

Splitting Orders in Fragmented Markets

Evidence from Cross-Listed Stocks

Albert J. Menkveld^a

Abstract

A number of recent theoretical studies have explored trading in fragmented markets, e.g. Biais et al. (2000), a phenomenon increasingly witnessed in modern markets. The key assumption generating the results is that there is at least one liquidity demander exploiting access to all markets by optimally splitting orders across markets. This paper seeks to test this assumption in a natural experiment involving Dutch stocks that are traded both in Amsterdam and New York. The results confirm the presence of rational, order splitting traders. This explains the increased volume and relatively large and persistent price changes for the overlapping period.

a) Tinbergen Institute, Erasmus University Rotterdam
Visiting Researcher GSB Stanford University

JEL Codes: G1, G15, G14, G12

Address correspondence, preferably via E-mail, to Albert J. Menkveld, Erasmus University Rotterdam, Tinbergen Institute, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands, phone: + 31 10 4088900, fax: + 31 10 4089031, E-mail: menkveld@few.eur.nl. I am particularly grateful to Marshall Blume for our stimulating discussions. Gratefully acknowledged is the support from the Tinbergen Institute, KLM Royal Dutch Airlines, the Rodney L. White Center at the Wharton School, the Graduate School of Business at Stanford University, the IIE for a Fulbright grant, the Amsterdam Exchanges and the NYSE. Financial support is gratefully acknowledged from the D.N. Chorafas Foundation, Erasmus Center for Financial Research, Fundatie van de Vrijvrouwe van Renswoude, Moret Fonds, Stichting Dr Hendrik Muller's Vaderlandsch Fonds, Stichting Fonds voor de Geld- en Effectenhandel, Stichting Hoogeschool-fonds, Stichting Organisatie van Effectenhandelaren te Rotterdam.

Splitting Orders in Fragmented Markets

Evidence from Cross-Listed Stocks

Abstract

A number of recent theoretical studies have explored trading in fragmented markets, e.g. Biais et al. (2000), a phenomenon increasingly witnessed in modern markets. The key assumption generating the results is that there is at least one liquidity demander exploiting access to all markets by optimally splitting orders across markets. This paper seeks to test this assumption in a natural experiment involving Dutch stocks that are traded both in Amsterdam and New York. The results confirm the presence of rational, order splitting traders. This explains the increased volume and relatively large and persistent price changes for the overlapping period.

JEL Codes: G1, G15, G14, G12

Most classic paradigms in market microstructure start from a single centralized market. This setting is justified by a common belief that fragmented markets have a general tendency to consolidate. In practice, however, the trend appears to be in the other direction. The increase in fragmented trading is illustrated by the quadruple growth over the last decade in the number of non-US companies cross-listed on the NYSE, 394 at the end of 1999. This trend has triggered a number of recent studies that prove the existence of an equilibrium in a multiple market setting under certain conditions, e.g. Chowdry and Nanda (1991), Berhardt and Hughson (1997), Biais et al. (2000). One of these conditions is that liquidity demanders exploit multiple markets by splitting orders and simultaneously feeding them to all markets. The intuition is that the total price impact of a split order is smaller than the impact of the entire order sent to one market only.

This paper seeks to test whether liquidity demanders indeed exploit a multiple market setting as predicted by theory. An ideal natural experiment involves two markets that satisfy (i) synchronicity, (ii) liquid trading in the same security, (iii) simultaneous accessibility by at least one trader, and (iv) an equal level of transparency. The last condition is imposed to prevent the trader from routing orders to the least transparent market. The optimality and use of such strategy has been documented by many recent studies (see, e.g., Bloomfield and O'Hara (2000)). Although this fourth condition narrows the domain considerably, it at the same time keeps the focus on those fragmented markets that have the ability to survive. Probably the best-known example of pressure on regulators to create equal transparency is the competition for order flow between the London Stock Exchange and the Paris Bourse in the early 1990s. London was the less transparent market and order flow started to gravitate that way. The Paris Bourse was, as a result, forced to change their reporting rules (Gemmill (1996)).

The natural experiment studied in this paper is the trading of four stocks in Amsterdam and New York: KLM Royal Dutch Airlines, Philips Electronics bv, Royal Dutch and Unilever. These shares are amongst the first non-US stocks introduced on the NYSE. They are New York Registered Shares as opposed to American Depositary Receipts (ADRs) and can therefore be regarded as more comparable to common US shares traded on the NYSE. More importantly, both types of shares can be exchanged for the underlying share and vice versa at a small fee of approximately 15 basis points at the Depository Bank. Simultaneous trading of these securities in New York and the underlying shares in the domestic market can thus be regarded as fragmented trading. The four shares studied in this paper are simultaneously traded one hour each day. On both sides of the Atlantic trading is highly liquid, in particular during the overlapping hour. Volume in New York is at least 30% of Amsterdam volume. High liquidity in New York is further evidenced by the inclusion of two of the shares in the S&P500. Investors can trade in both markets simultaneously since they are open trading platforms with virtually complete access for foreign investors. Finally, the level of transparency in both markets is high since the best quotes and trades are disseminated in real time. Hence, all conditions for the experiment are satisfied and we should therefore expect to see traders split their orders across markets. A nice additional feature of this experiment is that both markets have a period in which the other market is closed. Hence, “privileged” traders that have access to both markets and need to trade an order on a particular day are likely to prefer the overlap. The non-overlap then serves as a benchmark period, enabling us to study the markets’ reaction to the arrival of these privileged traders.

To verify the presence of order splitting traders in the current experiment, we generate predictions based on theory. These are then tested using one year of intraday data on quotes and trades in both markets as well as intraday data on the Dutch Guilder / US Dollar exchange rate. The main

contribution of this paper is that it finds strong empirical evidence of traders who prefer to trade during the overlap and split their orders across markets. To the knowledge of the author this is the first paper documenting such behavior and it therefore creates an empirical basis for the results developed in theoretical papers on the subject.

It is worth noting that some of the findings, e.g. increased volume during the overlap, are consistent with a less involved, “classic” alternative hypothesis based on arbitrage. Under this hypothesis markets consist of liquidity demanders trading in one market only and arbitrageurs trading in both markets. It is shown that all else being equal, the presence of these arbitrageurs deepens the market on both sides since the impact of a trade is partially corrected if followed by an arbitrage trade. Any rational trader having to trade a relatively large order on a specific day might prefer to trade during the overlap, since he can feed the order to the market in a number of smaller orders and its price impact is partially neutralized by subsequent arbitrage trades. The “splitting orders” and “arbitrage” hypotheses can be discriminated by studying signed volume correlations across markets in sufficiently small intervals. Under the null hypothesis of “splitting orders” this correlation is positive whereas it is negative under the alternative hypothesis of “arbitrage.” Such a test is further developed in the paper and the results are strongly in favor of the “splitting orders” hypothesis.

The first section reviews related literature and further positions the paper. The setting is discussed in the second section and summary statistics are studied for the overlap and the non-overlap periods. The third section develops hypotheses based on theory. Appropriate, primarily non-parametric tests are developed and performed in the fourth section. The fifth section tries to identify the traders who split their orders across markets. The sixth and final section reviews the main findings and discusses the implications for the imminence of one global equity market.

I Review of Related Literature

Glosten (1994), Bernhardt and Hughson (1997) and Biais et al. (2000) model multiple markets with a focus on liquidity suppliers also referred to as market makers. All models involve multiple market makers that first post their price schedules and one liquidity demander (referred to as trader) who then optimally chooses how much to buy from or sell to each market maker. Glosten (1994) shows that an electronic open limit order book might be inevitable, because if no liquidity is supplied by the book any competing exchange would expect to lose money by staying open for trade. Bernhardt and Hughson (1997) analyze the duopoly case and show there is no equilibrium where market makers earn zero profits in a classic Kyle (1995) setting. They do not continue and prove the non-existence of a more general equilibrium in this setting. Instead, they recover equilibrium by modifying the setting such that liquidity demand is made endogenous and contingent on the price schedules set by market makers. They show that market makers now extract positive rents and this therefore raises trading costs. Biais et al. (2000) return to the Kyle setting and show that an equilibrium does exist in the oligopoly case. An analysis of this equilibrium shows that trading volume is lower than ex ante efficiency would require, liquidity suppliers make positive expected profits and do not bid the asset price to the expected value of the asset. Increasing the number of suppliers reduces these effects. All three studies should raise the interest of investors and regulators, since they have profound implications for market structures, trading costs and the allocative efficiency of markets.

Chowdry and Nanda (1991) build a richer model that not only involves multiple liquidity suppliers but also multiple liquidity demanders. The latter are categorized as informed traders who privately observe the innovation in the value of the security, “large” liquidity traders with access to all markets and “small” liquidity traders who only trade on one market. They show that it is optimal for

those who can to split their orders across markets. Their results support the intuitive notion that these traders are indeed better off in a fragmented market setting as compared to a single unified market. This is at the cost of the small liquidity traders who in a fragmented market setting are used to camouflage trades of those that split orders across markets. This concern has led to the *Security Acts Amendments of 1975*, which mandated the SEC to move rapidly to a truly nationwide competitive securities market. The increasingly global world in which the same securities are traded worldwide would require a similar initiative on a global scale.

The assumed fairness of a single securities market strived for by the SEC is the subject of academic debate. Blume (2000) argues that this assumption does not capture the realities of modern markets. Investors have different needs and thus rationally prefer different types of markets. Fragmentation therefore is a natural result of competition. Some investors might place a high premium on speed of execution, others might prefer to get the best net price, retail investors are likely to examine differences in commission structures, settlement procedures can be a reason to prefer one exchange to the other, etc. This world is consistent with Nanda and Chowdhry (1991) since some investors might choose to trade in one market for reasons other than the best net price while others might split their trades across markets to benefit from fragmented trading.

The discussion of recent literature shows that fragmented trading is viable theoretically and is likely to occur in modern markets. The “key assumption” in the words of Lawrence Glosten is that some traders split their orders across markets. It is this assumption that is the subject of this paper. Do we find such behavior in current markets?

The presence of traders who split orders across markets has implications for another strand of literature that seeks to model arbitrage. Amongst the most popular models are threshold autoregressive

(TAR) error correction models (see, e.g., Yadav et al (1994), Dwyer et al. (1996)). The assumption is that prices diverge until they hit a bound that triggers arbitrage. In the presence of large traders who split orders across markets, these models fail to capture the complete dynamics of the system. If in addition to splitting orders these traders rationally route more to the cheaper market, this mechanism exhibits error correction features. Markets then are more efficient than arbitrage would suggest. Inferring arbitrage bounds from TAR model estimates might lead you to underestimate the true arbitrage bounds.

II Trading in Amsterdam and New York

The volume of non-US shares grew to over 15% of total NYSE volume in 1999. European shares accounted for most of this volume- approximately one third. Not surprisingly UK shares accounted for most European volume followed by Dutch shares that generated more volume than French and German shares combined. The cross-listed Dutch shares studied in this paper are NY Registered Shares as opposed to ADRs but these are not regarded as materially different in the eyes of investors according to Citibank, one of the key players in the Depository Services industry. Most important is that both the NY Registered Share and the ADR can be changed for the underlying common share at a small fee of approximately 15 basis points.

The Amsterdam Stock Exchange and the New York Stock Exchange are both continuous, consolidated auction markets in the terminology proposed by Madhavan (2000). Both exchanges release quote and trade information in real time. The main difference, however, is that New York is a hybrid market in that orders can arrive at the floor both through brokers and the electronic Superdot system. Amsterdam is a pure electronic market in which orders are routed to a central market maker (“hoekman”) who manages a consolidated limit order book and makes sure that orders are executed

according to time-price priority. Although the market maker has an obligation to “make a market” at times of illiquidity this is not an issue for the blue chip stocks studied in this paper. This setting seems to provide fertile ground for “splitting orders” strategies, in particular because of the ability to route orders to both markets through electronic channels. As a matter of fact, the author has seen traders access both markets from a split screen at the trading floor of a major bank on Wall Street.

