	VOD	93969.44	150000

Notes: The table gives basic information on the components FT-30 index. The stocks are drawn from various deciles of liquidity within the FTSE-100 index. The second column indicates the stock mnemonics used in the tables that follow. The third column indicates the component stocks' market values (in millions of pounds as at December 1, 1999). The final column gives each stocks' "Normal Market Size" as calculated by the London Stock Exchange, a measure of the average size (in number of shares) of an institutional trade in the stock concerned, computed as 2.5% of average daily trading volume in a given stock, then rounded to the nearest of a set of pre-specified NMS bands.

Table 2: Summary statistics by order venue

Statistic	Dealer trades	Book trades
Number of trades	383,882	383,366
Total traded value	33103.10	26906.19
Average traded value	0.086	0.070
SD of traded value	0.372	0.108
Total traded quantity	43926.06	35036.42
Average traded quantity	0.114	0.091
SD of traded quantity	0.476	0.160
Min of trade quantity	5×10^{-6}	5×10^{-6}
Max of traded quantity	8.00	7.13
Empirical quantity distribution		
5^{th} percentile	0.0004	0.0025
25^{th} percentile	0.0020	0.0132
50^{th} percentile	0.0056	0.0397
75^{th} percentile	0.0221	0.1085
95^{th} percentile	0.5000	0.3333

Notes: the table reports statistics on the distribution of trade size in the dealer and order book segments of the market, after filtering out all trades with size greater than 8 NMS, all trades taking longer than 5 business days to settle and all trades reported outside SETS hours. Turnover by value statistics are in millions of pounds, traded quantity statistics are in terms of NMS as defined previously.

Table 3: Trade breakdown by stock and venue

Stock	Book trade prop^n			Dealer trade size			Book trade size		
	Num.	Quant.	Vol.	25 %	50 %	75 %	25 %	50 %	75%
ALLD	0.615	0.450	0.450	0.007	0.014	0.080	0.018	0.058	0.187
BA.	0.490	0.359	0.360	0.003	0.006	0.025	0.012	0.034	0.100
BAY	0.538	0.466	0.463	0.003	0.007	0.020	0.009	0.031	0.100
BCI	0.675	0.414	0.410	0.009	0.025	0.168	0.014	0.052	0.133
BG.	0.550	0.445	0.444	0.003	0.007	0.027	0.017	0.048	0.115
BOC	0.741	0.522	0.522	0.012	0.040	0.329	0.020	0.049	0.188
BOOT	0.431	0.336	0.337	0.004	0.010	0.027	0.014	0.031	0.102
BPA	0.604	0.567	0.568	0.002	0.005	0.025	0.017	0.050	0.125
BT.A	0.311	0.456	0.456	0.001	0.002	0.003	0.007	0.020	0.050
CBRY	0.601	0.461	0.463	0.009	0.021	0.117	0.029	0.071	0.200
DGE	0.594	0.426	0.426	0.003	0.008	0.043	0.013	0.037	0.093
EMI	0.674	0.282	0.281	0.008	0.027	0.333	0.012	0.026	0.120
GAA	0.501	0.414	0.413	0.009	0.019	0.078	0.032	0.085	0.206
GKN	0.511	0.361	0.360	0.006	0.010	0.030	0.013	0.032	0.105
GLXO	0.514	0.485	0.484	0.001	0.003	0.009	0.007	0.017	0.040
ICI	0.612	0.495	0.493	0.004	0.009	0.028	0.012	0.029	0.093
ISYS	0.587	0.394	0.392	0.004	0.009	0.050	0.012	0.037	0.100
LLOY	0.491	0.400	0.398	0.002	0.004	0.013	0.009	0.025	0.054
MKS	0.355	0.402	0.402	0.002	0.005	0.015	0.015	0.039	0.100
MNI	0.432	0.342	0.349	0.002	0.006	0.020	0.013	0.043	0.101
NWB	0.576	0.438	0.439	0.002	0.005	0.040	0.014	0.040	0.100
PO.	0.630	0.410	0.412	0.006	0.015	0.077	0.017	0.040	0.120
PRU	0.523	0.379	0.379	0.003	0.006	0.021	0.011	0.031	0.080
RSA	0.627	0.489	0.491	0.009	0.026	0.182	0.030	0.093	0.266
RTR	0.560	0.404	0.401	0.004	0.010	0.050	0.012	0.034	0.100
SB.	0.515	0.477	0.479	0.002	0.004	0.013	0.009	0.025	0.059
SPW	0.541	0.509	0.510	0.004	0.008	0.025	0.017	0.050	0.180
TATE	0.557	0.480	0.483	0.020	0.051	0.247	0.048	0.186	0.400
TSCO	0.496	0.410	0.410	0.005	0.011	0.049	0.022	0.061	0.160
VOD	0.500	0.478	0.478	0.005	0.011	0.043	0.025	0.091	0.247

Notes: we have filtered all trade with size greater than 8 NMS, all trades taking longer than 5 business days to settle and all trades reported outside SETS hours. The first three columns report the proportion of trades executed on the SETS order book, by number of shares traded, quantity (in NMS), and value traded. The remaining six columns report statistics on the distribution of trade size in the dealer and order book segments of the market, expressed in NMS.

