Volatility regimes and the provision of liquidity in order book markets

Helena Beltran

Alain Durré

Pierre Giot*

December 22, 2004

ABSTRACT

We analyze whether the liquidity provision in a pure order book market during normal market conditions (low volatility regime) differs from what is observed when the market is under stress (high volatility regime). We show that the static relationship between liquidity and volatility is resilient to regime changes in volatility. Nevertheless, we do find that it is more costly to trade when volatility is large. A VAR analysis shows that the liquidity dynamics is similar in the low and high volatility regimes, although the drop in liquidity subsequent to volatility shocks is larger in the high volatility regime. Finally, the market is more resilient to volatility or liquidity shocks in periods of turmoils.

Keywords: order book, volatility, liquidity.

JEL classification: G10, C32.

---

1. Introduction

Modelling and appraising liquidity in financial markets has been of paramount importance for central banks, regulators and practitioners for the last decade. The perceived liquidity decrease during the financial crisis of 1998 has led many to question the functioning of stock markets during stress periods (see Borio, 2004, and references given therein). Moreover, the well-publicized problems of large hedge funds such as LTCM have also pointed out that liquidity could dry out rapidly during crisis periods, hence normal market conditions do not offer much information regarding what happens during volatile periods. As pointed out in the empirical and theoretical literature, liquidity depends crucially on the market structure. In price-driven markets (e.g. at the NASDAQ in its pre-1996 version or in bond and FOREX markets), a market maker ensures the continuity and viability of the trading process by quoting firm bid and ask prices whatever the market conditions. Thus, the inside spread (i.e. the difference between the best buy and sell prices) and depth at the best quotes seem to be good measures of the available liquidity, that is on an ex-ante basis. Ex-post, the liquidity of an exchange is often assessed by computing measures such as the effective or realized spread, or VWAP (volume-weighted average price) measures. Note that measures related to the liquidity displayed by the order book (which is fully described below) refer to pre-trade liquidity, and will correspondingly be referred to as ex-ante liquidity measures. Examples of such measures are the quoted spread and bid/ask depths. Measures computed with transaction data refer to realized trading costs, thus called ex-post liquidity measures. Effective spread is a well-known example of an ex-post liquidity measure.

In pure order-driven markets, no market maker stands ready to trade. Liquidity is thus provided by limit orders entered throughout the day by 'patient' or liquidity supplier investors (often value investors), and orders are executed only when prices match, i.e when liquidity is demanded by 'impatient' or liquidity demander investors. Examples of impatient traders include traders who wish to transact near the close of the trading session (so that the price of their trade is not far from the official closing price), see Cushing and Madhavan (2000), or momentum traders who are keen on entering immediate long or short positions (Keim and Madhavan, 1997). Therefore, the inside spread is not as relevant as in price-driven markets and depth outside the quotes (i.e. the complete state of the order book) and times between order entry and execution (the immediacy component) become crucial. As shown in Handa and Schwartz (1996), and discussed below, there exists a

dynamical equilibrium between limit order and market order trading which strongly determines the available liquidity of the order book.

While in a price-driven market the market makers ensure the continuity of the price process, no investors have to provide liquidity in order book markets.[1] As such, in order book markets, the available liquidity is the outcome of the yet non-executed limit orders posted by patient traders. It has been well documented that traders enter limit orders at a price that takes into account the so-called picking-off risk. Indeed, a limit order is more likely to be executed against an order submitted by an informed trader, or after adverse price changes. In the second case, the risk stems from the lack of monitoring of the trader's positions in the book, i.e. the trader did not cancel her order and new public information was disclosed after the submission of the order.[2] When trading becomes hectic, picking-off risk may rise to such a level that no market participant has an incentive to supply liquidity to the market. For example, Goldstein and Kavajecz (2004) highlight such a behavior at the NYSE. After two market-wide trading halts, the book was empty and most of the trading was redirected to the floor crowd. However, their results may not readily be extended to a pure limit order book platform as the presence of specialists at the NYSE offers an alternative to the book; this is not possible in a pure order book market.[3] Finally, events such as the Asian crisis show that the market response to high volatility can be markedly different from what is observed in normal periods, which suggests that the liquidity-volatility relationship may be non-linear. This is consistent with a picking-off risk that does not increase linearly with respect to price volatility. Therefore, it is not inconceivable that order book systems could break down in times of stress because the dynamical equilibrium of Handa and Schwartz (1996) between limit orders and market orders is disrupted.

In this paper we analyze the liquidity behavior *across different volatility regimes* for some stocks traded on the Euronext trading platform during the December 2, 2002 - April 30, 2003 time period. More precisely, we study whether liquidity supply and demand are more responsive to price changes when volatility is large. If this is indeed the case, then liquidity drops during volatile periods should be more pronounced that during normal market conditions and the behavior of market liquidity under regular market conditions cannot be extended to stress periods (and even less to extreme events). Thereafter, a period with a high level of volatility will be referred to as a 'stress pe-

---

[1]For example the specialists at the NYSE, which is best characterized as an hybrid market (i.e. a combination of a price-driven market, an order book market and a floor market), are required by the exchange to maintain an 'orderly market'.

[2]Or previously private information was made public.

[3]Indeed the authors stress that, for the smallest cap stocks in their sample, the liquidity drain from the book to the floor trading was not so pronounced, probably because the floor crowd is smaller for the smallest cap stocks.

riod', while low volatility periods are referred to as 'normal periods'. Note that, in contrast to most papers (e.g. Goldstein and Kavajecz (2004)) dealing with high volatility periods, we do not focus on one (or succession of) extreme event(s). In our empirical analysis, using high-frequency trade and order book data, we analyze the liquidity and volatility exhibited by some large- and mid-cap Brussels traded stocks on the Euronext platform. We study more particularly the contemporaneous relationships between volatility and ex-ante/ex-post liquidity to ascertain if the high volatility was associated with decreases in liquidity and large trading costs. From an econometric point of view, the low and high volatility regime states are determined according to an endogenous classification rule based on Markov switching models. Besides the ex-ante and ex-post assessment of liquidity, we also estimate VAR models for some of the variables measured on an intraday basis. Thereafter, we assess the impulse response functions derived from these estimated VAR models and analyze the dynamics of liquidity. Because we choose stocks for which there are no market makers, we thus shed light on the ex-ante and ex-post liquidity vs volatility relationships in a pure automated auction market.

The results (which also include an event study for one of the stock) indicate that while ex-ante or ex-post trading costs somewhat increase with volatility, the (static) relationship between liquidity and volatility does not undergo fundamental changes and the trading process does not break down. The dynamical analysis based on the VAR model offers a balanced view according to which the volatility regime bears moderately on the dynamics of the liquidity provision. Liquidity dynamics in the low and high volatility regimes are similar, but traders 'overreact' in the high volatility regime as the decrease in liquidity in response to volatility shocks is larger in the high volatility regime. Finally, the market is more resilient to volatility or liquidity shocks in periods of turmoils. As such, our results seem to indicate that there is no real important deterioration in the provision of liquidity when volatility increases, although we do find that it is more costly to trade when volatility is high and that the market dynamics are somewhat affected.

The rest of the paper is structured as follows. After this introduction, we present a review of the literature in Section 2. The Euronext trading system and the dataset are discussed in Section 3. The first part of the empirical analysis is presented in Section 4, where we also provide an event study for the Delhaize stock. The second part of the empirical analysis (trading dynamics and VAR analysis) is presented in Section 5. Finally, Section 6 concludes.

2. Review of the literature

2.1. Automated auction markets, liquidity and volatility

The literature on market microstructure has traditionally focused on dealership markets. Indeed, most of the models surveyed in O'Hara (1995) focus on the behavior of market makers or deal with fixed costs, inventory costs or asymmetric information costs models in the framework of market maker based trading systems. Because of the growing popularity of automated auction systems in European countries or in the electronic trading systems in the United States, there is now a rapidly evolving literature on order book markets.[4] Most of the empirical studies in that field focus on the provision of liquidity in automated auction markets. Indeed, as no market makers stand ready to buy and sell the traded assets in this setting, the viability of pure electronic order book markets and the ability to trade at all times are far from ascertained. Crucially, the provision of liquidity in times of crisis is of paramount importance. We thereafter survey some of the recent empirical work that focuses on the provision of liquidity in order book markets, the relationship between volatility and liquidity and the characteristics of automated auction markets in times of crisis.

In an important extension of pure dealership markets, automated auction markets allow a relatively easy ex-ante characterization of liquidity beyond the inside bid-ask spread. Because the state of the order book is usually fully or partially made available to market participants, price impact curves (i.e. unit bid and ask prices for a given volume, also called costs of buy and sell trades by Irvine, Benston, and Kandel, 2000) can be computed which allow the computation of extended liquidity measures such as the cost of buy or sell trades. These measures, popularized in Irvine, Benston, and Kandel (2000), Martinez, Tapia, and Rubio (2000), Coppejans, Domowitz, and Madhavan (2002) or Beltran, Giot, and Grammig (2004), aggregate the status of the order book at any given time and offer a relatively accurate picture of the available ex-ante liquidity, i.e. before the submission of a buy or sell trade.[5]

In a now seminal paper, Biais, Hillion, and Spatt (1995) provide one of the first empirical analysis of a limit order book market (the Paris Bourse). They study the joint dynamics of the order flow (placement of market or limit orders) and the order book: investors place limit (market) orders

---

[4]See the book by Harris (2002).

[5]In dealership markets, the ex-ante available liquidity often reduces to the best bid and ask prices (or quoted spread), and the available depth at these prices. Effective spreads or realized spreads are ex-post liquidity measures as they are computed after the submission of the buy or sell trade.

when the bid-ask spread is large (small) or the order book is thin (thick). Therefore, "investors provide liquidity when it is valuable to the marketplace and consume liquidity when it is plentiful". They also show that there is a strong competition among traders (who monitor the state of the order book) to provide liquidity as the flow of order placements is concentrated at or inside the bid-ask quote and the corresponding limit orders are placed in quick succession. For stocks traded on the pure electronic limit order platform of the Hong Kong stock exchange, Ahn, Bae, and Chan (2001) investigate the 'ecological' nature of the pure order driven market such as put forward in Handa and Schwartz (1996). They show that there exists a dynamical equilibrium between limit order trading and volatility: market depth rises subsequent to increases in volatility and volatility declines subsequent to increases in market depth. Indeed volatility attracts the placement of limit orders (instead of market orders) which therefore add liquidity to the order book. They also show the need to separate volatility at the ask and bid sides of the order book: when volatility arises from the ask (bid) side, investors submit more limit sell (buy) orders than market sell (buy) orders. On a related topic and for NYSE stocks, Bae, Jang, and Park (2003) show that it is important to distinguish between transitory and informational volatility: "a rise in transitory volatility induces a new placement of limit orders. A rise in informational volatility appear neither to increase nor decrease the placement of limit orders relative to market orders". Using a Probit model applied to Swiss stocks traded on the Swiss Stock Exchange, Ranaldo (2004) presents quite similar results: orders are more aggressive (i.e. traders submit more marketable limit orders than just plain limit orders) when the order queue on the incoming trader's side of the book is larger. For example, buyers then face a smaller execution probability and have to raise their order aggressiveness. The opposite is true for sellers. Moreover, volatility and larger spreads imply weaker trading aggressiveness. Frino, McInish, and Toner (1998) study the intraday pattern of the spread for two competing structures (a floor and an order book) where the same asset is traded (German Bund futures) and conclude that the order book provides more liquidity than the floor, although the performance of the automated auction market deteriorates when volatility increases (however they only focus on one liquidity variable, the quoted spread). Note that these studies do not focus on times of crises and it is thus not clear whether they would get similar results when trading is hectic.