The data set used in this study consists of trade and quote data from the Amsterdam Stock Exchange and the NYSE for July 1, 1997 through June 30, 1998. Although ten Dutch stocks were cross-listed in New York at the time, four of them were selected for this study: KLM Royal Dutch Airlines, Philips Electronics bv, Royal Dutch and Unilever. These stocks showed highly liquid trading on both sides of the Atlantic. An intraday data set on the Dutch Guilder / US Dollar exchange rate for the same period enabled us to bring price data to one currency and make meaningful comparisons across markets.

CET	9:30	15:30	16:30	22:00
EST	3:30	9:30	10:30	16:00
	Amsterdam opens	New York opens	Amsterdam closes	New York closes

The timetable for Amsterdam and New York trading during the sample period shows that there is a one hour trading overlap each day. In 1997 daylight savings time ended at the same day for The Netherlands and the US. In 1998, however, The Netherlands changed to daylight savings time one week before the US. This week was removed from the sample.

To illustrate price discovery, figure 1 shows bid and ask prices for Royal Dutch on October 27, 1997- a random day in the data set. This figure leads to a few interesting observations. First, quoted spreads in both markets seem to be competitive. Second, the spread in Amsterdam widens when the New York market opens. Third, the overlap appears to be more volatile than the non-overlap. And, perhaps most importantly, volatility during the overlap is high as compared to the quoted spread. Two independent stochastic processes with this level of volatility and relatively small spreads should yield many “arbitrage” opportunities. The realization shown in figure 1 is highly unlikely under the assumption of independent price discovery.

Figure 2 shows the intraday volume pattern for Royal Dutch based on the entire sample period. Fifteen minute averages and their confidence intervals show that Amsterdam volume jumps by a significant 90% when New York opens. It then stays at these levels during the subsequent two fifteen minute intervals. Still higher volume in the last fifteen minutes of the day reflects heavy last minute trading. Not surprisingly New York shows highest volume on the opening. This is the result of a large opening trade reflecting order build-up during the pre-open period. The volume then drops and stays at the same level during the rest of the overlap. As soon as Amsterdam closes the New York volume drops by a significant 18% and appears to stay at this level in the subsequent fifteen minutes. Although increases in the afternoon and decreases in the morning are to be expected given the well-known stylized fact of intraday U-shape patterns (Goodhart and O’Hara (1997)) this figure shows that the 90% jump and the 18% drop are stronger than a U-shape would suggest. Royal Dutch therefore shows increased trading volume in both markets during the overlapping period.

The notions developed based on price and volume figures for Royal Dutch can be generalized to all four stocks. Table 1 shows averages for volume, number of trades, volatility and spread for every

half hour from 9:00 to 11:00 EST thus including a period for both exchanges during which the other exchange is closed. Both markets share a consistent pattern across all stocks of increased volume, an increased number of trades, increased volatility and virtually unchanged effective spreads during the overlapping period. An aggregate pattern based on seven Dutch stocks cross-listed on the NYSE is estimated in Hupperets and Menkveld (2000). The only difference between the table 1 patterns and the aggregate pattern is that the latter shows a significant 5% jump in spreads in Amsterdam when New York opens and another significant 5% jump in New York when Amsterdam closes. A thorough discussion of the aggregate pattern is beyond the scope of this paper but can be found in Hupperets and Menkveld (2000). A brief discussion of the values and patterns of each of the four trading variables, however, is useful and is presented in the following few paragraphs.

The average, five-minute volume in the overlapping period is highest for Royal Dutch-amounting to more than 70,000 shares traded in Amsterdam and half that number in New York. With an average price of \$55 this means that more than \$385,000 worth of shares changes hands every five minutes in Amsterdam and half that amount in New York. These numbers make Royal Dutch the most liquid share both in terms of absolute volumes and in terms of relative volume in New York. The numbers for Philips and Unilever are approximately 40,000 shares for Amsterdam and one-fourth that amount for New York. KLM is the smallest in terms of trading volume but still trades a five-minute average of 10,000 shares in Amsterdam and one-third that amount in New York. For all stocks, volume jumps significantly in Amsterdam when New York opens, ranging from a 51% jump for Philips to an 80% jump for Royal Dutch. This number for Royal Dutch deviates slightly from the number reported in figure 2, because of a different interval length.

The statistics on the number of trades reconfirm the results for volume with a few minor but nonetheless interesting differences. First, note that during the overlap the average number of trades for five-minute intervals in Amsterdam ranges from 4.2 for KLM to 14.1 for Royal Dutch. Comparable numbers for New York are 2.4 and 15.2. Second, note that when comparing New York to Amsterdam on this measure as compared to the volume measure, New York consistently shows a better relative performance. Apparently the average trade size in New York is smaller. Third, the intraday jumps and drops, although still significant, are smaller than those reported for volume. The average trade size must therefore be higher for both markets during the overlap. This turns out to be an important finding and will be addressed in further detail at a later stage in the paper.

Increased volatility for the overlapping period potentially indicates that the extra volume arriving at the market is informative. Volatility as measured by average squared returns jumps by a significant 71% in Amsterdam upon the New York open and drops by a significant 30% in New York upon the Amsterdam close. The jump in Amsterdam could in theory be due to the information revealed through the New York open. This, however, cannot explain why it is that volatility stays at these higher levels throughout the entire overlapping period (see Hupperets and Menkveld (2000)).

The spreads in both markets are competitive and do not show significant change when comparing the overlap to the non-overlap. The spread measure is the effective spread defined as twice the difference between the prevailing midquote and the transaction price. It is scaled by the midquote to obtain the relative spread. This ex-post measure of spread is preferred to the quoted spread measure, primarily because the latter is a flawed proxy for cost of trade, in particular for the NYSE where approximately one third of the orders are executed inside the quoted spread (Lee and Ready (1991)).

The low spreads during the overlap, ranging from 10.8 to 30.6 basis points, are further evidence of market liquidity.

A rigorous microscopic analysis of trading and price discovery during the overlap requires high data density. Too many intervals without observations can hamper statistical inference. Although table 1 shows that we have an overall average of seven to eight trades for five-minute intervals the data could still show many periods with no trades due to the tendency of trades to cluster in time (Engle and Russell (1998)). Table 2 shows the fraction of intervals containing at least one observation for both quotes and trades for different interval lengths. For intervals as short as one minute an average 74% of all intervals have at least one trade for Amsterdam and 60% for New York. These numbers do not show extreme variation across the four stocks. KLM is at the low end with 49% and 34% and Royal Dutch is at the high end with 89% and 86% respectively. When looked at five-minute intervals these numbers improve dramatically with an average of 98% containing at least one trade for Amsterdam and 91% containing at least one trade for New York. The quote data are even better in terms of density across all stocks and interval lengths. For one-minute intervals an average of 81% contain at least one quote in Amsterdam and 67% in New York. Comparable five-minute numbers are 99% and 91%. These results show that this data set indeed enables us to study intervals as short as one minute in a meaningful way.

Now that we have a basic understanding of trading during and outside the overlap and have set the stage for analysis based on short time intervals, we can start studying the notion of high market efficiency as was suggested by figure 1. This figure showed that midquotes on both sides of the Atlantic appeared to move in lockstep for Royal Dutch on October 27, 1997. From table 1 it is apparent that the standard deviation of five-minute returns is roughly equal to the size of the relative effective

spread. This observation leads one to believe *ex ante* that midquote returns in both markets cannot be uncorrelated, since that would inevitably lead to arbitrage opportunities given that there are 12 five-minute returns in the overlapping hour. Table 3 documents correlation in midquote returns for different interval lengths. Correlations are significantly positive for all stocks and all interval lengths. The one-minute interval correlations range from 0.10 for KLM to 0.25 for Royal Dutch. These are conservative estimates for true correlations since intervals without quote updates are registered as zero return intervals. The underlying efficient price might have changed but is not revealed since no new quotes were recorded. The correlations for five-minute intervals range from 0.39 for KLM to 0.72 for Royal Dutch. Both markets are therefore shown to move in lockstep.

Higher volatility during the overlap suggests that a disproportionate amount of information is revealed during this period. But, can high volatility not be the result of disproportionate noise? Alternatively, can price changes be large but non-informative since they are corrected for at a later stage? To study this question we show the autocorrelation function for one-minute returns with up to five lags in table 4. All significant autocorrelations are positive and we therefore do not see error correction in price changes on either market. This, however, does not preclude corrections over longer periods of time. To study this we decided to calculate a ratio with the n-minute return variance in the numerator and n times the one-minute return variance in the denominator, with n a positive integer. To show why this will tell us whether price changes are transient or persistent we develop the numerator as

$$\text{var}(r) = \text{var}(r(1) + r(2) + \dots + r(n)) = \sum_i \text{var}(r(i)) + \sum_{\substack{i,j \\ i \neq j}} \text{cov}(r(i), r(j)) \equiv A + B \quad (1)$$

where r is an n -minute return,

$r(i)$ is the return in the i 'th one-minute subinterval and

$\text{var}(\cdot)$, $\text{cov}(\cdot, \cdot)$ are the variance and covariance operators.

The denominator is represented as the sum of two terms A and B where A is equal to the denominator since it equals n times the average variance and B is the sum of all off-diagonal elements in the variance matrix for r . A ratio higher than 1 corresponds with positive B and a ratio lower than 1 with negative B . The interpretation is that in the first case positive covariances dominate negative covariances or, alternatively, “persistence” dominates “correction.” The alternative case of a ratio lower than 1 corresponds to net “correction” effects which would be expected in case of noise. The variance ratios are calculated for 5-, 15-, 45- and 60-minute intervals and shown in table 4.

Significance of these ratios is judged by comparing them with critical values of the test statistic under the naïve assumption of independent increments. Although the distribution of this statistic resembles an F-distribution it is not the same. It has thinner tails because the variance estimates in the numerator and the denominator are based on the same data and therefore not independent. The critical values are found through simulations. The results show that the ratios are all significantly larger than 1 for Amsterdam. Most of them are significantly larger than 1 for New York with the exception of KLM for which we find two ratios that are significantly lower than 1. The overall 60-minute return ratios show that variance over the entire overlapping period is at least 27% higher and up to 99% higher than predicted by the variance of one-minute returns under the naïve assumption of independent increments. What all these results show is that price changes are not only larger during the overlap as evidenced by higher volatility, they appear to be persistent as well.

This section has sought to show that the markets in Amsterdam and New York are competitive when it comes to trading in the same security. Both markets attract considerable volume and effective spreads are comparable. This sets the stage for the privileged traders of Chowdhry and Nanda (1991), the “large” liquidity traders and informed traders who have access to both markets and exploit the situation by splitting orders and feeding them to both markets simultaneously. The empirical findings thus far are entirely consistent with such behavior. Volumes are significantly higher during the overlap. Volatility is significantly higher and price changes are persistent, which indicates that information is revealed through the activity of informed traders. Effective spreads are not significantly different when comparing the overlap to the non-overlap, thus encouraging the “privileged” traders to trade during the overlap using the “splitting orders” strategy. These findings, however, are also consistent with an alternative explanation and the next section seeks to develop direct tests for the “splitting orders” hypothesis and discriminates it from a more straightforward alternative explanation.

III Splitting Orders or Arbitrage?

The rationale for splitting orders and feeding them to multiple markets simultaneously is that in this way trading cost is minimized because the price concession is limited. This intuitive notion is the key assumption on which virtually all theoretical models developed for a one-security-multiple-markets framework depend. This notion can be formalized using a simple “reduced form” market microstructure model with no commissions, zero spreads and market depth equal to one, or, in formula terms:

$$\begin{aligned} dP &= m * dQ, \\ m &= 1 \end{aligned} \tag{2}$$

where dQ is signed volume, positive for a buy order, negative for a sell order,

dP is the price impact of the order and

$(1/m)$ is the market depth.