Table 4: Baseline panel regression analysis of the share of the order book in total order flow

Variable	Coefficient
Lagged volatility	-0.2003
	(1.90)
Volatility interaction dummy	-1.4702
	(1.18)
Opening dummy	-0.0364
	(2.40)*
Total value traded	0.0010
	(3.33)**
High value traded interaction dummy	-0.0014
	(5.44)**
Low value traded interaction dummy	0.0355
	(16.93)**
Mean trade value	-0.6460
	(40.07)**
Imbalance	-0.0036
	(3.50)**
Spread	-2.6557
	(14.33)**
Depth	0.0029
	(7.86)**
V small trades dummy	-0.0619
	(13.00)**
Tiny trades dummy	-0.1311
	(18.08)**
Constant	0.5964
	(110.82)**
Observations	19040
Number of stocks	28
R-squared	0.1733

Notes: the table reports the coefficients estimated from a fixed-effects panel model, where the dependent variable is the market share of the order book (by number of shares traded). Laggedvol is our lagged volatility regressor. We add a volatility interaction dummy for the first half-hour of trading, and a dummy variable which is unity for the first interval of each trading day (Opendum). TradedVal is the total value traded regressor. MeanVal is the mean transaction value regressor. Imbalance is our measure of order flow imbalance. Spread is the time-weighted book spread and Depth the time-weighted book depth. "V small trades dummy" and "Tiny trades dummy" are dummy variables for intervals containing large numbers of small trades. See Section 3.1.1 for detailed definitions of these regressors. Absolute value of t-statistics in parentheses. A * denotes that the coefficient is significantly different from zero at the 5% level for a two-tailed test and a ** denotes significance at the 1% level.

Table 5: Panel IV analysis of the share of the order book in total order flow

Variable	Coefficient
Lagged volatility	0.0270
	(0.24)
Opening dummy	0.1026
	(5.50)**
Total value traded in NMS	0.0000
	(0.14)
High value traded interaction dummy	-0.0006
	(2.07)*
Low value traded interaction dummy	0.0321
	(14.51)**
Mean trade value	-0.6589
	(39.01)**
Imbalance	-0.0038
	(3.50)**
Spread	-10.0258
	(12.92)**
Depth	0.0024
	(6.00)**
V small trades dummy	-0.0575
	(11.51)**
Tiny trades dummy	-0.1120
	(14.34)**
Constant	0.6433
	(88.95)**
Observations	19040
Number of stocks	28
R-squared	0.1182

Notes: the table reports the coefficients from a fixed-effects panel model, where the dependent variable is the market share of the order book (by number of shares traded). See the notes to Table 4 and Section 3.1.1 for definitions of the regressors. Estimation is by instrumental variables, instrumenting spreads and depth using their first two lags. Absolute value of t-statistics in parentheses. A * denotes that the coefficient is significantly different from zero at the 5% level for a two-tailed test and a ** denotes significance at the 1% level.

Table 6: Principal components analysis and market models for order book depth, spreads and market share

Panel (a): Principal Components Analysis

Variable	PC1	PC2	PC3
Time-weighted spreads	0.492	0.081	0.061
Total depth	0.142	0.056	0.048
Order book share of order flow	0.116	0.047	0.044

Panel (b): Market model regressions

Variable	$\hat{\beta}$	t-value	R^2
Time-weighted spreads	0.959	28.453	0.452
Total depth	0.751	6.140	0.083
Order book share of order flow	0.700	6.867	0.058

Panel (a) of the table presents, for order book spreads, depth, and market shares (as defined in Section 3.1.1), the proportion of total variation explained by the first three principal components. These are calculated as the ratio of the largest three eigenvalues of the covariance matrix of the given variable to 28, the numbers of stocks we study. All variables have had their repetitive intra-day patterns filtered prior to analysis and have been rescaled so that their variance is unity. Panel (b) presents results from stock-by-stock time-series regressions of liquidity and order book usage variables on market-wide liquidity or order book usage. The latter is defined as the cross-sectional average of the relevant variable, excluding the stock currently under analysis from the cross-section. The table presents the cross-sectional average of the slope coefficient, it's t-value and the regression R^2 for each variable.