Danielsson and Payne (2001) study the dynamics of liquidity supply and demand in the Reuters D2000-2 order book trading system.[6] They focus on the interaction between market and limit orders and show that the probability of a limit buy (sell) order is relatively low after a market sell

---

[6]The Reuters D2000-2 system is an electronic order book system designed for inter-dealer FOREX trades.

(buy). Therefore, there could be strong fluctuations in the provision of liquidity because of the complex interplay between market and limit orders (what they call dynamic illiquidity). In agreement with Foucault (1999), they show that the fraction of limit orders in total order arrivals increases with volatility (which increases liquidity), although the bid-ask spread also increases with volatility (which decreases liquidity). Hence, increases in volatility yield wider bid-ask spreads and lead to the increased placement of limit order relatively far from the quote mid-point. They also show that market participants react strongly to the unanticipated component of volume (predictable volume increases liquidity, unpredictable decreases liquidity). This hints at the importance of asymmetric information in automated auction markets and suggests the need for extensions of the models by Glosten and Milgrom (1985) and Easley and O'Hara (1987).

Goldstein and Kavajecz (2004) focus on the liquidity provision at the New York Stock Exchange during extreme market crises. Indeed, they deal with the very short time period that surrounds October 27, 1997. On that day, the Dow Jones Industrial Average lost 554 points (which triggered the circuit breakers) and on October 28, 1997 the index shot up by 337 points. They examine the liquidity supplied by the limit orders (routed by the SuperDOT order book trading system) and by the NYSE market participants (specialists and floor brokers). They show that a substantial liquidity drain occurred on the day after the market crash (i.e. on October 28, 1997) as the order book exhibited continuous large spreads and poor depth. However, the overall market liquidity did not drop dramatically as the specialists and floor brokers fulfilled their functions of liquidity providers and thus ensured good overall depth and low spreads at the NYSE. This hints at the adequacy of hybrid[7] market structures and shows that the viability of pure automated auction markets in times of crisis can be threatened by the significant drop in liquidity due to the substantial fall in the number of limit orders entered in the trading system.[8] Finally, Venkatamaran (2001) also stresses the merits of hybrid trading structures which lead to reduced trading costs.

Most empirical studies thus conclude that the 'ecological' nature of the pure order driven market works quite well: traders enter limit orders when liquidity is needed and are more impatient when liquidity is plentiful. Automated auction markets are also quite cheap to run, and bid-ask spreads for small to medium trades are quite low (see also Degryse, 1999). It is however not clear whether these results hold in all circumstances. Indeed, almost all studies (Goldstein and Kavajecz (2004) being

---

[7]A hybrid trading structure combines features of order book markets (the existence of a centralized order book run by a computer system) and of dealership markets (the existence of market makers or floor brokers). A good example of such a structure is the NYSE, see for example Bauwens and Giot (2001) or Sofianos and Werner (2000).

[8]Note however that it is 'easy' for traders to avoid the order book as they know that they can rely on the specialist in a hybrid trading structure.

the exception) on automated auction markets focus on the provision of liquidity in normal periods, i.e. not in times of crisis. In that latter case, liquidity could rapidly deteriorate if the sole provision of liquidity comes from limit orders (i.e. in the absence of hybrid systems that allow some provision of liquidity by market makers).

2.2. The regulator's point of view

During the nineties, the growing concern in monetary economics has been the opportunity for monetary authorities to react to stock market crashes and financial distress. The key question is whether Central Banks should have a prudential role in targeting financial stability. This is clearly relevant as history is plentiful of periods where financial instability involved macroeconomic instability. According to this paradigm, a financial crisis, because it acts for instance on the solvability of financial intermediaries, could affect the activity of firms through credit rationing, also called the credit crunch.[9] See Bernanke and Gertler (1995) and Bernanke and Gertler (1999). A key underpinning of this rationale is that price and financial stability are symbiotic in order to maintain a sustainable non inflationary growth. Solow argued that Central Banks should aim for financial stability as a larger risky asset volatility increases the probability of failure for financial institutions. If the Central Bank does not include the financial stability criterium as a monetary policy target, an increasing number of failures is to be expected which would be costly for the economy.[10] In the same vein and because the potential vulnerability of financial systems increases the probability of huge and costly crises (Borio, 2003), Borio and Lowe (2002) call for the inclusion of a financial target along the usual macroeconomic targets. Of course, this issue is controversial because it implies an ability to determine an equilibrium level for financial prices (Cecchetti, Genberg, Lipsky, and Wadhwani, 2000 and Borio and Lowe, 2002) discuss this problem). Note also that theoretical models and discussions of this problem (Bernanke and Gertler, 1999 and Borio and Lowe, 2002) yield ambiguous results and raise the issue of the choice of the financial asset whose price must be monitored.

Recently, Mishkin and White (2002) analyze fifteen historical episodes of stock market crashes in the US and examine the aftermath of these crises. Interestingly, their study suggests that stock market crashes by themselves do not involve financial instability. They show that the state of the financial system and the nature of stress in financial markets seem to be important. In particular,

---

[9]An example is the constraint in the real activity in Japan from 1992 onwards.

[10]See Bernanke and Gertler (1999), Cecchetti, Genberg, Lipsky, and Wadhwani (2000) or Durré (2003) for a discussion.

rapidly falling stock prices in conjunction with decreasing liquidity may be particularly destabiliz-ing.[11] Therefore, as a first step towards a prudential role of monetary authorities, a comprehensive understanding of the dynamics of liquidity during stress periods is warranted. Moreover, a closer look at the relationship between liquidity and volatility in times of crisis sheds light on key issues relevant to Central Banks and regulators in the future.

3. The Euronext platform and the dataset

3.1. Trading on the Euronext platform

Euronext encompasses five exchanges, namely the Amsterdam, Brussels, Lisbon, Paris exchanges and the LIFFE. Euronext aims to put forward a unique electronic trading platform for all financial assets. This is already the case for equities trading, as the same trading platform is now used by all exchanges. Trading on the Euronext platform takes place from 9 to 5.25 p.m CET. Limit or-ders are matched according to the standard price and time priority rules. Market orders (also called marketable limit orders) are executed against the best (in terms of price) prevailing order on the opposite side of the book. If there is not enough volume to fully execute the incoming order, the remaining part of the order is transformed into a limit order at the best price. Traders can also use more sophisticated orders, e.g. fill-or-kill orders (the limit order is either fully executed or can-celled), must-be-filled orders (the market order is completely executed, whatever the price), iceberg orders (part of the volume is not displayed in the book),... Block trading is allowed for large volume trades (the size of these trades is larger than the stock specific minimum block size, called "Taille normale de bloc"). Although the block trade formally takes place outside the book (akin to the up-stairs market at the NYSE), the transaction price is actually constrained by the available liquidity in the book. Indeed, Euronext displays throughout the day the hypothetical prices for a sell and a buy order with a volume equal to the minimum block volume. No blocks can be traded at a price outside these limits. Besides block trades, Euronext also allows so-called iceberg (or hidden) orders. As the name suggests, a hidden limit order is not (fully) visible in the order book. This implies that if a market order is executed against a hidden order, the trader submitting the market order may receive an unexpected price improvement. As on other automated auction exchanges (XETRA, Toronto

---

[11]Mishkin and White (2002) point out for instance that the responses of the Federal Reserve during the stock market declines in 1929 and 1987 were more appropriate than during the recent decline which began in 2000.

stock exchange,...), iceberg orders have been allowed to heed the request of investors who were reluctant to see their (potentially large) limit orders openly revealed in the order book.[12]

At the start of the trading day and before the regular continuous trading, a pre-opening auction takes place: limit orders are submitted and a start-of-day auction sets the opening price; all orders not executed at the end of the opening period remain in the order book.[13] Throughout the trading day, achievable trade prices are bounded by a static and a dynamic price limit. The static bounds are set immediately after the opening auction: they are equal to the auction price +- 10%. During the day, if a trade takes place outside these static bounds, trading is stopped and a new auction takes place (for a time period of 5 minutes). This auction final price defines new static bounds, used thereafter. The second type of bounds are dynamic: a trade cannot take place at a price larger (smaller) than the last trade price plus 2% (minus 2%). If orders can be matched at a transaction price outside the dynamic bounds, the trade is not executed and trading is stopped. A new auction takes place and defines new static and dynamic bounds. A final auction occurs between 5.25 and 5.30 p.m., followed by an additional 10-minute period where traders can trade at the price set by the end-of-day auction.

Note that, depending on the stock, two different Euronext members are involved in the trading process: brokers (called "Négociateurs") and market makers (called "Animateurs de marché"). All stocks do not feature a market maker. Indeed, stocks belonging to the Euronext 100 index (the first 100 Euronext stocks which have the largest market capitalizations) don't feature any market maker. Nevertheless, market makers are still allowed to enter orders for these stocks, but then they are considered as simple brokers.

3.2. The dataset

We were granted access to two historical datasets (for Brussels-traded stocks over a period ranging from December 2, 2002 to April 30, 2003) by Euronext. The first dataset contains the limit order book (LOB) as available to market participants who are not formally Euronext members, i.e. the historical real-time feed of the 5 best orders (price, total volume at that price and number of standing limit orders at that price) on the bid and ask sides of the order book. Indeed, all order book events (order entry, cancellation,...) are time-stamped to the second and lead to a potential order book

---

[12]See D'Hondt, De Winne, and Francois-Heude (2002) for a description of hidden orders on the Euronext trading platform.