In this setting no order can be traded at the prevailing price since each infinitesimally small order changes the price linearly with factor m . The cost of trading in this framework is solely the result of the price concession and, for a buy order, this is equal to:

total price concession “benchmark case” =

$$\int_{x=0}^Q (p(x) - p(0)) dx = \int_{x=0}^Q (m * x) dx = 0.5 * m * Q^2 = 0.5 * Q^2. \tag{3}$$

Due to symmetry the result for a sell order is equal to that for a buy order. In the case of two markets with market depth equal to 1, this order of size Q can be split in two and fed to both markets such that the total price concession is smaller, since

total price concession “splitting orders” =

$$2 * \int_{x=0}^{0.5*Q} (p(x) - p(0)) dx = 2 * \int_{x=0}^{0.5*Q} (m * x) dx = 2 * 0.5 * m * (0.5 * Q)^2 = 0.25 * Q^2 \tag{4}$$

Both these situations are illustrated in figure 3 in the first and second graph. The third graph in this figure illustrates an alternative and perhaps more straightforward explanation, which is entirely based on arbitrage. In this explanation liquidity demanders trade in one market only. They might have had discretion over which market to trade, but we study the situation in which they have chosen and trade on one market only. In addition, there are professional arbitrageurs that have instant access to both markets and immediately trade on any arbitrage opportunity that arises. In this way prices on both markets will move in lockstep and, more importantly, it is rational for those who trade sizeable orders to choose to trade during the overlap. The rationale is that these traders know of the presence of arbitrageurs and cut their order into smaller ones that they will then feed to the market one at a time. In the same setting used to illustrate optimality of a “splitting orders” strategy, the order of size Q can be split in two equal sized orders and fed to the market subsequently. The execution of the first order triggers an arbitrage trade of size $\frac{1}{4}Q$ in the opposite direction, hence half the price impact of the first order is neutralized and the trader can start executing the second order starting at a lower price than the price that would have prevailed had he sent the complete order at once. The total price concession in this case amounts to

total price concession “arbitrage” =

$$\int_{x=0}^{0.5*Q} (m * x) dx + \int_{x=0}^{0.5*Q} (0.25 * Q + m * x) dx = 2 * 0.5 * (0.5 * Q)^2 + 0.25 * Q * (0.5 * Q) = 0.375 * Q^2 \quad (5)$$

Although the total price concession appears to be higher than in the “splitting orders” setting, it can be shown that by splitting up the order into infinitely many small orders and subsequently feeding them to the market the total price concession is the same as in the “splitting orders” case. In this world of

arbitrageurs, as was the case in the world of traders splitting orders, the results of increased volume and volatility during the overlap can be explained by optimal behavior of market participants. In the course of this paper these two explanations will be referred to as the null hypothesis of a world in which traders are splitting orders versus the alternative hypothesis of a world where arbitrageurs are keeping markets efficient. Note that although both mechanisms can co-exist, it is interesting to see which one is most likely to have caused the empirical results reported thus far.

Both the “arbitrage” and the “splitting orders” world generate three testable predictions, two of which are identical. The first prediction is that although market depth could be worse during the overlap when compared to the non-overlap, this deterioration is not unbounded. In both worlds those who prefer to trade during the overlap only do so in order to minimize trading costs. If, for example, market depth outside the overlap is m and inside the overlap less than $\frac{1}{2} m$, it is optimal in both worlds for those who have discretion over timing to trade outside the overlap. This prediction is formalized as:

H1: Market depth during the overlap as compared to the non-overlap cannot be such that those who have discretion over timing prefer to trade during the non-overlap.

The second prediction generated by both hypotheses is that the data should show coordinated, simultaneous trading in both markets. In the world of “arbitrage” this reveals the arbitrage trades, a sell in one market and a buy in the other. In the “splitting orders” world this reflects traders splitting their orders and feeding them to both markets simultaneously.

H2: There is coordinated, simultaneous trading across markets.

This coordinated trading suggests a stronger prediction that can discriminate between the “arbitrage” and the “splitting orders” hypothesis.

H3(0): Correlation of signed volume is positive (“splitting orders” hypothesis)

H3(a): Correlation of signed volume is negative (“arbitrage” hypothesis)

Developing an appropriate test for all of the predictions is not as straightforward as it seems. The next section discusses the issues, proposes appropriate tests and generates results.

IV Empirical Evidence

We return to trading in Royal Dutch on October 27, 1997 to develop intuition for what will later be shown using primarily non-parametric econometric tests. Figure 4 plots the bid ask prices for the overlapping hour on the right axis and adds one minute signed volumes on the left axis. The graph seems to support a world of “splitting orders,” since volume for both markets seems to cluster in the same periods and has the same sign. The periods A, B and C correspond to periods of buy, sell and sell orders respectively in both(!) markets. The price changes seem indeed to be driven by signed volume supporting the “reduced form” market microstructure model in which volume drives price. Although this graph appears to support a world of “splitting orders,” one should be careful inferring from this graph that little or no arbitrage opportunities have occurred. The reason is that bid and ask prices in this graph are one-minute “snapshots.” If arbitrage trades occur immediately when an opportunity arises, this wipes out the opportunity, which means that it is hard to observe those opportunities using a time

grid- however fine it may be. That this particular day did not seem to have had arbitrage opportunities is evident from the absence of intervals with simultaneous volume in both markets with opposite sign.

H1: Market Depth during the Overlap

The first prediction that any deterioration of market depths during the overlap should be limited is verified by estimating a linear “reduced form” market microstructure model that can be formalized as:

$$\Delta \log(\text{Midquote}(t)) = \alpha_0 + \alpha_1 * \text{Signed_Volume}(t) + \varepsilon(t) \quad (6)$$

where $\varepsilon(t)$ is an i.i.d. distributed error term with zero mean and

$(1/\alpha_1)$ is the market depth.

To allow for different market depths for different times of day, the day is split in N periods and the model adjusted such that:

$$\Delta \log(\text{Midquote}(t)) = \alpha_0 + \sum_{j=1}^N (I^j(t) * \alpha_j * \text{Signed_Volume}(t)) + \varepsilon(t) \quad (7)$$

where $I^j(t)$ takes the value one if t falls in the j’th period in the day

$(1/\alpha_j)$ is the average market depth at the j’th period in the day.

This model is estimated for Amsterdam and New York based on one-minute intervals and N equal to 3. The periods considered for Amsterdam are 9:00-9:30 before the NYSE open, 9:30-10:00 and 10:00-10:30 during the overlap. The periods for New York are 9:30-10:00, 10:00-10:30 and 10:30-11:00 with the last period being the period just after the Amsterdam close. The model estimates are reported in

table 5. The model fits remarkably well for a “simple” linear regression model on almost 20,000 one-minute returns for a variety of days. The R^2 is between 0.25 and 0.31 for Amsterdam and between 0.06 and 0.11 for New York. It is therefore not surprising that all coefficients are highly significant.

Comparing market depths across markets shows that both markets appear to be competitive when it comes to market depth with KLM being an exception since for that stock the market in Amsterdam is almost twice as deep. The prediction in both the “arbitrage” and “splitting orders” world is that market depth cannot be worse during the overlap to the extent that traders would be better off trading outside the overlap. These results show that market depth in Amsterdam worsens for all stocks immediately after the New York open and then improves again in the last half-hour of trading. The factor by which the market worsens is less than 15% for all stocks. In New York, on the other hand, results are mixed, ranging from a 10% improvement in market depth for KLM and a 30% reduction in market depth for Philips after the Amsterdam close. These changes, however, do not render the strategy of trading during the overlap sub-optimal and hence confirm the first prediction.

H2: Coordinated, Simultaneous Trading?

The second prediction says that the data should show coordinated, simultaneous trading in both markets. The intuitive test is to study contemporaneous correlation in Amsterdam and New York volume during the overlap. This correlation is positive if there are traders, be it “large” liquidity traders, informed traders or arbitrageurs who simultaneously trade in both markets. The first column in table 6 shows this correlation for both one- and five-minute intervals. It has the right sign and is highly significant for all stocks ranging from 0.09 to 0.13 for one-minute intervals and from 0.20 to 0.38 for five-minute intervals. Although this undeniably shows that periods of high volume in Amsterdam coincide with periods of high volume in New York, this could be for reasons exogenous to the

“arbitrage” or “splitting orders” world. It is, for example, very likely that the days(!) of high volume in Amsterdam coincide with days of high volume in New York. Market sentiment in New York has been shown to determine market sentiment in markets worldwide. This effect causes positive correlation in volume across markets and therefore constitutes an alternative explanation for the positive correlation observed in the data. To correct for this effect volume in both markets is scaled by daily volume. The contemporaneous correlations for scaled volume are in the second column of table 6 and show that correlation has indeed dropped for all shares and interval lengths. The resulting correlations are still significantly positive ranging from 0.04 to 0.10 for one-minute intervals and 0.11 to 0.16 for five-minute intervals. These correlations, however, are likely to underestimate the effect of coordinated trading in both markets, because of the U-shape in volume. The market in Amsterdam is at the end of the day during the overlapping period and therefore shows increases in average volume as time progresses, whereas the New York market is at the start of the day and shows decline in average volume. This trend difference negatively affects volume correlations and the current estimates, therefore, are likely to underestimate coordinated trading. To correct for this effect volume, after it is scaled by daily volume, is demeaned by subtracting the mean for the time of day. The third column in table 6 shows that the correlations for this adjusted volume is indeed higher for all stocks across all interval lengths. The resulting correlations are significantly positive ranging from 0.04 to 0.11 for one-minute intervals and 0.13 to 0.24 for five-minute intervals. This confirms the third prediction since periods of high volume in Amsterdam coincide with periods of high volume in New York.

H3: Splitting Orders or Arbitrage?

Now that we have shown that there appears to be coordinated trading, a test based on signed volume should reveal whether this is the result of traders splitting orders across markets or arbitrageurs

exploiting arbitrage opportunities. An important consideration for designing a test is that splitting orders across markets is most likely when price differences are small, whereas arbitrage only occurs when price difference are sufficiently large. This suggests that we condition on the price difference at the start of the interval when analyzing correlation in signed volume. To be more precise, an arbitrage opportunity only exists when the bid price in one market exceeds the ask price in the other market by an amount larger than the cost of arbitrage with the conversion fee (+/- 15 basispoints) being the lower bound. Based on this observation we create the variable “Arb_Opp” in the following way:

Arb_Opp (t) = the bid price in Amsterdam -/- the ask price in New York, if this is positive,
= the ask price in Amsterdam -/- the bid price in New York, if this is negative,
= 0 otherwise,

where bid and ask prices are the prevailing prices at time t.

For each interval $[t(i), t(i+1)]$ we take the value of “Arb_Opp” at the start of the interval and condition on this variable when calculating signed volume correlations. For “Arb_Opp” equal to zero there is no *a priori* reason to expect arbitrage and we therefore expect positive sign in signed volume correlation in the “splitting orders” world and zero correlation in the “arbitrage” world. When looked at the intervals where “Arb_Opp” is nonzero and larger in absolute value than some value x interpreted to be the cost of arbitrage, these predictions change. In a world of “splitting orders” the correlation is still positive, whereas in a world of “arbitrage” these intervals are expected to show negative correlation. As mentioned before an arbitrage strategy is only successful when trades are executed immediately after the arbitrage opportunity arises. If not, the arbitrageur runs the risk of trading when the opportunity has

disappeared due to other trades arriving at the market. It is for this reason that we have to consider small intervals when studying arbitrage. We decided to look not only at five- and one-minute intervals but also at intervals of 20 seconds for which only those are considered with nonzero volume in both markets.