Table 7: Fixed-effects panel estimation of the share of the order book in total order flow including market-wide activity and imbalance variables

Variable	Coefficient
Lagged volatility	-0.2373
	(2.26)*
Opening dummy	-0.0796
	(8.28)**
Total value traded	0.0006
	(2.03)*
High value traded interaction dummy	-0.0014
	(5.38)**
Low value traded interaction dummy	0.0373
	(17.70)**
Mean trade value	-0.603
	(37.16)**
Imbalance	-0.0030
	(2.90)**
Spread	-2.6790
	(14.50)**
Depth	0.0032
	(8.59)**
V small trades dummy	-0.0625
	(13.10)**
Tiny trades dummy	-0.1249
	(17.15)**
Market activity	3.58×10^{-5}
	(11.05)**
Market imbalance	-1.77×10^{-5}
	(3.22)**
Constant	0.5804
	(100.61)**
Observations	19040
Number of stocks	28
R-squared	0.1826

Notes: the table reports the coefficients estimated from a fixed-effects panel model, where the dependent variable is the market share of the order book (by number of shares traded). Regressors down to "Tiny trades dummy" are defined as previously (see the notes to Table 4 and Section 3.1.1 for detailed definitions). "Market activity" and "Market imbalance" measure total, market-wide trading activity and imbalance in order flow, respectively. See Section 4 for detailed definitions. Absolute value of t-statistics in parentheses. A * denotes that the coefficient is significantly different from zero at the 5% level for a two-tailed test and a ** denotes significance at the 1% level.

Figure 1: Intra-day Pattern in Share of Transaction Activity on SETS

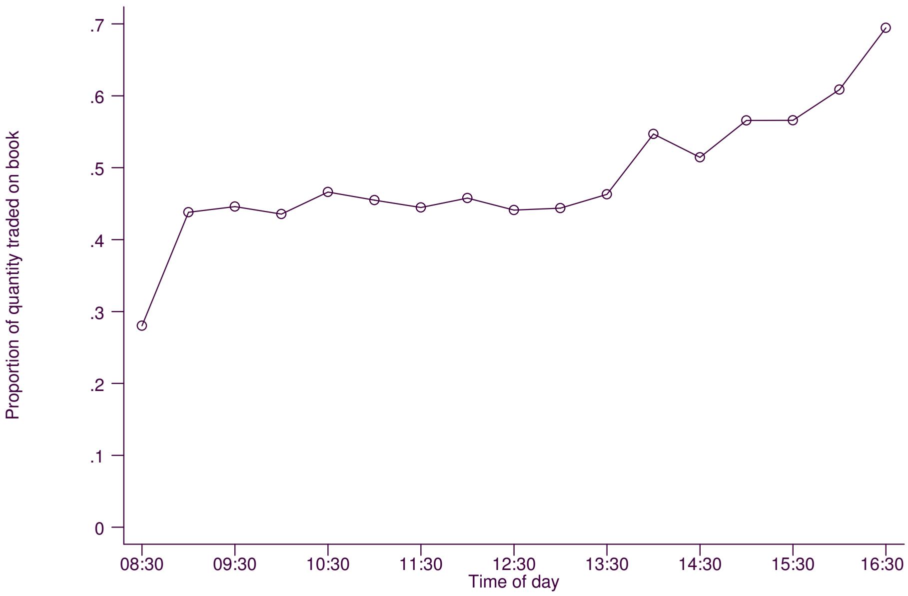

(a) Book share in number of trades

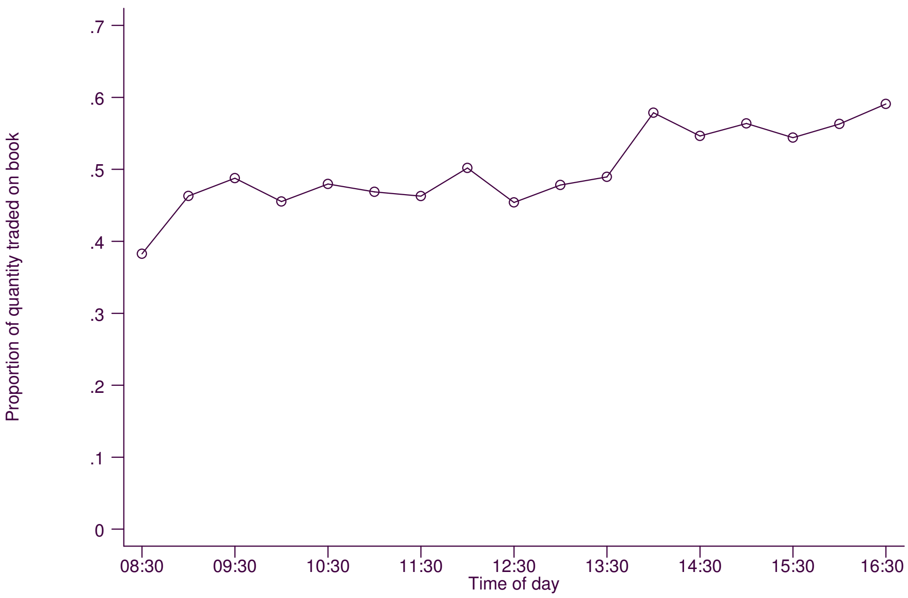

(b) Book share in traded value