[13]See also Biais, Hillion, and Spatt (1999).

modification, which is recorded in real-time by Euronext. We thus have snapshots of the 5 best bid and ask limit orders in real-time over the historical period we work with. It should however be stressed that the hidden portion of the iceberg orders is not included in the dataset. As discussed below, this will impact some of our conclusions (regarding the available ex-ante liquidity in the order book for example), while others should not be affected (the ex-post assessment of trading costs for example). The second dataset contains all transactions, more specifically the prices and volumes of the trades time-stamped to the second. Moreover, we also know if the orders matched in the transaction were so-called client or proprietary orders (the two most frequent cases), or market maker orders (a third possibility).[14] Note that the LOB dataset sometimes contains errors as the ordering of prices is not always enforced (e.g. the best ask price is sometimes larger than the ask price ranked second). These errors amount to less than 2% percent of all LOB observations and are removed from the dataset. Furthermore, the trades dataset did not give any information on the side (buy or sell) from which the trade originated. By using the LOB data, we are however able to determine rigourously the sign of the trade, as trades can only occur at the prices displayed in the book. Thus we did not have to rely on the Lee and Ready (1991) algorithm as used by most authors who work with NYSE data.

In this study we focus on five large-cap Belgian stocks (Dexia, Electrabel, Fortis, Interbrew and KBC) and five mid-cap Belgian Stocks (Agfa, Colruyt, Delhaize, Solvay and UCB).[15] The first five stocks are characterized by a large trading activity and are well-known blue-chip stocks widely held by individual and institutional investors. The five mid-cap stocks are also quite actively traded stocks. All ten stocks are members of the BEL20 stock index (which features the most 'representative' stocks of the Belgian economy) and no market maker (also called "animateur de marché") is involved in the trading of any of these stocks. Descriptive characteristics for the selected stocks are given in Table I. The stock price for Fortis (whose pattern is representative of what can be observed for the other stocks, save for the Colruyt stock which has been strongly trending upwards since the end of 2002) is plotted in the top left of Figure 2. This figure shows that the stock price was in a sharp downtrend till the start of the second Gulf war and that the market sharply recovered thereafter. In the remaining of the paper, we will focus on two representative stocks, Fortis and Colruyt, and for all the other stocks 'on average'. Colruyt (Fortis) is the smallest (largest) cap stock

---

[14]A so-called client order is an order routed to a Euronext member for execution by an outside investor. A proprietary order is executed by a Euronext member for his own trading account.

[15]This classification of large- and mid-cap stocks is relevant for average European investors. Some US investors would consider all these stocks to be only mid-cap stocks, and some of these even almost small-cap stocks.

in our sample. The results for the other stocks are similar except when we specifically state that there is a difference.

4. Empirical analysis

4.1. The importance of intraday seasonality

Most empirical studies on high-frequency data (Engle and Russell, 1998; Bisière and Kamionka, 2000; Bauwens and Giot, 2001; Bauwens and Giot, 2003) stress the need to correctly model the intraday seasonality exhibited by this kind of data. Indeed, when modelling the volatility, the traded volume, or the spread on an intraday basis, it is of paramount importance to proceed along a four-step procedure: (1) define regularly time-spaced measures of interest (e.g. working at the 15-minute frequency, the 15-minute return volatility, the 15-minute traded volume, the average effective spread over the 15-minute interval,...); (2) compute the time-of-day pattern for each measure; (3) deseasonalize each measure by its respective time-of-day pattern; (4) model the deseasonalized variable using an econometric model. Failure to recognize the importance of steps (2) and (3) often lead to incorrect model estimations (see Andersen and Bollerslev, 1997, for an application to the modelling of intraday volatility). Moreover there is also an economic justification in the modelling of the intraday seasonality. Because market participants are actively involved in the day-to-day market action, they know and expect a given pattern of activity (or volatility, spread,...) and are only affected by deviations (or surprises) from what was expected. A well-known example is the reaction of economic agents to news announcements: by itself, the news (e.g. the CPI number in the US) is not really relevant; what matters is the difference between the actual number and the expected number (see e.g. Bauwens, Ben Omrane, and Giot, 2003 or Andersen, Bollerslev, Diebold, and Vega, 2003).

As an illustration, we plot in Figures 1 and 2 the time-of-day pattern for the annualized volatility of the 15-minute returns, the 15-minute traded volume, the ask inside depth, the current and effective spread and the price impacts for the bid side. Current spreads, effective spreads and the ask inside depth are defined as usual, see Harris (2002).[16] Price impacts capture the premium paid by traders

---

[16]Note that, because we deal with an automated auction market, the effective spread can be larger than the current spread, as some transactions walk up the book and thus transaction prices are larger than the quoted spread.

when the transaction is executed against standing limit orders beyond the best quotes. Formally, the average price paid per share for a sell of a v shares at time t is

$$b_t(v) = \frac{\sum_{i=1}^{k} p_i v_i}{\sum_{i=1}^{k} v_i} \tag{1}$$

with $\sum_{i=1}^{k} v_i = v$, and $p_i(v_i)$ the ith bid price (volume) available in the book. The bid price impact is then defined as

$$bp_t(v) = \frac{b_t(1) - b_t(v)}{b_t(1)} * 100. \tag{2}$$

The same formula is used for the ask side. A bid price impact of , say, 2% for a trade of 10,000 shares means that each share costs 2% more than if all shares had been traded at the best bid price. By construction, the larger the transaction (i.e. the larger v), the larger the price paid as the market order hits more and more limit orders and is likely to walk up further in the book. We compute the price impacts for a volume v equal to 0.5, 1, 1.5 and 2 times a reference volume (the corresponding price impacts are labelled price impact level 1 to 4). The reference volume corresponds to a transaction for a nominal amount of 30,000 euros divided by the average price over the sample; this provides an easy comparison across stocks. Price impacts measure liquidity as offered by the book on an ex-ante basis, i.e before the transaction takes place. Finally, we also compute the trade aggressiveness, which measures how much traders use the book. It is computed as the volume-weighted average of the trades which are matched by standing limit orders strictly beyond the best quotes.[17] When trading volume rises, trade aggressiveness can remain low if the book provides more liquidity to the market (thus quoted depths rise).

The first figure is for Colruyt (the smallest cap stock in our sample and the least active in terms of average number of transactions per day), while the second figure is for Fortis (the largest cap stock in our sample and the most active in terms of average number of transactions per day). While both stocks are markedly different in terms of market cap and trading activity (see Table I), the time-of-day patterns are, broadly speaking, quite similar. As suggested by either Figure 1 or 2 (which are representative of the patterns exhibited by the other eight stocks, smallest cap stocks are closest to Colruyt, largest cap stocks are more similar to Fortis), trading activity tends to be large at the start of the day, slows down around lunch time and picks up in the middle of the afternoon session. This latter afternoon activity is consistent with the well-known influence of the pre-opening and opening of the US stock markets on the dynamics of the European markets. At the start of the trading session

---

[17]For example, a buy trade must be matched with at least one standing limit order above the best ask price to be characterized as being 'aggressive'.

on Euronext, the volatility, along with the traded volume for some stocks, is particularly high. For most of the stocks, traded volume increases at the end of the day while the increases in volatility appear more subdued. As far as the order book is concerned, it provides a reasonable amount of liquidity throughout the day. Although the spreads and price impacts are larger at the opening, they remain small and stable throughout the trading day, with a slight deterioration near the close of the trading session.[18] Depths at the quotes are up by roughly 50% in the afternoon compared with the morning. Moreover and although trading volumes are larger after 2 pm, transaction costs, as measured by the effective spread, are quite low and constant, except at the start of the trading session where traders have to pay twice the price they pay during the rest of the day. Indeed, despite the increase in traded volume in the afternoon, the provision of liquidity by limit orders in the book seems to avoid a sharp increase in transaction costs. These empirical facts are consistent with the previous findings reported in the literature, although it is well known that exchanges somewhat differ in their time-of-day pattern at the close of trading (see Biais, Hillion, and Spatt (1995) for the Paris Bourse, Beltran, Giot, and Grammig (2004) for the Frankfurt XETRA platform, Hamao and Hasbrouck (1995) for the Tokyo Stock Exchange, Chan, Christie, and Schultz (1995) for the NASDAQ or Brockman and Chung (1998) and Ahn and Cheung (1999) for the Hong Kong stock exchange).

4.2. The contemporaneous relationship between volatility and the liquidity measures

The main goal of the paper is to study how market conditions and liquidity are affected by volatility. As discussed above, what matters for market participants are deviations from expected volatility, hence the need to focus on the deseasonalized volatility and its contemporaneous relationship with (deseasonalized) liquidity measures. While the raw data are first pre-sampled at the 15-minute frequency (to define the 15-minute returns and to compute the time-of-day patterns as given above for example), we thereafter focus on 4 sub-intervals which span one trading day: [9h:11h], [11h:13h], [13h:15h] and [15h:17h30]. The [9h:11h] interval is just after the market open, [11h:13h] ends with the traders' lunchtime, [13h:15h] ranges from the start of the afternoon trading up to the New York pre-open and [15h:17h30] should capture the increased activity due to the opening of the US markets and ends with the close of trading on the Euronext platform. Besides, the switch from 15-

---

[18]Note that the figures for the ask side of the book are very similar to those presented for the bid side. Hence they are not given here but are available on request.

minute intervals to 2-hour intervals is consistent with the notion of realized volatility (see below) as a volatility measure computed from the 'aggregation' of really high-frequency squared returns. As such, estimation results (see the log-log regressions below) from models where the volatility is the independent variable should be less noisy. With respect to these 4 intervals, we thus compute the realized volatility, aggregated effective spread, aggregated quoted spread, aggregated trade aggressiveness, aggregated ask and bid depths and different measures of the aggregated ask and bid price impacts (as defined above). We now proceed with the definition of these aggregated measures computed from the data sampled at the 15-minute frequency.

First and following Andersen and Bollerslev (1998) or Giot and Laurent (2004), we define the realized volatility as the sum of the intraday squared returns which pertain to the required intervals. As shown in Andersen and Bollerslev (1998), the realized volatility measure provides a model-free estimation of return volatility over a given time interval (provided that high-frequency returns are available). For example, with 15-minute returns and for the [11h-13h] time interval, the realized volatility on December 2, 2002 is computed as:

$$RV_{13h,2/12/02} = r^2_{11h15} + \ldots + r^2_{13h} \tag{3}$$

where r_{11h15} is the 15-minute return for the [11h-11h15] time interval and r_{13h} is the 15-minute return for the [12h45-13h] time interval on December 2, 2002. The aggregated effective spread, quoted spread, trade aggressiveness, ask and bid depths and price impacts are respectively the mean effective spread, mean quoted spread, mean trade aggressiveness, mean depths and mean price impacts (computed from the 15-minute intervals) averaged over the [9h:11h], [11h:13h], [13h:15h] and [15h:17h30] intervals.