The correlations of signed volume conditioned on the presence of potential arbitrage opportunities at the start of the interval are reported in table 7 and show overwhelming evidence of a world of traders splitting orders across markets as opposed to a world of arbitrageurs. For five-minute intervals that do not have arbitrage opportunities at the start, “Arb_Opp” being zero, the correlation in signed volume is significantly positive for all stocks ranging from 0.11 for KLM to 0.34 for Royal Dutch. When decreasing the length of the interval to one minute these correlations remain significantly positive for all stocks, ranging from 0.06 for Philips to 0.09 for KLM. Even for intervals as short as twenty seconds the correlation is significantly positive for Unilever and Royal Dutch. This strongly indicates the presence of traders that split orders across markets. We now start to look at intervals with potential arbitrage opportunities. For five-minute intervals that have a varying degree of arbitrage opportunities at the start of the interval correlations remain positive for Philips and KLM, most of them significant. For Royal Dutch and Unilever they do turn negative for opportunities larger than 30 basis points, although insignificant. Changing the interval length to one minute KLM and Philips continue to show primarily positive correlations and the negative correlation for Unilever disappears. Royal Dutch continues to show negative correlation, still insignificant though, for arbitrage opportunities larger than 30 basis points. We lose almost all significance when turning to intervals of 20 seconds, but now see arbitrage for Royal Dutch as evidenced by the significant -0.12 correlation in signed volume when the arbitrage opportunity is larger than 30 basis points at the start of the interval.

The correlation analysis shows positive significant positive correlation in signed volume for all stocks indicating the presence of order splitting traders. Only for Royal Dutch do we find significant negative correlation in signed volume for opportunities larger than 30 basis points. Add to this that these opportunities only appeared in 3% of all intervals, whereas the evidence generated in favor of the “splitting orders” hypothesis was based on 85% of the intervals for KLM, 74% for Philips, 97% for Royal Dutch and 96% for Unilever. These are those one-minute intervals that showed significant positive correlation in order flow. This is therefore compelling evidence of not only the presence of order splitting traders but also of them driving the increased volume and volatility during the overlap.

A believer in the “arbitrage” world might not give up and could argue that the results of positive correlation for intervals with zero arbitrage opportunities at the start can, in fact, be due to arbitrage under certain conditions. To develop intuition, we return to the “reduced form” market microstructure model and consider the case that both markets have equal depth and two unrelated orders of size Q are arriving simultaneously, one to market A and the other to market B. These orders can be either sell or buy orders. We thus have to consider four equally likely cases: a buy in market A and a sell in market B, a sell in A and a buy in B, a sell in both markets and a buy in both markets. In the latter two cases the arbitrageur does not take action since the markets have moved in the same direction, but in the first two cases the arbitrageur does trade size Q orders of opposite sign immediately in both markets. The net order flow observed is therefore $(0,0)$, $(0,0)$, $(-Q,-Q)$ and $(+Q,+Q)$ respectively. This shows that under these conditions signed volume can indeed be positively correlated due to arbitrage. The test results reported thus far have tried to deal with this problem by evaluating small interval lengths assuming that, in practice, arbitrageurs cannot react immediately to arbitrage opportunities. A stronger and more convincing test, however, is to screen the entire interval on arbitrage opportunities and only

consider those intervals that did not show any arbitrage opportunity. Thus far, we have conditioned on “Arb_Opp” at the start of the interval, but we can condition on this variable being zero throughout the interval. For these intervals we can exclude the arbitrage explanation and would expect zero correlation if the arbitrage effect was driving the results. The correlations are reported in the last eight rows of table 7 and show that the positive correlation not only remains significantly positive, but appears to be even stronger!

V Who is Splitting Orders across Markets?

The fragmented trading in Amsterdam and New York appears to be exploited by traders that split their orders across markets. But, who are these traders? And, if they are rational shouldn't we expect them to send more volume to the cheaper market, the market with the highest bid in case of a sell order or the market with the lowest ask in case of a buy order? These questions are explored in this section.

To start with the last question, a rational order splitting trader with a large buy order will send relative more volume to the market with the lower ask. Although we cannot identify his trades, his behavior implies that the difference in aggregate buy volume in both markets is negatively correlated to the difference in the ask price. This is a testable implication. Table 8 contains correlation estimates based on one-minute intervals and shows that the correlations are indeed significantly negative for all stocks ranging from -0.03 for Philips to -0.09 for KLM and Royal Dutch. The same can be done for sell volume and bid price. The only difference is that we now expect a positive correlation since it is optimal to sell in the market with the highest bid. The correlations have the right sign and are significant for three out of four stocks. KLM correlation is 0.01 and insignificant. The others range from 0.04 for Philips to 0.09 for Royal Dutch.

What can we say about the identity of the trader splitting orders across markets? Chowdry and Nanda (1991) postulate that it is the “large” liquidity traders and the informed traders that split orders across markets. Taking one step back, we could say that, in practice, this behavior only makes sense for those that are trading orders that would need major price concessions when traded at one market, in other words the very large orders. Even when those orders are split they are likely to still be larger than the average orders in the market. This implies that the distribution of order size shifts in favor of large orders in Amsterdam when New York opens and in favor of small orders in New York when Amsterdam closes. This is tested by evaluating the change in the average number of orders for five different size categories. The results are shown in table 9. All four stocks for both markets show indeed largest changes for orders in the highest size categories when comparing the overlap to the non-overlap. For Royal Dutch, for example, orders larger than 5,000 shares jump by a significant 83% when New York opens, orders between 1,000 and 5,000 shares jump by a significant 67% and orders smaller than 1,000 shares jump by no more than 12%. Comparable numbers for New York are drops of 26%, 22% and less than 6% respectively on the Amsterdam close. This pattern is consistent across all stocks and shows that indeed order flow composition is different for the overlap in favor of large orders. But, can we say in addition whether it is the “large” liquidity traders or the informed traders splitting these large orders across markets? In other words, can we discriminate between the two? One potentially fruitful starting point is to assume that heavy, liquidity motivated trading is likely to be market wide, whereas trading based on private, stock-specific information is not. Although buying or selling one stock in large amounts might very well be liquidity motivated, it could be argued unlikely since investors seeking exposure to the stock market for liquidity reasons should prefer to hold “index” portfolios for reasons of risk diversification. Table 10 shows five-minute contemporaneous correlations

in signed volume across stocks in Amsterdam as well as in New York. In Amsterdam we find evidence of market wide trades reflected in all correlations being significantly positive, ranging from 0.16 for KLM and Philips to 0.46 for Unilever and Royal Dutch. This result holds for the overlap as well as for the non-overlap. The results for New York, on the other hand, only show a significantly positive correlation of 0.07 for Royal Dutch and Unilever during the overlap, which loses significance outside the overlap but still amounts to 0.06. To study whether it is these “large” liquidity traders trading the “index” in Amsterdam who exploit access to both markets we proceed by decomposing order flow for each stock in two components, one in line with and the other orthogonal to market order flow. The latter is interpreted to be order flow emanating from privately informed traders with stock-specific information. This decomposition is carried out by estimating the following equation:

$$Signed_Volume(j,t) = \alpha_0 + \sum_{\substack{i=1 \\ i \neq j}}^N \alpha_i * Signed_Volume(i,t) + \varepsilon(t) \quad (8)$$

where $Signed_Volume(j,t)$ is the signed volume in the j 'th stock at time t and

$\varepsilon(t)$ is an i.i.d. distributed error term with zero mean.

Based on the model estimates the order flow for the j 'th stock can be decomposed as follows

$$Signed_Volume(j,t) = f(Signed_Volume(j,t)) + e(t)$$

$$f(Signed_Volume(j,t)) = \hat{\alpha}_0 + \sum_{\substack{i=1 \\ i \neq j}}^N \hat{\alpha}_i * Signed_Volume(i,t) \quad (9)$$

where $\hat{\alpha}$ denotes the parameter estimates,

$f(Signed_Volume(j,t))$ is the signed volume forecast conditional on the signed volume observed for the rest of the market and

$e(t)$ is the residual.

The first term in the equation is the forecast of signed volume for some stock j conditional on signed volume observed for the other stocks and can thus be interpreted as liquidity motivated order flow in line with the market. The second term is the residual and is by construction the component of order flow orthogonal to the market. This decomposition is performed for signed volume in Amsterdam and contemporaneous correlation across markets is studied for both components and compared to the original signed volume correlation. Table 11 contains the results showing that both order flow components are positively correlated to order flow in New York, but the orthogonal component correlations are consistently stronger. The correlation results for the order flow in line with the market are positive for all stocks but only significant for Unilever and Royal Dutch with values of 0.08 and 0.20 respectively. The results for the orthogonal component are all significantly positive ranging from 0.13 for KLM to 0.23 for Royal Dutch. These results can be considered evidence for both “large” liquidity traders and informed traders splitting orders across markets with stronger indications for the presence of the latter ones.

VI Conclusion

The aim of this paper is to test, using a natural experiment, whether traders act rationally by splitting orders across markets. It therefore verifies the “key” assumption underlying results generated by recent theoretical studies on the subject. The experiment concerns four stocks that are traded in Amsterdam and New York during a one-hour overlapping period. Both markets are highly liquid, easily accessible and competitive. In this setting, theory suggests that traders should split their orders across markets to limit price concession and therefore reduce the cost of trading. This paper develops predictions based

on this conjectured optimal behavior and tests them using intraday data on trades, quotes and the exchange rate for the period from July 1997 through June 1998.

Intraday patterns show that volume during the overlap is significantly higher than is suggested by the well-known U-shape. Volume in Amsterdam is on average 68% higher for the overlap, in New York 29%. That this volume is informative is evidenced by (i) significantly higher volatility during the overlap and (ii) price changes being persistent rather than transient. These results are consistent with the “splitting orders” hypothesis, since “privileged” traders with access to both markets have reasons to prefer to trade during the overlap. It is, however, also consistent with an alternative hypothesis based on a “classic” arbitrage argument that involves liquidity demanders who only trade in one market and arbitrageurs who trade in both markets.

Both the “splitting orders” and “arbitrage” hypotheses are explored in order to generate positive predictions. The first two predictions do not discriminate between the two hypotheses. Both are consistent with evidence generated by comparing market depth in and outside the overlap and exploring whether intervals of high volume in Amsterdam and New York coincide. The most important prediction discriminates between the two hypotheses and is based on contemporaneous correlation in signed volume across markets. To demonstrate the presence of order splitting traders we condition on the absence of arbitrage opportunities during the interval and find significantly positive correlation in signed volume for all stocks. To find evidence of arbitrage we condition on the presence of a potential arbitrage opportunity at the start of the interval. Correlation in signed volume should be negative for those intervals. Such correlation is only found for Royal Dutch concerning price differences exceeding 30 basis points. Add to this the fact that such differences only existed for 3% of all intervals, whereas the evidence generated in favor of the “splitting orders” hypothesis was based on at least 49% of all

intervals for all stocks, and we see compelling evidence that the presence of order splitting traders is the cause of increased volume and volatility during the overlap.

Further analysis tries to identify the traders who are splitting orders across markets. The order flow composition during the overlap being skewed towards larger orders supports the notion that those who pursue an order splitting strategy are the ones who trade larger orders in the first place. This is perfectly intuitive since they are the ones to gain most from splitting their orders across markets. Additionally, the Amsterdam market shows significantly positive intra-market correlation in signed volume indicating heavy trading for liquidity reasons, in other words selling or buying part of an entire portfolio. By decomposing signed volume for each stock into a component in line with market order flow and a component orthogonal to market order flow, we are able to discriminate between order flow for reasons of liquidity and order flow potentially based on private stock-specific information. These two types of order flow can be interpreted as emanating from “large” liquidity traders and informed traders respectively, the two types of order splitting traders identified in theory (Chowdhry and Nanda (1991)). Both components of signed volume are positively correlated across markets indicating the presence of both types of traders. The results are consistently stronger for the component orthogonal to market order flow, which is evidence in favor of the presence of privately informed traders with stock-specific information.

The empirical results strongly indicate that liquidity demanders that typically trade the larger orders indeed exploit their privilege of multiple markets access. Add to this that fragmented markets are theoretically viable and likely to emerge in a world of exchanges that not only compete in terms of best net price but also on a number of other attributes such as speed of execution, commission

structures, settlement procedures, etc, (Blume (2000)) and then ask the question: “Is one, worldwide equity market imminent?”

References

- Bernhardt, D. and Hughson, E. 1997, “Splitting Orders”, *The Review of Financial Studies*, 10(1), 69-101.
- Biais, B. Martimort, D. and J.C. Rochet, 2000, “Competing Mechanisms in a Common Value Environment”, *Econometrica*, 68(4), 799-837.
- Bloomfield, R. and M. O’Hara, 2000, “Can transparent markets survive?”, *Journal of Financial Economics*, 55, 425-459.
- Blume M.E., 2000, “The Structure of US Equity Markets”, Rodney L. White Center for Financial Research Discussion Paper #017-00, The Wharton School.
- Chowdry B., and V. Nanda, 1991, “Multimarket Trading and Market Liquidity”, *Review of Financial Studies* 4, 485-511.
- Dwyer, G.P.J., Locke, P., Yu, W., 1996, “Index Arbitrage and Nonlinear Dynamics between the S&P 500 Futures and Cash”, *Review of Financial Studies*, 9, 301-332.
- Engle, R.F. and J.R. Russell, 1998, "Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Rates", *Econometrica*, 66, 1127-1162.
- Gemmill, G., 1996, “Transparency and Liquidity: a Study of Block Transactions in the London Stock Exchange under Different Publication Rules”, *Journal of Finance*, 1765-1790.
- Glosten, L., 1994, “Is the Electronic Open Limit Order Book Inevitable?”, *Journal of Finance*, 49, 1127-1161.

- Goodhart, C.A.E., and M. O'Hara, 1997, High Frequency Data in Financial Markets: Issues and Applications, *Journal of Empirical Finance* 4, 73-114.
- Hupperets, E.J. and A.J. Menkveld, 2000, "Intraday Analysis of Market Integration: Dutch Blue Chip Stocks traded in Amsterdam and New York", *Journal of Financial Markets*, forthcoming.
- Kyle, 1985, "Continuous Auctions and Insider Trading", *Econometrica*, 53, 1315-1335.
- Lee, C.M.C., and M.J. Ready, 1991, "Inferring Trade Direction from Intraday Data", *Journal of Finance*, 46, 733-746.
- Madhavan, A., 2000, Market microstructure: A survey, *Journal of Financial Markets* 3(3), 205-258.
- Yadav, P.K., Pope, P.F., Paudyal, K., 1994, "Threshold Autoregressive Modeling in Finance: The Price Differences of Equivalent Assets", *Mathematical Finance*, 4, 205-221.

Table 1: Summary Statistics and Intraday Jumps

5 minute intervals		Amsterdam			New York			Δ AMS on	Δ NY on
(σ)	EST	9:00- 9:30	9:30- 10:00	10:00-10:30	9:30- 10:00	10:00-10:30	10:30-11:00	NY Open	AMS Close
Volume (in shares)	KLM	6,089 (419)	10,769 (419)	11,326 (438)	3,064 (202)	3,270 (175)	1,783 (124)	77%*	-45%*
	Philips	24,793 (1,080)	37,431 (1,080)	44,631 (1,131)	11,995 (588)	12,780 (495)	8,260 (350)	51%*	-35%*
	Royal Dutch	39,161 (1,638)	70,483 (1,638)	77,345 (1,717)	38,614 (1,125)	35,292 (1,006)	26,878 (712)	80%*	-24%*
	<u>Unilever</u>	<u>21,534</u> (983)	<u>35,062</u> (0)	<u>39,520</u> (1,029)	<u>13,532</u> (530)	<u>13,407</u> (461)	<u>9,045</u> (326)	<u>63%*</u>	<u>-33%*</u>
	Mean	22,894	38,436	43,205	16,801	16,187	11,492	68%*	-29%*
Number of Trades	KLM	2.6 (0.1)	4.2 (0.1)	4.3 (0.1)	2.5 (0.1)	2.4 (0.1)	2.0 (0.1)	61%*	-20%*
	Philips	7.6 (0.2)	10.3 (0.2)	11.6 (0.2)	7.2 (0.2)	6.2 (0.1)	4.6 (0.1)	36%*	-25%*
	Royal Dutch	9.7 (0.2)	13.3 (0.2)	14.1 (0.2)	15.2 (0.2)	14.0 (0.2)	12.2 (0.1)	37%*	-13%*
	<u>Unilever</u>	<u>6.1</u> (0.1)	<u>8.5</u> (0.1)	<u>9.1</u> (0.1)	<u>6.3</u> (0.1)	<u>6.3</u> (0.1)	<u>4.8</u> (0.1)	<u>41%*</u>	<u>-23%*</u>
	Mean	6.5	9.1	9.8	7.8	7.2	5.9	40%*	-18%*
Volatility (in basis_points ²)	KLM	328 (67)	452 (67)	584 (70)	521 (46)	557 (40)	302 (28)	38%	-46%*
	Philips	259 (25)	393 (25)	506 (26)	454 (28)	378 (23)	265 (16)	52%*	-30%*
	Royal Dutch	158 (23)	403 (23)	517 (24)	399 (27)	428 (24)	366 (17)	155%*	-14%
	<u>Unilever</u>	<u>121</u> (14)	<u>232</u> (14)	<u>274</u> (15)	<u>225</u> (16)	<u>255</u> (14)	<u>195</u> (10)	<u>91%*</u>	<u>-23%*</u>
	Mean	217	370	470	400	404	282	71%*	-30%*
Effective Spread (in basis_points)	KLM	24.9 (0.5)	25.9 (0.5)	25.6 (0.5)	30.6 (0.8)	28.2 (0.7)	27.3 (0.5)	4%	-3%
	Philips	18.5 (0.3)	18.3 (0.3)	17.7 (0.3)	14.0 (0.3)	12.8 (0.3)	12.7 (0.2)	-1%	-1%
	Royal Dutch	15.0 (0.2)	16.2 (0.2)	15.7 (0.2)	13.3 (0.2)	13.1 (0.2)	13.2 (0.1)	8%	0%
	<u>Unilever</u>	<u>14.5</u> (0.3)	<u>15.0</u> (0.3)	<u>16.1</u> (0.3)	<u>11.7</u> (0.3)	<u>10.8</u> (0.3)	<u>12.0</u> (0.2)	<u>3%*</u>	<u>11%*</u>
	Mean	18.2	18.8	18.8	17.4	16.2	16.3	3%	0%

*: Significant at a 99% confidence level

This table shows averages based on five minute intervals for half hour time periods from 9:00-11:00 EST for those days that both exchanges are open. Given that the overlapping period is 9:30-10:30 these figures contain a 'benchmark' period of trading during which the other exchange is closed. The intraday jump in Amsterdam on the New York open is based on comparison of the period 9:00-9:30 and 9:30-10:00, the intraday drop for New York on 10:00-10:30 and 10:30-11:00. Significance is tested based on the difference in means for these periods. Standard deviations are in brackets and calculated after correcting for differences in daily volume.

Table 2: Data Density

	Fraction of Intervals Containing Trades				Fraction of Intervals Containing New Quotes		
	1 min	5 min	15 min		1 min	5 min	15 min
Amsterdam							
KLM	49	92	100		61	96	100
Philips	81	100	100		87	100	100
Royal Dutch	89	100	100	0	92	100	100
<u>Unilever</u>	<u>77</u>	<u>100</u>	<u>100</u>		<u>84</u>	<u>100</u>	<u>100</u>
Mean	74	98	100		81	99	100
New York							
KLM	34	77	97		43	83	99
Philips	63	97	100		62	90	99
Royal Dutch	86	97	100		87	96	99
<u>Unilever</u>	<u>59</u>	<u>94</u>	<u>100</u>		<u>76</u>	<u>95</u>	<u>100</u>
Mean	60	91	99		67	91	99

This table shows data density during the overlapping period from 9:30-10:30 EST for those days that both exchanges are open. The data set for New York included many quotes from regional exchanges, these were removed from the set because they were not competitive. The quotes originating in New York were virtually always inside these quotes. Quote density for New York is calculated based on the remaining quotes for the period starting at the time of the first quote until 10:30.

Table 3: Correlation in Returns Across Markets

	$\rho (r_{\text{midquoteAMS}}, r_{\text{midquoteNY}})$			
(σ)	KLM	Philips	Royal Dutch	Unilever
Interval Length				
1 min	0.10* (0.01)	0.21* (0.01)	0.25* (0.01)	0.19* (0.01)
5 min	0.39* (0.02)	0.60* (0.02)	0.72* (0.02)	0.64* (0.02)
15 min	0.65* (0.04)	0.82* (0.04)	0.85* (0.04)	0.80* (0.04)

*: Significant at a 99% confidence level

This table shows correlation in midquote returns for both markets during the overlapping period from 9:30-10:30 EST. Standard deviations are in brackets.

Table 4: Transient or Persistent Price Changes?

(σ)	Amsterdam				New York				
	KLM	Philips	Royal Dutch	Unilever	KLM	Philips	Royal Dutch	Unilever	
1 min autocorrelation function									
1	0.03*	0.07*	0.01	0.02*	-0.02	0.10*	0.02	-0.01	
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	
2	0.03*	0.06*	0.05*	0.03*	0.00	0.07*	0.03*	0.00	
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	
3	0.01	0.06*	0.03*	0.04*	0.00	0.07*	0.01	0.02	
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	
4	0.02	0.02	0.00	0.00	0.01	0.02	0.00	0.00	
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	
5	-0.02	0.01	0.01	-0.01	-0.01	0.04*	0.00	0.00	
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	
N	13,796	13,980	14,100	14,150	10,255	10,392	10,484	10,617	
$\left(\frac{5 \text{ min var}}{5 * (1 \text{ min var})} \right)$	1.17*	1.22*	1.12*	1.10*	0.99	1.28*	1.07*	1.02	
$(F_{0.01}^a, F_{0.99}^a)$	(0.98,1.04)	(0.98,1.04)	(0.98,1.04)	(0.98,1.04)	(0.98,1.1)	(0.98,1.1)	(0.98,1.1)	(0.98,1.1)	
$N1^b$	2,758	2,796	2,820	2,830	2,011	2,042	2,091	2,111	
$\left(\frac{15 \text{ min var}}{15 * (1 \text{ min var})} \right)$	1.15*	1.46*	1.09*	1.14*	0.89*	1.61*	1.00	0.98	
$(F_{0.01}^a, F_{0.99}^a)$	(0.95,1.04)	(0.95,1.04)	(0.95,1.04)	(0.95,1.04)	(0.98,1.03)	(0.98,1.03)	(0.98,1.03)	(0.98,1.03)	
$N1^b$	918	932	940	943	643	656	693	696	
$\left(\frac{45 \text{ min var}}{45 * (1 \text{ min var})} \right)$	1.38*	1.94*	1.45*	1.57*	0.90*	2.04*	1.11*	1.12*	
$(F_{0.01}^a, F_{0.99}^a)$	(0.9,1.05)	(0.9,1.05)	(0.9,1.05)	(0.9,1.05)	(0.95,1.03)	(0.95,1.03)	(0.95,1.03)	(0.95,1.03)	
$N1^b$	228	233	235	235	224	228	232	234	
$\left(\frac{60 \text{ min var}}{60 * (1 \text{ min var})} \right)$	1.28*	1.99*	1.27*	1.47*					
$(F_{0.01}^a, F_{0.99}^a)$	(0.91,1.04)	(0.91,1.04)	(0.91,1.04)	(0.91,1.04)					
$N1^b$	228	233	235	235					

*: Significant at a 99% confidence level

^a: Although similar, the true distribution of the test statistic is not an F distribution. It has thinner tails because the variance estimates in the numerator and the denominator are based on the same data and thus not independent. The values shown here are critical values of the true distribution found through simulations.