In a second step and for each interval, the time-of-day pattern of each measure is computed. Next we compute the deseasonalized variables by dividing each measure by its respective time-of-day. We then assess the contemporaneous relationships between the liquidity variable and the volatility by estimating the following log-log regressions:

$$ln(X_i) = \beta_0 + \beta_1 ln(RV_i) + \varepsilon_i, \qquad i = 1 \ldots N, \tag{4}$$

where X_i is successively S_i, Q_i, TA_i, DB_i, DA_i, BPI_i and API_i (respectively the deseasonalized aggregated effective spread, quoted spread, trade aggressiveness, bid depth, ask depth, bid price im-

14

pact and ask price impact), RV_i is the deseasonalized realized volatility and N is the total number of observations. Because we use log-log regressions, β_1 can be interpreted as an elasticity that 'links' the deseasonalized variables. The interpretation of these elasticities is as follows. For the $ln(S_i) = \beta_0 + \beta_1 ln(RV_i) + \varepsilon_i$ regression for example, a β_1 of 0.3 would imply that a 100% increase in the level of realized volatility (with respect to its expected level based on the time-of-day) would yield a 30% increase in the effective spread (with respect to its expected level based on the time-of-day).

Estimation results are given in the top panel of Table II. To save space, we present the full results for the Colruyt and Fortis stocks only. Next to these full results, we also report the results for 'all stocks' on average, that is we give the mean elasticity across stocks. As an illustration, we also plot the relationship between the deseasonalized aggregated effective spread and the deseasonalized realized volatility in Figure 3. The results indicate that the elasticity for the effective spread - realized volatility relationship is around 10% for Fortis, around 28% for Colruyt, and a bit more than 22% for the ten stocks on average. For the trade aggressiveness - realized volatility relationship, the elasticities range from a low of 7.4% (Fortis) to a high of 32.9% (Colruyt), with the average across stocks being equal to 15.7%. Figure 3 (and a similar figure for the trade aggressiveness, not reported here) also shows that there is no sharp deterioration in market liquidity when volatility increases sharply. Indeed a positive relationship between both effective spread and trade aggressiveness vs realized volatility is at play (which is expected from the market microstructure literature), but this positive dependence is somewhat muted (see below for additional discussions). As expected, the smallest cap stock (Colruyt) exhibits a more pronounced upward trending relationship than the largest cap stock (Fortis). Plots for the other stocks are quite similar to what is given in Figure 3.

The analysis for the quoted spread and depths yields similar results. Table III shows that the elasticity for the quoted spread - realized volatility relationship is around 12% for Fortis, and is equal to 24% on average for all the stocks. Furthermore, while there is a negative relationship between the depth (for both sides of the order book) and the realized volatility, it is not significant for most stocks. We also look at the relationship between the deseasonalized aggregated price impact (level 1 to 4) and the deseasonalized realized volatility (this last analysis thus uses information provided by the full limit order book dataset). The outputs of these log-log regressions are also given in Table III but we only present the level 3 price impacts to save space. In contrast to the bid and ask depths results, these level 3 elasticities are significant and are close to 15% on average for all stocks, with a

low around 9% to 10% for Fortis and close to 20% for Colruyt. The level 1, 2 and 4 elasticities and the detailed results for all the stocks deliver the same qualitative results.

4.2.1. High and low volatility regimes

Up to now we analyzed the whole bunch of observations put together, i.e. we did not deal separately with high volatility and low volatility time periods. Thereafter we split the deseasonalized measures defined on the [9h:11h], [11h:13h], [13h:15h] and [15h:17h30] intervals into a low volatility and high volatility subset. To construct the two sub-datasets, we apply a two-state Markov switching model (such as introduced by Hamilton, 1989) to the series of deseasonalized realized volatility. Using the smoothed transition probabilities, we can then immediately determine which observations belong to the low volatility regime and which ones can be put into the high volatility sub-dataset. This procedure is run for each stock, i.e. each stock features its own endogenous Markov switching classification. An alternative would be to use the same classification for all stocks, but then it would no longer be endogenous to the appropriate deseasonalized realized volatility. Moreover, there are relatively few cases when all stocks switch states at the same time (save around the start of the second Gulf war for the five largest cap stocks).

More formally, we assume that the deseasonalized realized volatility RV_i switches regime according to an unobserved variable s_i: regime 1 ($s_i = 1$) is the low volatility state, while regime 2 ($s_i = 2$) is the high volatility state. At time i, the volatility state is thus $s_i \in \{1, 2\}$ and the dynamics of s_i is governed by a Markov process: $P(s_i = 1|s_{i-1} = 1) = p_{11}$, $P(s_i = 2|s_{i-1} = 1) = 1 - p_{11}$, $P(s_i = 2|s_{i-1} = 2) = p_{22}$ and $P(s_i = 1|s_{i-1} = 2) = 1 - p_{22}$, where p_{11} (p_{22}) is the probability of being in the low volatility (high volatility) state at time i given that the low volatility (high volatility) state is observed at time $i - 1$. In state m, the deseasonalized realized volatility is equal to μ_m, with variance σ_m^2. We estimate the parameters of the model using the MSVAR package (maximum likelihood, EM algorithm) of H.-M. Krolzig in the OX 3.3 econometric framework, which also computes the smoothed transition probabilities. Finally, these are used to separate the observations into the two sub-datasets. We then re-run the log-log regressions.

Estimation results for these regressions are given in the middle and bottom panels of Tables II and III. Let us consider first the effective spread and trade aggressiveness (Table II). A comparison of the elasticities in both regimes for all stocks on average indicates that the numerical values are close to one another both for the effective spread (21% versus 18%) and for the trade aggressiveness

16

(18% versus 15%). In all cases, the effective spread - realized volatility and trade aggressiveness - realized volatility elasticities do not significantly change when volatility switches from the low- to the high volatility state.[19] In other words, these relationships (which focus mainly on the ex-post liquidity or actual cost of trading) do not seem to significantly deteriorate in times of high volatility. These results are in agreement with the estimates for the limit order book dataset (quoted spread, bid and ask depths, bid and ask price impacts) provided in Table III. As reported, most elasticities are not significant and only the quoted spread elasticity is really significant during stress periods.[20] This suggests that the Euronext system provides adequate liquidity in both low- and high volatility regimes as the slopes of these key relationships do not change in a meaningful way. In contrast, a trading system with poor liquidity would be characterized by increasing elasticities as volatility increases, indicating that liquidity dries up in high volatility regimes.

If high volatility regimes do not significantly impact the elasticities of the effective spread - realized volatility and trade aggressiveness - realized volatility relationships, they do affect the *mean* (or expected value) of the effective spread and trade aggressiveness. These results are reported in Table IV. For Colruyt (the 'worst case' in terms of deterioration of liquidity during the high volatility regime), the effective spread goes up by 91% and trades are more aggressive (+32%) despite the decrease of liquidity in the book; price impacts (bid side, level 3) surged by 66% on average. This suggests that traders were somehow reluctant to enter large orders given the low liquidity offered by the book. As expected, figures for the largest cap stock (Fortis) indicate that the changes in the means are more moderate. Note that Dexia is the most liquid stock as the effective spread only increases by 10%. Broadly speaking and looking at the reported results for all the stocks (in detail and on average), the decrease in liquidity seems very reasonable when compared with the increase in the average volatility between the low- and high volatility regimes (nearly 400%). Moreover and given that the amount (in share volume) of the hidden orders (not featured in our database) on the Euronext trading platform is estimated at 30% of the total book (see D'Hondt, De Winne, and Francois-Heude (2002)), the argument according to which there is a sufficient liquidity provision seems to be valid.

---

[19]This was tested using regression analysis and appropriately defined dummy variables. Note that there is however one exception: the Delhaize stock, for the trade aggressiveness.

[20]A regression analysis with appropriately defined dummy variables again confirms that these elasticities are not significantly different in the two regimes, save for the Colruyt, Delhaize, Fortis and KBC stocks (ask and bid depths only).

4.2.2. Additional results and robustness checks

For the aggregated effective spread and aggregated trade aggressiveness, we also re-estimate some of the log-log regressions allowing for a quadratic effect, i.e. we include the squared independent variable as an additional explicative variable. We thus estimate:

$$ln(S_i) = \beta_0 + \beta_1 ln(RV_i) + \beta_2 (ln(RV_i))^2 + \varepsilon_i, \tag{5}$$

and

$$ln(TA_i) = \beta_0 + \beta_1 ln(RV_i) + \beta_2 (ln(RV_i))^2 + \varepsilon_i \tag{6}$$

where the variables are defined as before. For the 10 stocks and for both liquidity measures, the β_2 coefficient is however almost never significant (full numerical results are available on request). Finally, in a previous version of the paper, we also considered an exogenous volatility criteria: the low volatility subset featured the measures for which the realized volatility was within one standard deviation of its expected value ('average volatility' group) while the high volatility subset featured the intervals for which the realized volatility was beyond one standard deviation of its expected value ('above-average volatility' group). The estimation results were quite close to those shown above for the volatility criteria based on the Markov switching process and are therefore not included in this version of the paper. We also implemented a classification of volatility in 5 regimes, based on an exogenous criteria. Again, the results were similar to those presented here.

Last, it is well-known that traded volume and market activity are closely linked to observed market volatility. Therefore, it could perhaps be argued that our volatility regimes simply correspond to activity regimes. To assess this issue, we classified all periods according to a Markov switching process applied to the deseasonalized trading activity instead of the deseasonalized realized volatility. Contingency tables computed per stock show that both types of classification are not highly correlated. In this alternative specification, the left- and right-hand side variables in the log-log regressions are defined as before. We therefore estimate the $ln(X_i) = \beta_0 + \beta_1 ln(RV_i) + \varepsilon_i$ regressions in a state of low deseasonalized trading activity and high deseasonalized trading activity. The estimated elasticities are not only very similar to those given above for the switch based on the deseasonalized realized volatility but they are also very similar across the newly-defined regimes. This thus leads us to conclude similarly and lends further credence to the results given above.

4.3. An event study: the Delhaize stock on March 13, 2003

While the 10 stocks under review were sometimes quite volatile during the 5-month period (and in particular around the start of the second Gulf war), it would however be wrong to classify this time period as being 'extreme'. This has already been highlighted in the introduction of the paper, as our study is thus different from e.g. Goldstein and Kavajecz (2004) who focus on extreme events that led to (or preceded) temporary market shutdowns. There is however one market event in our five-month period that could be considered as being extreme and which could be related to Goldstein and Kavajecz (2004). More specifically, we shall thereafter discuss the behavior of the Delhaize stock on March 13, 2003, which provides a detailed case study of the book dynamics when faced with extremely large intraday price changes and buy-sell pressure.