^b: $N1$ is the number of observations used to calculate the numerator, $N2$ is easily inferred from $N1$

This table documents the autocorrelation function up to five lags for one-minute midquote returns. To study persistence or correction for periods longer than five minutes variance ratios are calculated where the x minute return variance is in the numerator and x times the one minute variance in the denominator. All estimates are based on midquote returns for the overlapping period from 9:30-10:30 EST for those days that both markets are open.

Table 5: Market Depth during and outside Overlap

1 minute intervals		Amsterdam				New York				
EST	9:00- 9:30	9:30- 10:00	10:00- 10:30	R ²	N	9:30- 10:00	10:00- 10:30	10:30- 11:00	R ²	N
(σ)										
KLM	1.07* (0.03)	1.10* (0.02)	1.06* (0.02)	0.25 (0.00)	18,657	2.35* (0.09)	1.92* (0.07)	1.73* (0.10)	0.09	17,929
Philips	0.37* (0.01)	0.41* (0.01)	0.37* (0.01)	0.29	19,410	0.41* (0.02)	0.27* (0.01)	0.36* (0.02)	0.07	18,107
Royal Dutch	0.24* (0.01)	0.27* (0.00)	0.26* (0.00)	0.29	19,612	0.29* (0.01)	0.28* (0.01)	0.32* (0.01)	0.11	18,993
Unilever	0.31* (0.01)	0.33* (0.01)	0.33* (0.01)	0.31	19,596	0.33* (0.02)	0.34* (0.02)	0.34* (0.02)	0.06	18,819

*: Significant at a 99% confidence level

This table shows the results of estimates of a "reduced form" market microstructure model regressing one-minute midquote returns on signed volume. By allowing for different coefficients for different times of day we measure market depth for half hour intervals from 9:00 to 11:00 EST. This period includes the overlapping period as well as periods during which either Amsterdam or New York is the only market open. Standard deviations are in brackets.

Table 6: Correlation in Volume Across Markets

(σ)	Volume		
	Unchanged	Scaled by Daily Volume	(i) Scaled by Daily Volume and (ii) Demeaned by Time of Day
1 minute			
KLM	0.10* (0.01)	0.10* (0.01)	0.11* (0.01)
Philips	0.11* (0.01)	0.04* (0.01)	0.18** (0.01)
Royal Dutch	0.09* (0.01)	0.08* (0.01)	0.09* (0.01)
Unilever	0.13* (0.01)	0.07* (0.01)	0.08* (0.01)
5 minute			
KLM	0.38* (0.02)	0.19* (0.02)	0.21* (0.02)
Philips	0.26* (0.02)	0.11* (0.02)	0.13* (0.02)
Royal Dutch	0.20* (0.02)	0.16* (0.02)	0.24* (0.02)
Unilever	0.30* (0.02)	0.13* (0.02)	0.18* (0.02)

*: Significant at a 99% confidence level

This table shows correlation in volume across markets during the overlapping period from 9:30-10:30 EST. The first column shows the correlation in original volume. The second column scales volume by the volume witnessed for the entire day. This corrects for the fact that high volume days in Amsterdam are likely to coincide with high volume days in New York. This creates an upward bias in the test statistic that seeks to verify the presence of simultaneous trading in both markets within(!) the overlapping hour. The third columns not only scales by daily volume but also corrects for a downward bias due to a different trend in the intraday volume pattern on both exchanges during the overlap. Although both exhibit a U-shape this causes an upward trend for Amsterdam since the overlap is at the end of the day and a negative trend for New York since the overlap is at the start of the day. This effect is corrected for by demeaning volume by time of day. Standard deviations are in brackets.

Table 7: Splitting Orders or Arbitrage?

EST 9:30 - 10:30 (σ)	Fraction of Intervals			$\rho(\text{Arb_Opp}, \text{SignVolume}_{\text{AMS}} -$			
	1 min	5 min	N	1 min	N	20 sec ^a	N
Prediction Theory 'Splitting Orders' ('Arbitrage')							
Arb_Opp = 0		+ (0)		+ (0)		+ (0)	
Arb_Opp > x, x in basis points		+ (-)	0	+ (-)		+ (-)	
Signed Volume Correlation Across Markets Conditional on Arbitrage Opportunity at Start of Interval							
KLM							
Arb_Opp = 0	70%	0.11*	1,705	0.09*	8,624	-0.01	7,846
		(0.02)		(0.01)		(0.01)	
0< Arb_Opp <=15	15%	0.20*	344	0.08*	1,811	-0.02	2,132
		(0.05)		(0.02)		(0.02)	
15< Arb_Opp <=30	9%	0.15	245	0.02	1,162	0.01	1,326
		(0.06)		(0.03)		(0.03)	
30< Arb_Opp	5%	0.36*	114	0.01	662	0.01	1,028
		(0.09)		(0.04)		(0.03)	
Philips							
Arb_Opp = 0	74%	0.16*	1,750	0.06*	8,907	0.01	16,545
		(0.02)		(0.01)		(0.01)	
0< Arb_Opp <=15	19%	0.18*	442	0.04	2,252	0.03	4,601
		(0.05)		(0.02)		(0.01)	
15< Arb_Opp <=30	5%	0.27*	120	-0.04	631	-0.02	1,460
		(0.09)		(0.04)		(0.03)	
30< Arb_Opp	3%	0.19	64	0.06	326	-0.01	876
		(0.13)		(0.06)		(0.03)	
Royal Dutch							
Arb_Opp = 0	68%	0.24*	1,735	0.08*	8,841	0.02*	21,410
		(0.02)		(0.01)		(0.01)	
0< Arb_Opp <=15	23%	0.36*	620	0.11*	2,983	0.05*	7,727
		(0.04)		(0.02)		(0.01)	
15< Arb_Opp <=30	6%	0.34*	135	0.11*	781	0.02	2,195
		(0.09)		(0.04)		(0.02)	
30< Arb_Opp	3%	-0.03	66	-0.04	361	-0.12*	1,268
		(0.12)		(0.05)		(0.03)	
Unilever							
Arb_Opp = 0	76%	0.16*	1,899	0.07*	9,635	0.03*	15,130
		(0.02)		(0.01)		(0.01)	
0< Arb_Opp <=15	18%	0.22*	449	0.08*	2,259	-0.01	5,266
		(0.05)		(0.02)		(0.01)	
15< Arb_Opp <=30	4%	0.16	107	0.07	557	0.03	1,612
		(0.10)		(0.04)		(0.02)	
30< Arb_Opp	2%	-0.14	55	0.22*	272	0.03	1,101
		(0.13)		(0.06)		(0.03)	
No Arbitrage Opportunities Observed during Entire Interval, i.e. Max(Arb_Opp)=0							
KLM							
	63%	0.16*	1,525	0.15*	7,669	-0.01	7,276
		(0.03)		(0.01)		(0.01)	
Philips							
	60%	0.17*	1,414	0.05*	7,211	0.02	15,061
		(0.03)		(0.01)		(0.01)	
Royal Dutch							
	49%	0.30*	1,246	0.12*	6,340	0.03*	18,662
		(0.03)		(0.01)		(0.01)	
Unilever							
	53%	0.14*	1,347	0.06*	6,785	0.03*	13,131
		(0.03)		(0.01)		(0.01)	

*: Significant at a 99% confidence level

^a: Only those observations where volume is nonzero in both markets are included

This table documents correlation in signed volume across markets. The upper part of the table conditions on the value of "Arb_Opp" at the start of the interval. This variable indicates the presence of arbitrage opportunities since it is zero if there are none and it is equal to the signed relative difference between the bid in one market and the ask in the other in case there are. The lower part of the table conditions on the absence of arbitrage opportunities throughout the entire interval.

Table 8: Order Flow Skewed to Market with Best Price?

1 minute intervals	ρ (Volume Difference, Price Difference)	
(σ)	(Buy Volume, Ask Price)	(Sell Volume, Bid Price)
KLM	-0.09* (0.01)	0.01 (0.01)
Philips	-0.03* (0.01)	0.04* (0.01)
Royal Dutch	-0.09* (0.01)	0.09* (0.01)
Unilever	-0.06* (0.01)	0.08* (0.01)

*: Significant at a 99% confidence level

This table shows the correlation between, on the one hand, Amsterdam buy volume minus New York buy volume and, on the other hand, the best ask price in Amsterdam minus the best ask price in New York, where the latter is translated to Dutch Guilder using intraday exchange rates. These differences are based on one-minute intervals. The same is done for sell volume and bid price.

Table 9: Trade Frequency by Order Size

5 minute intervals		Amsterdam			New York			Δ AMS on	Δ NY on
(σ)	EST	9:00- 9:30	9:30- 10:00	10:00-10:30	9:30- 10:00	10:00-10:30	10:30-11:00	NY Open	AMS Close
KLM	V<=100	0.36 (0.02)	0.36 (0.02)	0.41 (0.02)	0.65 (0.03)	0.50 (0.02)	0.46 (0.02)	0%	-8%
	100<V<=1,000	1.07 (0.04)	1.64 (0.07)	1.63 (0.05)	1.67 (0.07)	1.42 (0.06)	1.17 (0.05)	53%*	-18%
	1,000<V<=5,000	0.92 (0.03)	1.72 (0.05)	1.72 (0.05)	0.56 (0.03)	0.46 (0.02)	0.28 (0.02)	87%*	-39%*
	5,000<V<=25,000	0.27 (0.02)	0.48 (0.02)	0.54 (0.03)	0.10 (0.01)	0.08 (0.01)	0.04 (0.01)	81%*	-50%
	<u>V>25,000</u>	0.01 (0.00)	0.02 (0.00)	0.01 (0.00)	0.01 (0.00)	0.01 (0.00)	0.00 (0.00)	<u>175%</u>	<u>-78%</u>
	All Trades	2.62 (0.07)	4.21 (0.11)	4.31 (0.10)	2.98 (0.11)	2.47 (0.08)	1.95 (0.07)	61%*	-21%*
Philips	V<=100	1.60 (0.05)	1.63 (0.05)	1.82 (0.06)	1.95 (0.09)	1.39 (0.04)	1.13 (0.03)	2%	-19%*
	100<V<=1,000	2.58 (0.06)	3.23 (0.07)	3.53 (0.07)	3.97 (0.14)	2.89 (0.07)	2.30 (0.06)	25%*	-21%*
	1,000<V<=5,000	1.93 (0.05)	3.16 (0.07)	3.56 (0.08)	1.52 (0.06)	1.36 (0.04)	0.88 (0.03)	64%*	-36%*
	5,000<V<=25,000	1.44 (0.05)	2.22 (0.06)	2.57 (0.07)	0.40 (0.02)	0.46 (0.02)	0.28 (0.02)	55%*	-39%*
	<u>V>25,000</u>	0.07 (0.01)	0.10 (0.01)	0.14 (0.01)	0.04 (0.01)	0.04 (0.01)	0.02 (0.00)	<u>48%</u>	<u>-43%</u>
	All Trades	7.61 (0.14)	10.35 (0.18)	11.61 (0.20)	7.88 (0.26)	6.15 (0.12)	4.61 (0.10)	36%*	-25%*
Royal Dutch	V<=100	1.74 (0.04)	1.71 (0.05)	1.87 (0.05)	3.46 (0.12)	2.65 (0.06)	2.55 (0.06)	-2%	-4%
	100<V<=1,000	3.90 (0.06)	4.37 (0.08)	4.53 (0.08)	7.36 (0.15)	5.71 (0.09)	5.34 (0.09)	12%*	-6%
	1,000<V<=5,000	1.69 (0.04)	2.82 (0.06)	2.89 (0.07)	4.41 (0.08)	3.96 (0.07)	3.08 (0.06)	67%*	-22%*
	5,000<V<=25,000	2.21 (0.05)	4.06 (0.08)	4.46 (0.09)	1.45 (0.05)	1.58 (0.04)	1.17 (0.04)	83%*	-26%*
	<u>V>25,000</u>	0.18 (0.01)	0.34 (0.02)	0.39 (0.02)	0.17 (0.01)	0.10 (0.01)	0.06 (0.01)	<u>83%*</u>	<u>-45%*</u>
	All Trades	9.73 (0.13)	13.30 (0.19)	14.14 (0.19)	16.85 (0.29)	14.00 (0.17)	12.21 (0.16)	37%*	-13%*
Unilever	V<=100	1.21 (0.04)	1.40 (0.04)	1.40 (0.04)	1.25 (0.04)	1.28 (0.04)	1.11 (0.04)	16%	-13%
	100<V<=1,000	2.37 (0.05)	2.99 (0.06)	3.17 (0.07)	3.43 (0.08)	3.12 (0.07)	2.42 (0.06)	26%*	-23%*
	1,000<V<=5,000	1.53 (0.04)	2.43 (0.06)	2.63 (0.06)	1.63 (0.05)	1.60 (0.05)	1.09 (0.04)	59%*	-32%*
	5,000<V<=25,000	0.93 (0.04)	1.66 (0.06)	1.85 (0.06)	0.31 (0.02)	0.32 (0.02)	0.22 (0.02)	78%*	-31%*
	<u>V>25,000</u>	0.03 (0.01)	0.06 (0.01)	0.07 (0.01)	0.02 (0.00)	0.01 (0.00)	0.01 (0.00)	<u>100%</u>	<u>6%</u>
	All Trades	6.08 (0.10)	8.55 (0.13)	9.12 (0.13)	6.65 (0.12)	6.34 (0.11)	4.85 (0.09)	41%*	-24%*