The price pattern of the Delhaize stock is shown in Figure 4 for the five days centered on March 13, 2003. This figure shows that the price of the stock increased strongly on March 13, 2003, with an intraday variation close to +50% (the intraday high was close to 18 Euros, the closing price on March 12, 2003 was close to 12 Euros). As a result, the trading process for the Delhaize stock was automatically suspended for some brief period of time because of the very large intraday price changes.[21] The sequence of news events on that day reads as:

- 9h58: Delhaize reports net earnings (for the 2002 fiscal year) which have increased 19.3% year-on-year;

- 10h00: Delhaize reports an increased solvability, actually much better than what was expected by the market;

- 10h01: Delhaize declares a net dividend per share of 0.66 Euros. This amount was much larger than what was expected by the market;

- Thereafter in the day, many analysts upgrade the stock from 'Sell' to 'Neutral' or 'Buy'.

The very large intraday price changes for Delhaize were due to heavy buy pressure, as many market participants unwounded short positions (in the period preceding March 13, 2003, there were even rumors that Delhaize could go bankrupt) or acted according to the revised analysts recommendations. For the sake of our study, it is therefore interesting to look at the 'behavior' of the order book on March 13, 2003 and assess if these very large price movements were the result of poor market liquidity. More precisely we look at the deseasonalized 15-minute traded volume, number

---

[21]Technically speaking, the trading process was thus briefly interrupted because the maximum intraday price change limit was reached. See Section 3.1.

of trades, quoted spread, effective spread, trade aggressiveness, quoted bid and ask depths and bid and ask price impacts (level 2 and 4). All these market liquidity variables are plotted for the five days centered on March 13, 2003 in Figures 5 and 6. As expected, there was a very large buy-sell disequilibrium on March 13, 2003 and there was a surge in the number of trades and traded volume (both quantities are much larger than what is usually observed). However, the sub-figures for the dynamics of the (deseasonalized) quoted spread, effective spread, trade aggressiveness, book depths and price impacts all tell the same story: the book liquidity remained high and the trading process was not characterized by a hectic behavior. In other words, the surge in trading activity and the buy-sell disequilibrium indeed led to large intraday price changes (which is consistent with the fact that the economic agents adjust their view on the Delhaize stock given the public news announcements made by the company), but the book liquidity was adequate. This is perfectly exemplified by the sub-figure for the effective spread, which measures the actual trading costs faced by the market participants: as far as the effective spread is concerned, March 13, 2003 was a 'normal' trading day.

5. Trading dynamics

In this section we analyze how the volatility level affects the interplay between the main liquidity components (spreads, price impacts, average volume per trade,...) and the relationships between liquidity and volatility. Because this analysis hinges on the investigation of the dynamics of liquidity, we use Vector Autoregression (VAR) models applied to the original data sampled at the 15-minute frequency. The VAR analysis will thus first be performed on the whole dataset, and then on the subsets defined by the low- and high volatility states identified by the Markov switching model.[22]

5.1. VAR models and impulse response functions

We model the dynamics between liquidity and volatility using a VAR model. VAR models are to some extent a-theoretical, in the sense that there is no need to specify the economic relationships. Hence, we need to impose some restrictions on the estimated coefficients to reconstruct the underlying structural model. In this paper, we consider a VAR(p) model of the following type:

---

[22]The use of VAR models to analyze high-frequency equidistantly time spaced data has been advocated by Joel Hasbrouck, see e.g. Hasbrouck (1999).

$$X_t = \Gamma_0 + \sum_{i=1}^{p} \Gamma_i X_{t-i} + \varepsilon_t \qquad (7)$$

where X_t is the vector of endogenous variables and ε_t is the usual error term. With respect to the application considered in this paper, we estimate a 4-lag VAR (the lag dynamics is thus roughly equal to one hour as we work with 15-minute intervals), with 7 variables (6 of the 7 variables are endogenous and the last one is exogenous, see below). These variables, which have all been previously deseasonalized by their respective time-of-day as described previously, are:

- Liquidity ex-ante: quoted spread and the price impact for a trade of 45,000 euros (average of the ask and bid sides);

- Liquidity ex-post: effective spread and trade aggressiveness;

- Activity variables: number of trades and average volume per trade;

- Volatility.

We first test for block exogeneity of each of the variables and ascertain that only trade aggressiveness is exogenous at the 5% level. Thereafter, we thus estimate a VAR with 6 endogenous variables: quoted spread, average price impact for a transaction of 45,000 euros, number of trades, average volume per trade, the effective spread, and volatility; trade aggressiveness is the only exogenous variable. Using the BIC criteria, we further reduce the dimension of the system as it indicates that a 2-lag structure is adequate. Finally we estimate the selected VAR(2) model three times: first with all the observations, then with the observations belonging to the low volatility regime, and finally with the observations which pertain to the high volatility regime. Using a Cholesky decomposition, we further decompose the residuals ε_t to get a structural model:

$$X_t = \sum_{i=1}^{\infty} C_i e_{t-i} \qquad (8)$$

This decomposition ensures that the individual shocks are orthogonal, i.e. that the variance-covariance matrix $V(e_t)$ is diagonal. It also allows the system analysis of the impact of a one-period shock to a given variable, also called impulse response functions. We compute the 20-lag (5 hours, about half a trading day) impulse response functions for the VAR model estimated first with all the data, and then with the data provided by the low- and high volatility regime classification. For the first VAR(2) model as for the low- and high volatility regime VAR(2) models, we tried several endogenous variable ordering to ascertain that the choice of ordering did not lead to different results. The

impulse responses exhibit remarkably similar shapes whatever the ordering. This is important as it implies that the correlation between the individual shocks e_{jt} (where j denotes the j-th variable) is small and thus does not appear as important as in many macroeconomic structural models. The main argument as to why cross-correlations between shocks are large in macroeconomic models is that the data is typically monthly/quarterly and thus lagged response to a single shock within the month are aggregated and consequently treated as a contemporaneous impact when dealing with monthly data. This suggests that the chosen 15-minute interval is small enough to avoid aggregation issues.

In both regimes, most of the impulse responses (detailed plots are available on request) are significantly different from zero (flat IR), but there are no marked differences between the low- and high volatility regimes (see below for additional discussion). Moreover, the confidence intervals for the high volatility regime are larger than for the low volatility regime; for many impulse responses, the confidence intervals for the low volatility regime lie within the ones for the high volatility regime.[23] In all cases the width of the confidence intervals strongly decreases after 4 periods on average, i.e. roughly one hour. Furthermore, for volatility shocks and the impulse responses of a variable to its own shock, impulse responses are significantly different between regimes at the 95% confidence level. To improve on the impulse response analysis, we compute two additional statistics: half-life times and cumulated impacts. These are defined as follows. First, the cumulated impact of a shock is defined as the sum of the impulse responses over all 20 periods, it is thus the long-run impact of a permanent shock. Secondly, the half-life is the time needed to achieve half of the cumulated impact, it thus measures the speed of return to equilibrium.

5.2. Impulse response functions and the dynamics of liquidity

All relevant empirical results are summarized in Tables V and VI, which thus supplement/summarize the description given below. We analyze the impulse response functions using a three-step procedure. First, we look at their shapes as these are similar in low and high volatility regimes (see Table V). Then we focus on the differences in the amplitude of the response to volatility shocks. Finally, we present results regarding the half-life (Table VI). We put the variables in three groups: volatility, liquidity, both ex-ante (price impacts and the quoted spread) and ex-post (the effective spread), and trading activity (average volume per trade and trading frequency as measured by the number of trades).

---

[23]The number of observations in the low volatility regime is roughly twice the number of observations in the high volatility regime.

The VAR analysis first shows that the impulse responses are remarkably similar in the low and high volatility regimes. Volatility shocks drain both ex-ante and ex-post liquidity from the market. Although liquidity drops, volatility shocks also foster trading activity. These results are consistent with the literature: as the winner's curse rises with volatility, limit order traders want to be better rewarded for their provision of liquidity and the spread increases (see e.g. Foucault (1999)). This explains why, in both regimes, we find that liquidity drops when volatility surges. Moreover, if volatility is a proxy for the arrival of information, an increase in volatility should be related to larger and more frequent market orders. However, the response of the price impact to a volatility shock is not clear-cut (in the high volatility regime only): for 4 stocks out of 10, the price impact decreases when volatility rises, that is liquidity beyond the best prices is cheaper after a volatility shock. We revisit this issue in the following paragraph.

In the previous section we showed that the impulse response functions in the low and high volatility regimes are not statistically different, save for the volatility shocks and the impulse response of a variable to its own shock. Let us now focus on the differences in the amplitude of the volatility shocks. A one-standard-deviation volatility shock immediately increases the quoted spread by less than 16% in the low volatility regime, and by 16% to 35% (depending on the stock) in the high volatility regime. In the long run this is also the case as the quoted spread after a volatility shock in the high volatility regime is larger than the quoted spread after a volatility shock in the low volatility regime. The results are not so clear-cut for the liquidity provided beyond the best prices. We previously stressed that, for some stocks, the price impacts decrease after a volatility shock in the high volatility regime, whereas they tend to increase in the low volatility regime. Therefore, it is not surprising to find that, in the long run and for two third of the stocks, the book is deeper after a volatility shock in the high regime than in the low regime. Therefore, in contrast to what might be perhaps be expected, the provision of liquidity does not break down in volatile periods.

Finally, resiliency is a crucial issue for the continuity of the trading process. Table VI shows that liquidity is more resilient to volatility shocks in the high volatility regime than in the low volatility regime: half of the shock is incorporated into the spread after half an hour in the low regime, while 45 minutes are needed in the high volatility regime. For the price impacts, we have respectively half an hour and one hour. The results are consistent with a more active monitoring of limit orders by traders in periods where the winner's curse is more important. Nevertheless, traders update their orders' position in the book, whether at the best quotes or further in the book, more quickly in response to a shock in trading activity. To sum things us, our VAR analysis thus demonstrates that,

although shocks have larger impacts on average in high volatility regimes, they are still absorbed quite rapidly. Therefore, the VAR analysis supports the evidence presented in the previous section, i.e. that liquidity does not drain from the market during high volatility regimes.

6. Conclusion

Using limit order book and transaction data for five large-cap and five mid-cap Belgian stocks traded on the Euronext system, we study the contemporaneous relationships between different liquidity measures and market volatility. We consider ex-ante (e.g. quoted spread, bid and price price impacts, depths) and ex-post (e.g. effective spread, trade aggressiveness) liquidity measures which are fully available once complete order book data is at hand. From an econometric point of view, we work with the realized volatility as popularized recently in Andersen and Bollerslev (1998), which allows us to define a low and high volatility regime based on Markov switching techniques. Thereafter, we thus assess the different relationships (mostly modelled as log-log regressions on the appropriately defined liquidity measures) with all the data, and then separately in the low and high volatility regimes. Our analysis thus sheds light on the behavior of automated auction markets in times of low and high volatility, and allows us to quantify the impact of the switch in volatility regimes on key ex-ante and ex-post liquidity measures relevant for traders and/or institutional investors.

Our results indicate that the provision of liquidity in the Euronext trading system seems to be quite resilient to increases in volatility. Indeed, the (static) slopes of these liquidity measures - volatility relationships (e.g. effective spread - volatility or trade aggressiveness - volatility relationships for example) do not significantly change when volatility switches from the low volatility to the high volatility regime. In contrast, the mean (or expected value) of each liquidity measure is usually significantly higher in the high volatility state, but this was expected from the market microstructure literature. Moreover, the dynamic analysis confirms that liquidity behaves similarly in both regimes, although the amplitude of the response and the resilience of the market differs. As such, the main empirical result of this study is that there is no real important deterioration in the provision of liquidity when volatility increases, although we do find that it is more costly to trade when volatility is high.

As shown in many theoretical studies, adverse selection increases when volatility increases, which results in a more costly provision of limit orders. As a consequence, market liquidity drops

when volatility increases. In periods of financial distress, this is the one of the main concerns for central banks since this behavior may lead to the collapse or near-collapse of financial markets (e.g. the October 1987 market meltdown and the LTCM failure in 1998, among others). As recently suggested by Mishkin and White (2002), a financial crisis combined with a large drop in liquidity can be avoided with prompt and adequate action by monetary authorities. The concern about a systemic drop in financial liquidity, shared by many studies (see e.g. Borio and Lowe (2002) and Borio (2003)), leads many academics and practitioners to suggest that central banks should perhaps play a regulatory role in financial markets. However and given the many ways stock markets can be set up (pure order book market, price-driven market, hybrid market), the first natural step is to understand the dynamics of liquidity in stress vs normal periods for each type of market. While this kind of study had already been done for some price-driven markets or for some hybrid markets, no empirical study had yet focused on that topic for pure order book platform such as Euronext. Regarding the behavior of liquidity in high volatility regimes, the results presented in this paper are particularly promising. Indeed, even if trading costs are larger in stress periods, the trading system does not seem to break down.

Our results pave the way for additional research linked to that topic. An obvious extension would be to assess our relationships on an extended dataset which would feature a much larger number of stocks sub-divided into smaller groups based on the firms' characteristics. In this extended setting, we could thus quantify the possible deterioration in the provision of liquidity according to the most salient characteristics of the stock (e.g. small-cap, mid-cap, large-cap; type of industry;. . .). It could also be argued that the Markov switching algorithm should be applied to the overall market volatility (for example the volatility of the index) to study systemic risk. A crucial question for central bankers and exchanges is to ascertain which trading platform best performs during high volatility regimes. Extending the present methodology to allows comparisons of the liquidity supply and trading costs in specialist markets vis-a-vis order book markets would shed additional light on the role of market makers in hybrid structures.

References

Ahn, H., K. Bae, and K. Chan, 2001, Limit orders, depth and volatility: evidence from the stock exchange of Hong Kong, *Journal of Finance* 56, 767–788.

Ahn, H., and Y. Cheung, 1999, The intraday patterns of the spread and depth in a market with market makers: the Stock Exchange of Hong Kong, *Pacific-Basin Finance Journal* 7, 539–556.

Andersen, T.G., and T. Bollerslev, 1997, Intraday periodicity and volatility persistence in financial markets, *Journal of Empirical Finance* 4, 115–158.

Andersen, T.G., and T. Bollerslev, 1998, Answering the skeptics: yes, standard volatility models do provide accurate forecasts, *International Economic Review* 39, 885–905.

Andersen, T.G., T. Bollerslev, F.X. Diebold, and C. Vega, 2003, Micro effects of macro announcements: real-time price discovery in foreign exchange, *American Economic Review* 93, 38–62.

Bae, K., H. Jang, and S. Park, 2003, Traders' choice between limit and market orders: evidence from NYSE stocks, *Journal of Financial Markets* 6, 517–538.

Bauwens, L., W. Ben Omrane, and P. Giot, 2003, News announcements, market activity and volatility in the euro/dollar foreign exchange market, *Forthcoming in Journal of International Money and Finance*.

Bauwens, L., and P. Giot, 2001, *Econometric modelling of stock market intraday activity*. (Kluwer Academic Publishers).

Bauwens, L., and P. Giot, 2003, Asymmetric ACD model: introducing price information in the ACD model, *Empirical Economics* 28, 1–23.

Beltran, H., P. Giot, and J. Grammig, 2004, Modelling order book commonalities, Mimeo, University of Namur.

Bernanke, B.S., and M. Gertler, 1995, Inside the black box : the credit channel of monetary transmission, *Journal of Economic Perspectives* 9, 27–48.

Bernanke, B.S., and M. Gertler, 1999, Monetary policy and asset price volatility, Federal Reserve of Kansas City Conference on New Challenges for Monetary Policy.

Biais, B., P. Hillion, and C. Spatt, 1995, An empirical analysis of the limit order book and the order flow in the Paris Bourse, *Journal of Finance* 50, 1655–1689.

Biais, B., P. Hillion, and C. Spatt, 1999, Price discovery and learning during the preopening period in the Paris Bourse, *Journal of Political Economy* 107, 1218–1248.

Bisière, C., and T. Kamionka, 2000, Timing of orders, orders aggressiveness and the order book at the Paris Bourse, *Annales d'Economie et de Statistique* 60, 43–72.

Borio, C., 2003, Towards a macroprudential framework for financial supervision and regulation?, BIS Working Papers 128.

Borio, C., 2004, Market distress and vanishing liquidity: anatomy and policy options, BIS Working Papers 158.

Borio, C., and Ph. Lowe, 2002, Asset prices, financial and monetary stability: exploring the nexus, BIS Working Papers 114.

Brockman, P., and D.Y. Chung, 1998, Inter- and intra-day liquidity patterns on the Stock Exchange of Hong Kong, *Journal of International Financial Markets, Institutions and Money* 8, 277–298.

Cecchetti, S.G., H. Genberg, J. Lipsky, and S. Wadhwani, 2000, *Geneva reports on the World Economy 2*. (Centre for Economic Policy Research).

Chan, K.C., W.G. Christie, and P.H. Schultz, 1995, Market structure and the intraday evolution of bid-ask spreads for NASDAQ securities, *Journal of Business* 68, 35–60.

Coppejans, M., I. Domowitz, and A. Madhavan, 2002, Liquidity in an automated auction, Mimeo.

Cushing, D., and A. Madhavan, 2000, Stock returns and trading at the close, *Journal of Financial Markets* 3, 45–67.

Danielsson, J., and R. Payne, 2001, Measuring and explaining liquidity on an electronic limit order book: evidence from Reuters D2000-2, Financial Markets Group, London School of Economics.

Degryse, H., 1999, The total cost of trading Belgian shares: Brussels versus London, *Journal of Banking and Finance* 23, 1331–1355.

D'Hondt, C., R. De Winne, and A. Francois-Heude, 2002, Market liquidity on Euronext Paris: nothing is quite it seems, Working paper, FUCaM.

Durré, A., 2003, *Essays on the interaction between monetary policy and financial markets*. (Presses Universitaires de Louvain).

Easley, D., and M. O'Hara, 1987, Price, trade size and information in securities markets, *Journal of Financial Economics* 19, 69–90.

Engle, R.F., and J. Russell, 1998, Autoregressive conditional duration; a new model for irregularly spaced transaction data, *Econometrica* 66, 1127–1162.

Foucault, T., 1999, Order flow composition and trading costs in a dynamic limit order marker, *Journal of Financial Markets* 2, 99–134.

Frino, A., T.H. McInish, and M. Toner, 1998, The liquidity of automated exchanges: new evidence from German Bund futures, *Journal of International Financial Markets, Institutions and Money* 8, 225–241.

Giot, P., and S. Laurent, 2004, Modelling daily Value-at-Risk using realized volatility and ARCH type models, *Journal of Empirical Finance* 11, 379–398.

Glosten, L., and P. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71–100.

Goldstein, M.A., and K.A. Kavajecz, 2004, Trading strategies during circuit breakers and extreme market movements, *Journal of Financial Markets* 7, 301–333.

Hamao, Y., and J. Hasbrouck, 1995, Securities trading in the absence of dealers: trades, and quotes on the Tokyo stock exchange, *The Review of Financial Studies* 8, 849–878.

Hamilton, J.D., 1989, A new approach to the economic analysis of nonstationary time series and the business cycle, *Econometrica* 57, 357–384.

Handa, Puneet, and Robert Schwartz, 1996, Limit order trading, *Journal of Finance* 51, 1835–1861.

Harris, L., 2002, *Trading and exchanges*. (Oxford University Press).

Hasbrouck, J., 1999, The dynamics of discrete bid and ask quotes, *Journal of Finance* 54, 2109–2142.

Irvine, P., G. Benston, and E. Kandel, 2000, Liquidity beyond the inside spread: measuring and using information in the limit order book, Mimeo, Goizueta Business School, Emory University, Atlanta.

Keim, D.B., and A. Madhavan, 1997, Transaction costs and investment style: an inter-exchange analysis of institutional equity trades, *Journal of Financial Economics* 46, 265–292.

Lee, C., and M. Ready, 1991, Inferring trade direction from intraday data, *Journal of Finance*.

Martinez, M.A., M. Tapia, and G. Rubio, 2000, Understanding liquidity: a closer look at the limit order book, Working Paper 00-52, Universidad Carlos III de Madrid, Departamento de Economia de la Empresa.

Mishkin, F.S., and E.N. White, 2002, U.S. stock market crashes and their aftermath: implications for monetary policy, NBER working paper 8992.

O'Hara, M., 1995, *Market microstructure theory*. (Basil Blackwell Oxford).

Ranaldo, A., 2004, Order aggressiveness in limit order book markets, *Journal of Financial Markets* 7, 53–74.

Sofianos, G., and I.M. Werner, 2000, The trades of NYSE floor brokers, *Journal of Financial Markets* 3, 139–176.

Venkatamaran, K., 2001, Automated versus floor trading: an analysis of execution costs on the Paris and New York Stock Exchange, *Journal of Finance* 56, 1445–1485.

Table I
Descriptive statistics.

	DEXIA	ELECTRABEL	FORTIS	KBC	INTERBREW
Market capitalization	11.2	12.6	19.8	9.1	8.2
Average price	10	228.3	15.2	29.4	18.9
Average volume per trade	1695.2	159.7	1339.6	502.1	1048.4
Average number of transactions per day	719.4	329.3	876.3	354.4	411.3
% of aggressive trades	16.3	20.1	21.5	24.3	22.3
	AGFA	COLRUYT	DELHAIZE	SOLVAY	UCB
Market capitalization	2.5	1.9	2.1	4.9	3.6
Average price	19.7	53.9	16.7	59.6	24.7
Average volume per trade	757.2	205.2	663.9	302.6	524.6
Average number of transactions per day	284.7	136.8	315.3	261.7	275.6
% of aggressive trades	20.3	18.4	22.5	21.4	25.5

Descriptive statistics for the 10 Brussels-traded stocks selected in the empirical analysis. An aggressive trade is defined as a trade which is matched with standing limit orders beyond the best quote. Market caps are expressed in billions of euros and are the average market caps for the given stock over our sample.