*: Significant at a 99% confidence level

This table documents the average order flow composition for different times of day. It shows the five-minute average number of trades for five different size categories. The averages are calculated for half hour intervals from 9:00 to 11:00 EST thus including the overlapping period. The intraday jump in Amsterdam on the New York open is based on comparison of the period 9:00-9:30 and 9:30-10:00, the intraday drop for New York on 10:00-10:30 and 10:30-11:00. Significance is tested based on the difference in means for these periods. Standard deviations are in brackets.

Table 10: Intramarket Signed Volume Correlation

5 minute intervals (σ)	9:00 - 9:30			9:30 - 10:30			10:30 - 11:30		
	Philips	Royal Dutch	Unilever	Philips	Royal Dutch	Unilever	Philips	Royal Dutch	Unilever
Order Flow Correlation Amsterdam									
KLM	0.16*	0.22*	0.27*	0.16*	0.22*	0.19*			
	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)			
Philips		0.24*	0.33*		0.35*	0.36*			
		(0.03)	(0.03)		(0.02)	(0.02)			
Royal Dutch			0.34*			0.46*			
			(0.03)			(0.02)			
Order Flow Correlation New York									
KLM				-0.02	0.00	-0.02	-0.01	0.00	-0.02
				(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Philips					0.01	0.03		0.00	0.01
					(0.02)	(0.02)		(0.03)	(0.03)
Royal Dutch						0.07*			0.06
						(0.02)			(0.03)

*: Significant at a 99% confidence level

This table shows contemporaneous correlations in signed volume for each market both during the overlapping period and outside the overlap. It is based on five minute intervals.

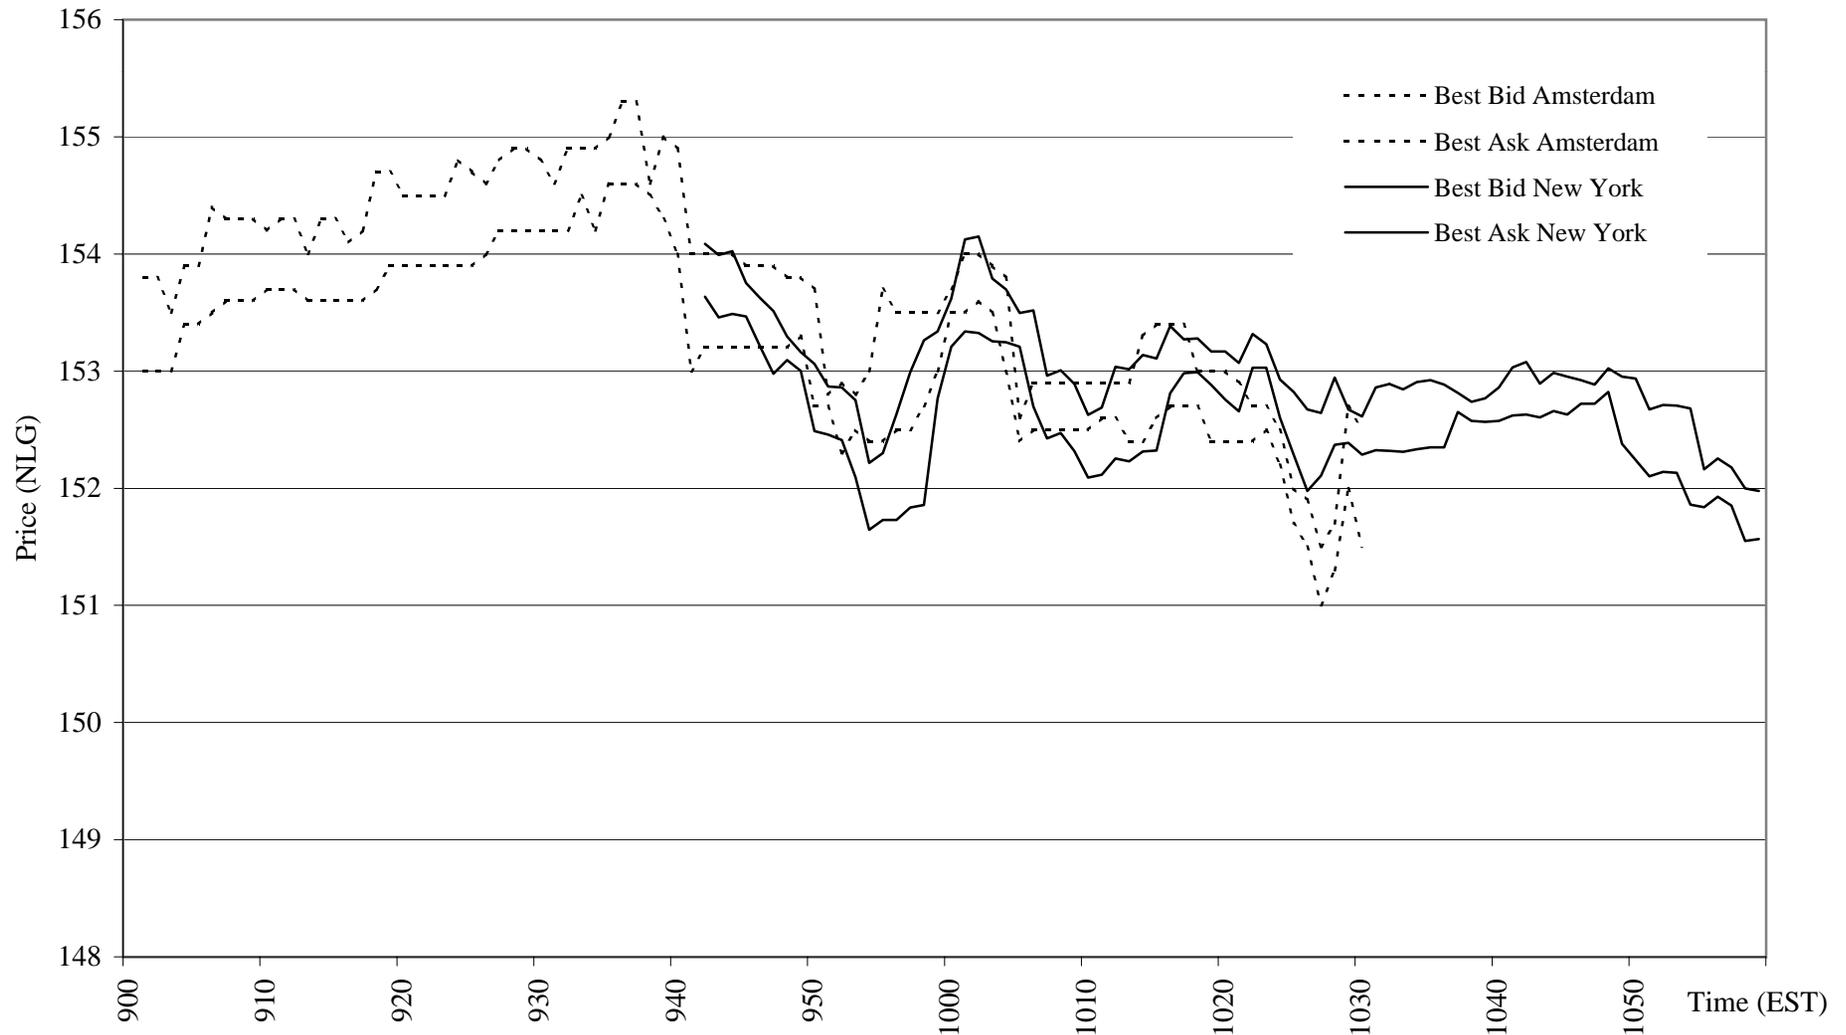
**Table 11: Intermarket Order Flow Correlation
Who is Splitting Orders?**

5 minute intervals	Order Flow Correlation Across Markets		
(σ)	Original Order Flow	= (i) Order Flow in Line with Market Order Flow	+ (ii) Order Flow Orthogonal to Market Order Flow
KLM	0.13* (0.02)	0.02 (0.02)	0.13* (0.02)
Philips	0.17* (0.02)	0.05 (0.00)	0.17* (0.02)
Royal Dutch	0.26* (0.02)	0.20* (0.02)	0.23* (0.02)
Unilever	0.16* (0.02)	0.08* (0.02)	0.15* (0.02)

*: Significant at a 99% confidence level

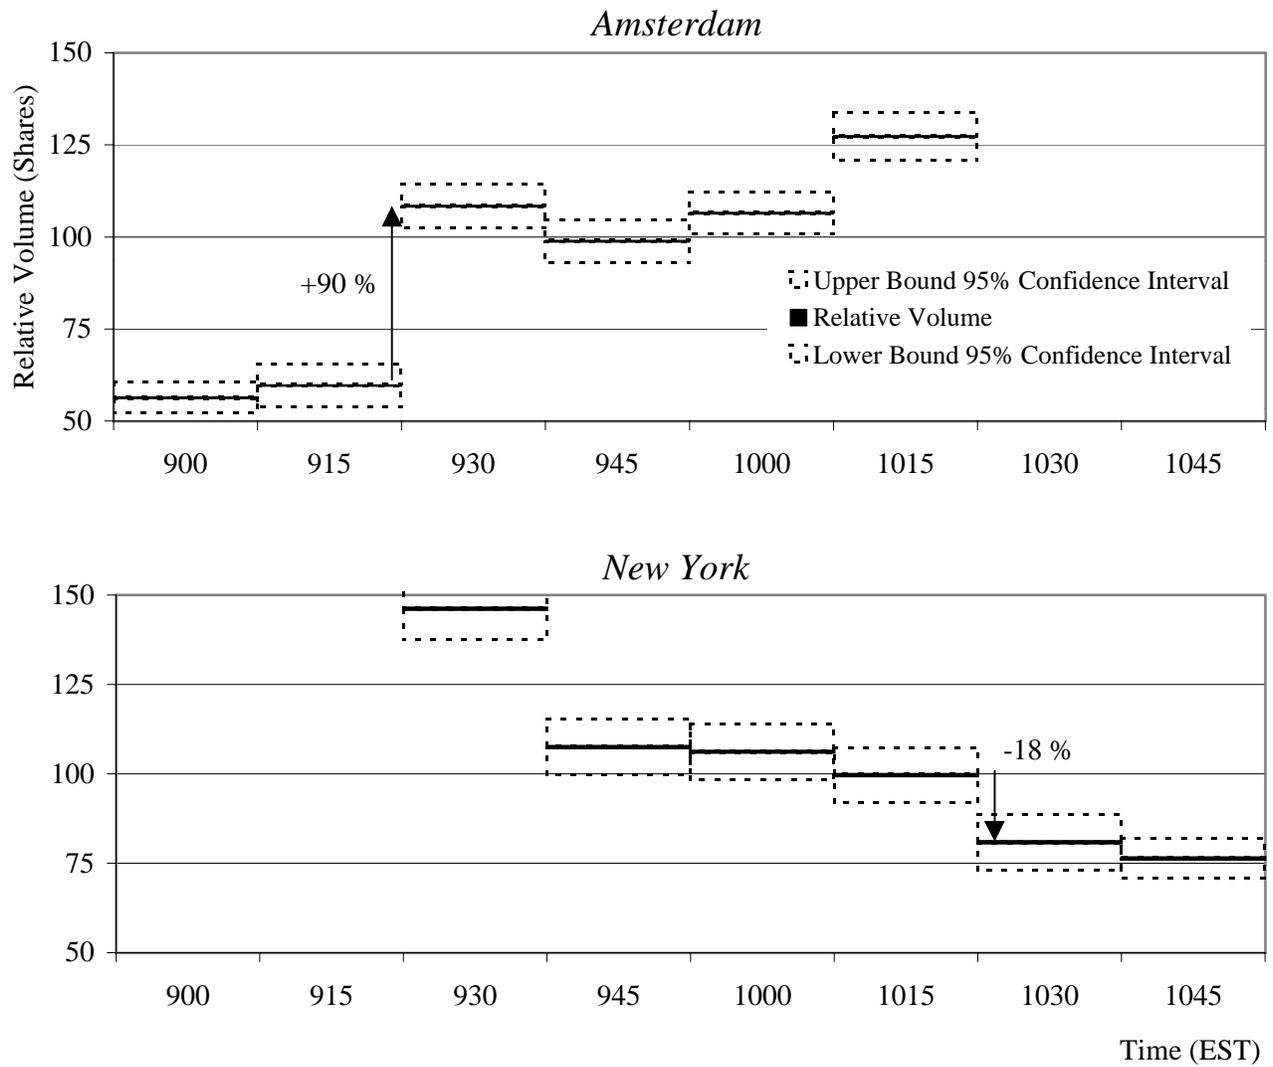
This table shows correlation in signed volume across market for five minute intervals. It decomposes signed volume in Amsterdam into (i) signed volume in line with the market or, alternatively, spanned by the market and (ii) signed volume orthogonal to the market. For both components the correlation with signed volume in New York is calculated.

**Figure 1: Intraday Price Discovery
Royal Dutch, October 27, 1997**



This figure shows the best bid and ask price in Amsterdam and New York for Royal Dutch. This figure reflects trading from 9:00 to 11:00 EST on October 27, 1997. This time period includes the overlapping period that runs from 9:30 to 10:30. The quotes are one minute snapshots. The New York quotes are translated to Dutch Guilders using the intraday exchange rate.

Figure 2: Intraday Volume Pattern for Royal Dutch



This figure depicts the results of least squares regressions that yield the intraday pattern in volume for Royal Dutch. These regressions are based on five-minute volume for the period from 9:00 to 11:00 EST, thus including the overlap. The dependent variable is volume scaled by the daily average.

Figure 3: The "Splitting Orders" and "Arbitrage" Hypotheses

Reduced Form Market Microstructure Model

Assumptions: (i) Spread equal to zero (ii) Market depth equal to one (iii) Markets identical

Definitions: ΔP = Price change, Q = Signed volume

Price Change due to Trade of $|Q|$ Shares: $\Delta P = 1 * Q$

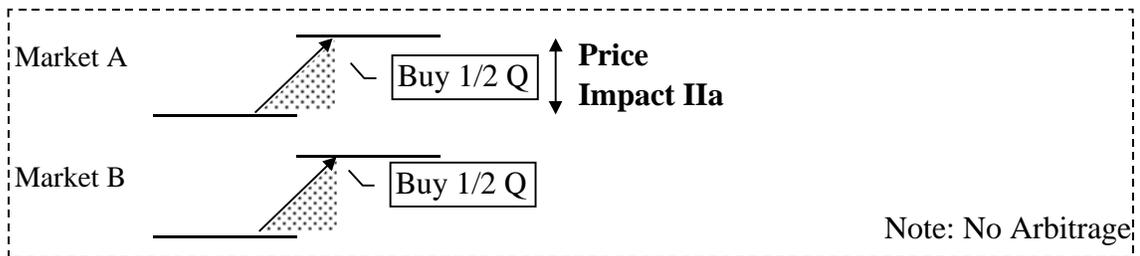
Traders' "Loss" due to price concession is the sum of the grey areas in each scenario

I Trading outside Overlap



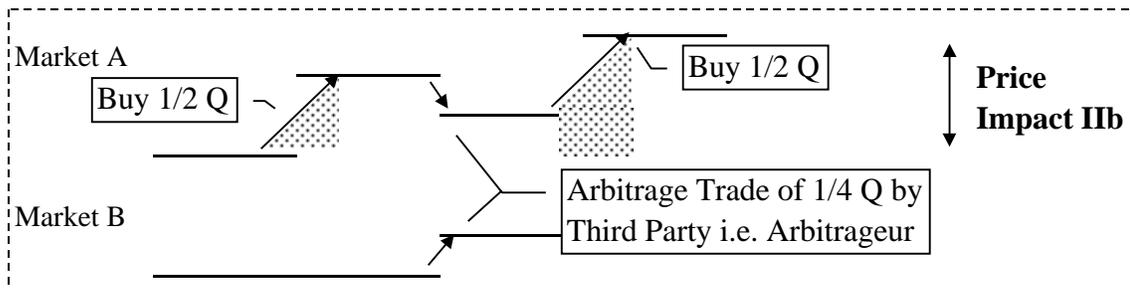
II Trading during Overlap

(a) "Splitting Orders": Limited price concession due to the trade being split across markets ...



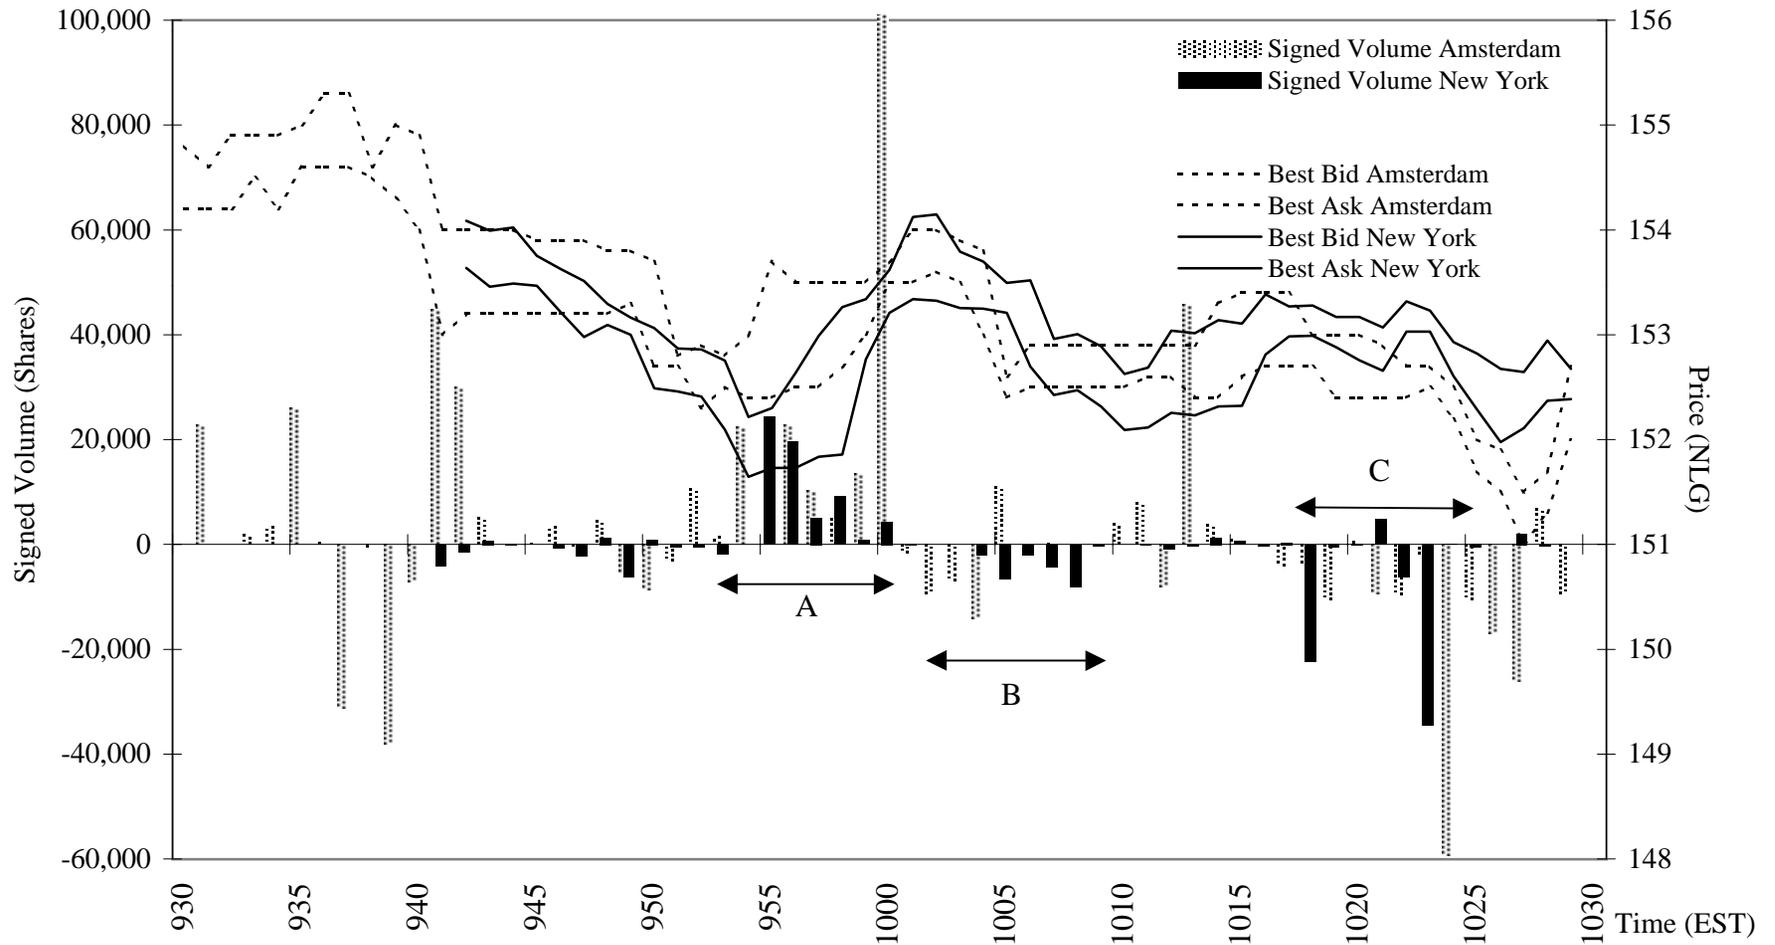
... or ...

(b) "Arbitrage": is the presence of an arbitrageur effectively deepening the market?



This figure illustrates the "splitting orders" and "arbitrage" hypotheses in a simplified setting involving a "reduced form" market microstructure model.

Figure 4: Intraday Price Discovery and Signed Volume
Royal Dutch, October 27, 1997



This figure shows best bid and ask prices and signed volume for Royal Dutch in Amsterdam and New York. This picture reflects trading during the overlap on October 27, 1997. The quote snapshots and signed volume are based on one minute intervals. The New York quotes are translated to Dutch Guilders using the intraday exchange rate.