Table II
Elasticities (I).

	N	Effective spread		Trade aggressiveness		Limit orders	
		All observations					
		Elasticity	Significant	Elasticity	Significant	Elasticity	Significant
COLRUYT	409	0.281	1/1	0.329	1/1	0.028	1/1
FORTIS	409	0.101	1/1	0.074	1/1	-0.005	0/1
ALL STOCKS	409	0.221	10/10	0.157	10/10	0.032	8/10
		Low volatility regime					
	N	Effective spread		Trade aggressiveness		Limit orders	
		Elasticity	Significant	Elasticity	Significant	Elasticity	Significant
COLRUYT	320	0.259	1/1	0.382	1/1	0.018	0/1
FORTIS	323	0.074	1/1	0.066	1/1	0.001	0/1
ALL STOCKS	-	0.213	10/10	0.184	10/10	0.032	7/10
		High volatility regime					
	N	Effective spread		Trade aggressiveness		Limit orders	
		Elasticity	Significant	Elasticity	Significant	Elasticity	Significant
COLRUYT	88	0.063	0/1	0.222	1/1	0.029	0/1
FORTIS	86	0.003	0/1	-0.041	0/1	-0.002	0/1
ALL STOCKS	-	0.175	7/10	0.150	4/10	0.005	0/10

Outputs of the log-log regressions where the dependent variable is successively the aggregated effective spread, the aggregated trade aggressiveness and the percentage of limit orders, the independent variable is the realized volatility in all cases. All measures are computed over the [9h:11h], [11h:13h], [13h:15h] and [15h:17h30] intervals and are deseasonalized (by their respective time-of-day) prior to running the regressions. The panel 'Low volatility regime' gives the outputs for the sub-datasets where the realized volatility is in the low volatility regime; the panel 'High volatility regime' gives the outputs for the sub-datasets where the realized volatility is in the high volatility regime. We provide results for the COLRUYT stock (smallest cap), FORTIS stock (largest cap) and for all stocks. In that latter case, we report the mean elasticity. The column 'Significant' reports the number of significant elasticities. The time period is December 2, 2002 to April 30, 2003.

Table III
Elasticities (II).

All observations

	N	Quoted spread Elasticity	Quoted spread Significant	Bid depth Elasticity	Bid depth Significant	Ask depth Elasticity	Ask depth Significant	Bid PI3 Elasticity	Bid PI3 Significant	Ask PI3 Elasticity	Ask PI3 Significant
COLRUYT	409	0.309	1/1	-0.041	0/1	-0.049	0/1	0.208	1/1	0.19	1/1
FORTIS	409	0.119	1/1	0.001	0/1	-0.061	1/1	0.104	1/1	0.083	1/1
ALL STOCKS	409	0.241	10/10	-0.024	2/10	-0.04	3/10	0.16	10/10	0.146	10/10

Low volatility regime

	N	Quoted spread Elasticity	Quoted spread Significant	Bid depth Elasticity	Bid depth Significant	Ask depth Elasticity	Ask depth Significant	Bid PI3 Elasticity	Bid PI3 Significant	Ask PI3 Elasticity	Ask PI3 Significant
COLRUYT	320	0.274	1/1	-0.12	1/1	-0.125	1/1	0.159	1/1	0.21	1/1
FORTIS	323	0.123	1/1	-0.059	0/1	-0.105	1/1	0.125	1/1	0.139	1/1
ALL STOCKS	-	0.235	10/10	-0.048	2/10	-0.058	3/10	0.162	10/10	0.135	7/10

High volatility regime

	N	Quoted spread Elasticity	Quoted spread Significant	Bid depth Elasticity	Bid depth Significant	Ask depth Elasticity	Ask depth Significant	Bid PI3 Elasticity	Bid PI3 Significant	Ask PI3 Elasticity	Ask PI3 Significant
COLRUYT	88	0.034	0/1	0.097	0/1	0.099	0/1	0.173	1/1	0.146	0/1
FORTIS	86	-0.021	0/1	0.14	1/1	0.141	1/1	-0.059	0/1	-0.137	0/1
ALL STOCKS	-	0.168	8/10	0.053	2/10	0.033	2/10	0.007	2/10	0.041	2/10

Outputs of the log-log regressions where the dependent variable is successively the quoted spread, the bid and ask depths, and the bid and ask price impacts (level 3), the independent variable is the realized volatility in all cases. All measures are computed over the [9h:11h], [11h:13h], [13h:15h] and [15h:17h30] intervals and are deseasonalized (by their respective time-of-day) prior to running the regressions. The panel 'Low volatility regime' gives the outputs for the sub-datasets where the realized volatility is in the low volatility regime; the panel 'High volatility regime' gives the outputs for the sub-datasets where the realized volatility is in the high volatility regime. We provide results for the COLRUYT stock (smallest cap), FORTIS stock (largest cap) and for all stocks. In that latter case, we report the mean elasticity. The column 'Significant' reports the number of significant elasticities. The time period is December 2, 2002 to April 30, 2003.

Table IV
% change in the mean of each deseasonalized variable when switching from the low- to the high volatility regime.

	Limit orders	Realized volatility	Effective spread	Trade aggressiveness	Quoted spread	Bid		Ask	
						Depth	PI3	Depth	PI3
COLRUYT	7.7	390.6	91	32.4	88.5	-1.8	66.1	-4.9	34.2
FORTIS	-1.7	337.3	23.6	11.8	21.3	2.8	16.5	-11.8	6.3
ALLSTOCKS	6.3	388.1	51.2	17.1	54	-3.2	34.1	-7.3	31.3

The table reports the % change in the mean of each deseasonalized variable when switching from the low- to the high volatility regime. Depth corresponds to the depth at the best prices. PI3 corresponds to the price impacts computed for a transaction of 45,000 Euros. The null hypothesis of the equality in means is rejected at the 5% level for all variables except for the bid depth.

Table V
The dynamics of liquidity.

shock	volatility		av. volu		nb tr		pi3		qsp		sp	
	L	H	L	H	L	H	L	H	L	H	L	H
volatility	+	+							+	+		
av. volu	+	+	+	+			-	-				
nb tr	+	+			+	+	-		-	-		
pi3	+	+	-				+	+	+			
qsp	+	+		+-	-	-	-+	-+	+	+	+	+
sp	+	+	+	+-	-	-			+	+	+	+

This table presents a summary of the VAR results (analysis of the dynamics of liquidity). A "+" ("-") means that the shock on the given variable (in the top row) has a positive (negative) and significant impact on the variable (in the first column). In a few cases, we report a "-+", which indicates that the shock is first negative and then positive. "H" refers to the high volatility state, while "L" refers to the low volatility state. Note that av. volu relates to the average volume per trade, nb tr, the number of trades, pi3, the level 3 price impact, qsp, the quoted spread, sp, the effective spread.

Table VI
Half-life of the impulse responses.

shock	volatility	av. volu	nb tr	pi3	qsp	sp
volatility	15				30	
av. volu	15	15		15		
nb tr	30 (45)		30	15	15	
pi3	1H (30 mn)	15		30 (15 mn)	1H30	
qsp	45 (30)	15	15	1H15	15	30
sp	1H (30)	15	15 (30)		45	15

This table reports the half-life of the impulse responses (averaged across the 10 stocks). All results are expressed in minutes except when there is an "H" (for hour). We only report results for the significant impulse responses. All results are for the low- and high volatility states, except when there is a number in parenthesis (high volatility state). Note that av. volu relates to the average volume per trade, nb tr, the number of trades, pi3, the level 3 price impact, qsp, the quoted spread, sp, the effective spread.
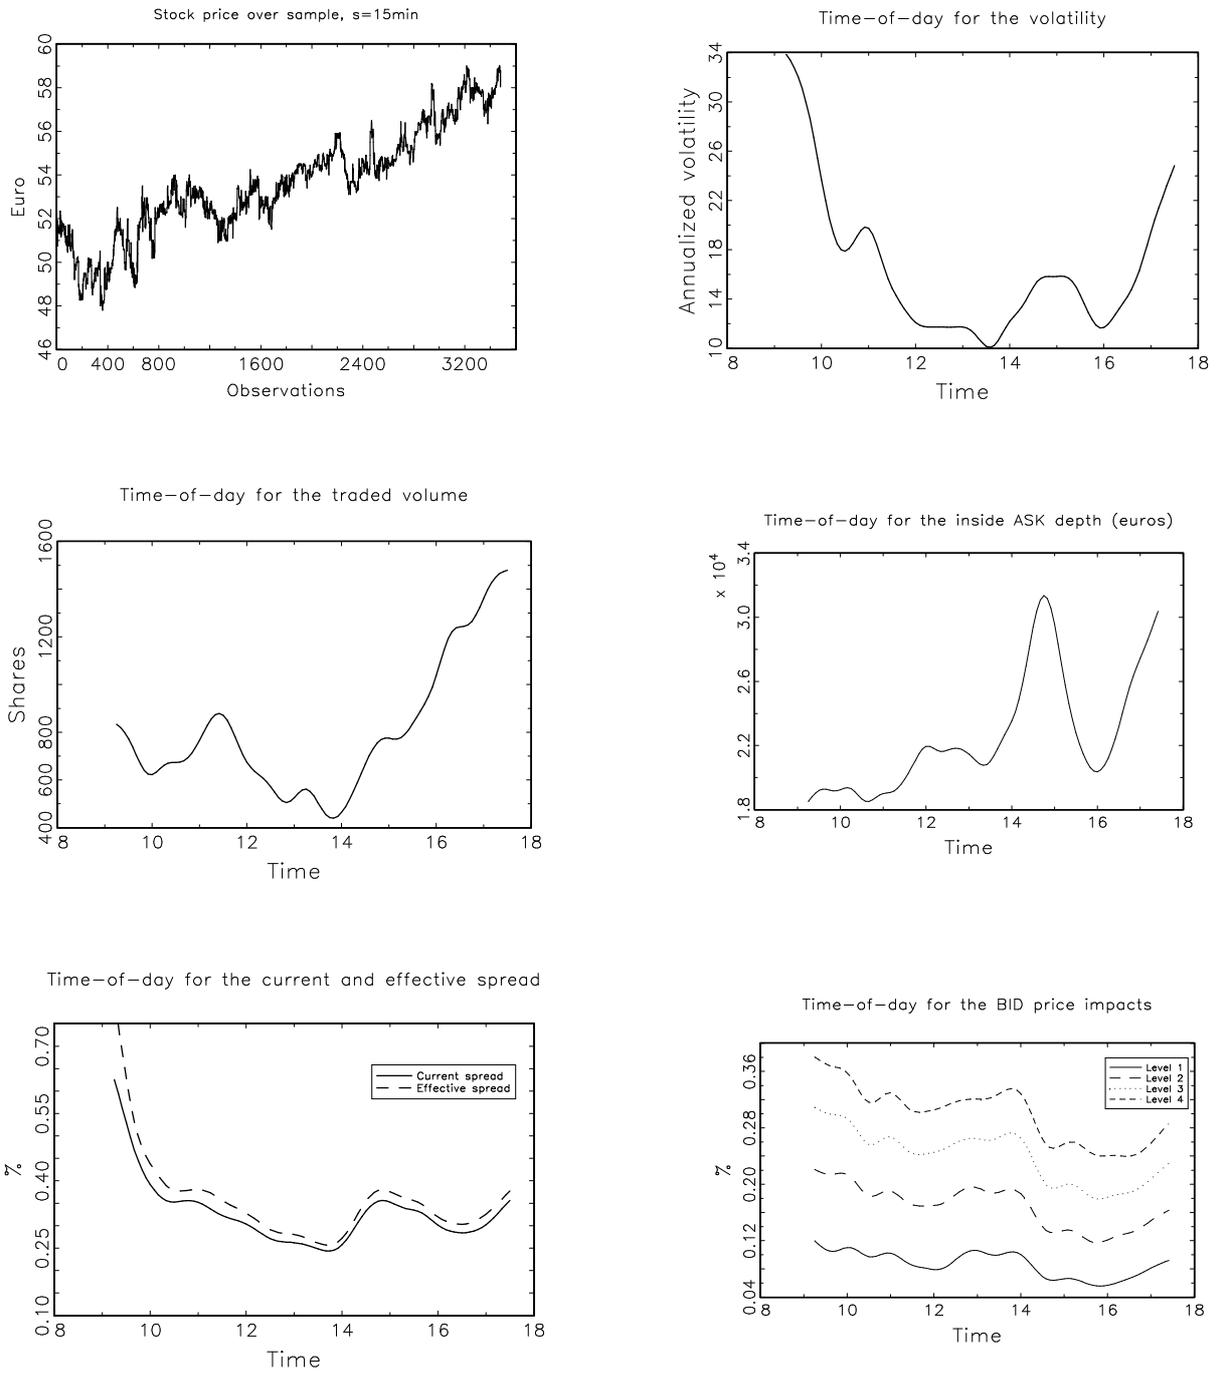
Figure 1. Colruyt. From top left to bottom right: stock price (sampled at 15-minute intervals), time-of-day for the volatility, time-of-day for the traded volume, time-of-day for the ask inside depth, time-of-day for the current and effective spread and time-of-day for the bid price impacts. The time period is December 2, 2002 to April 30, 2003.
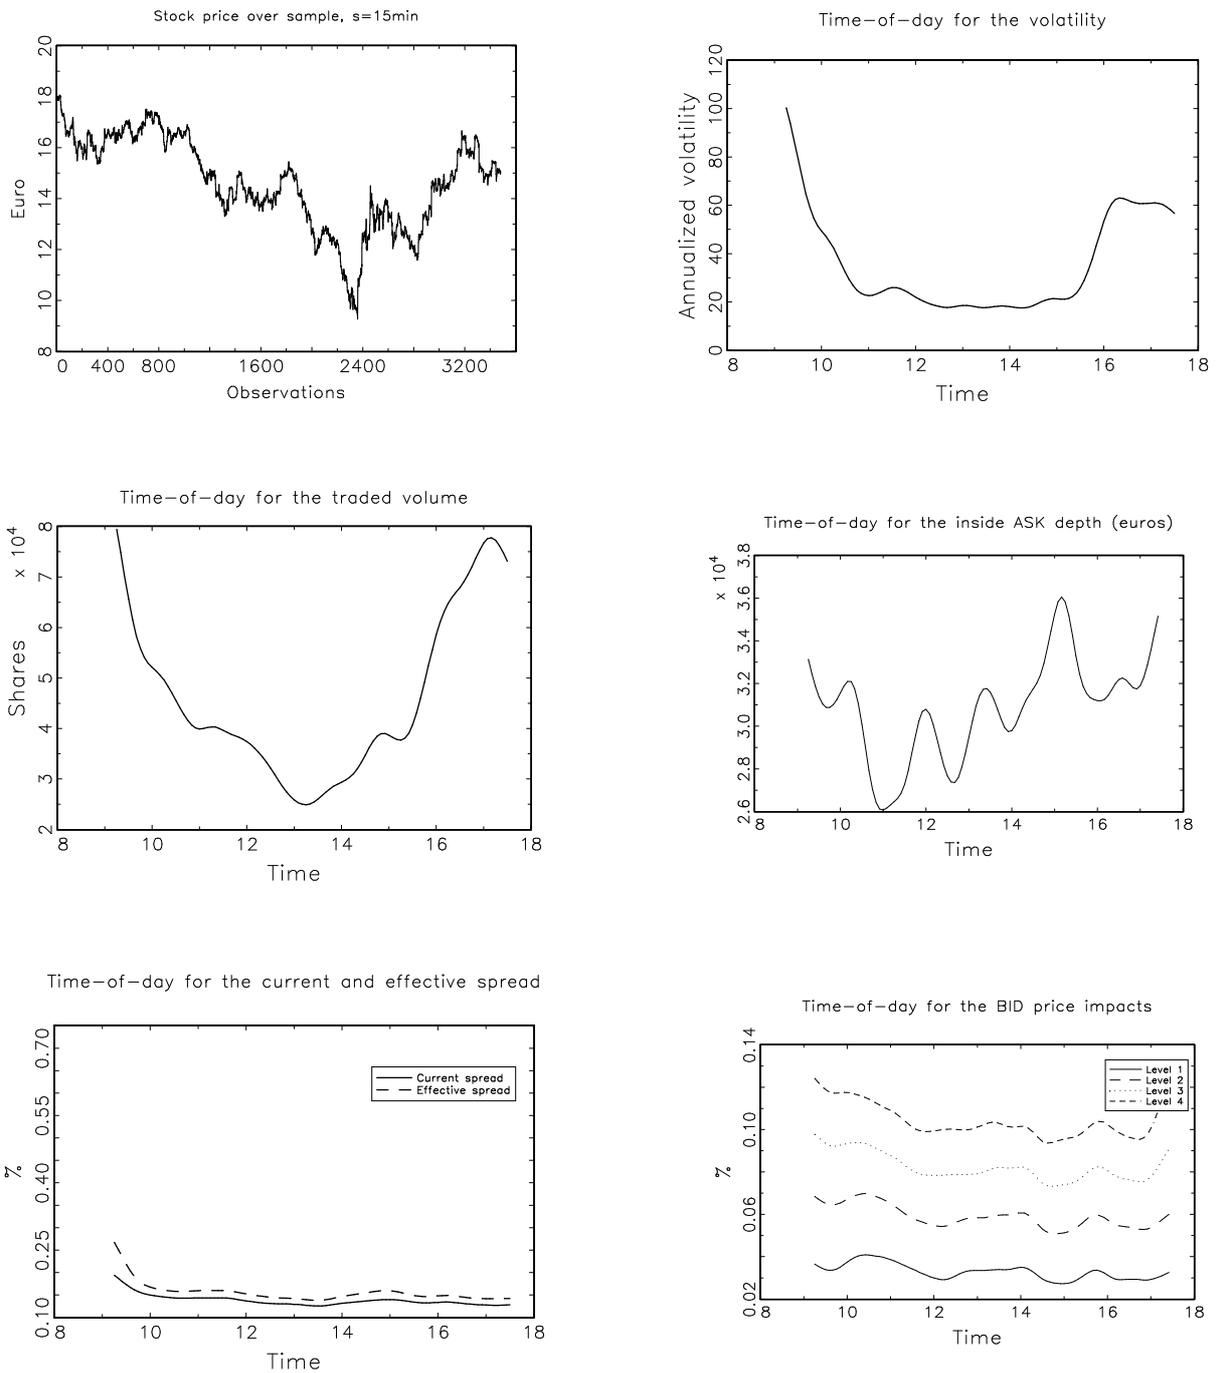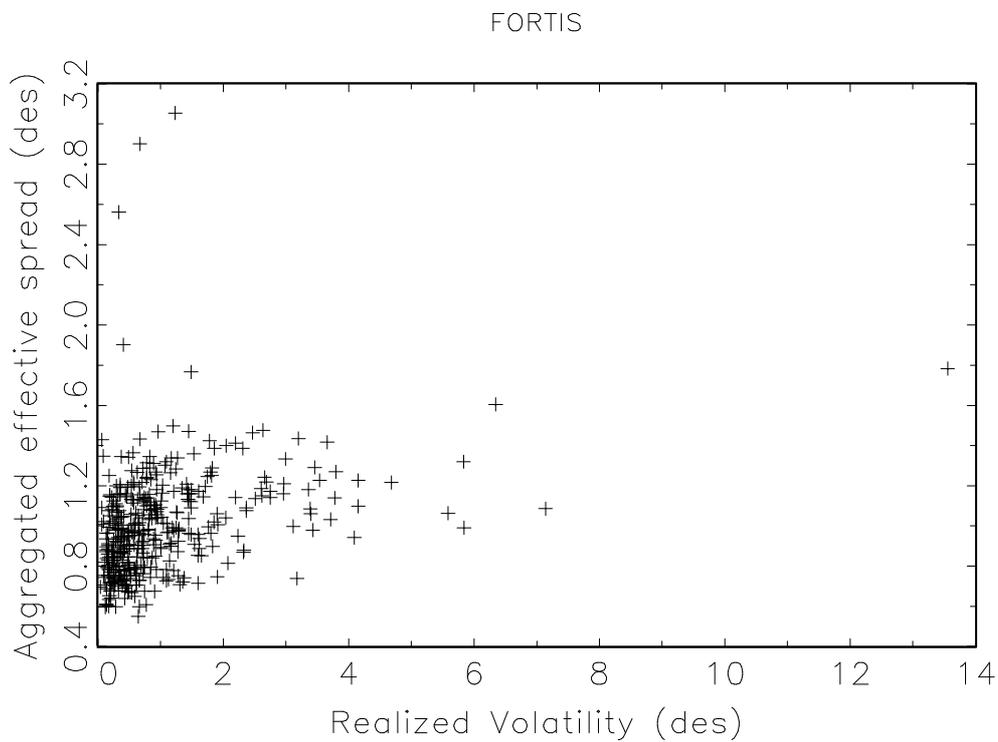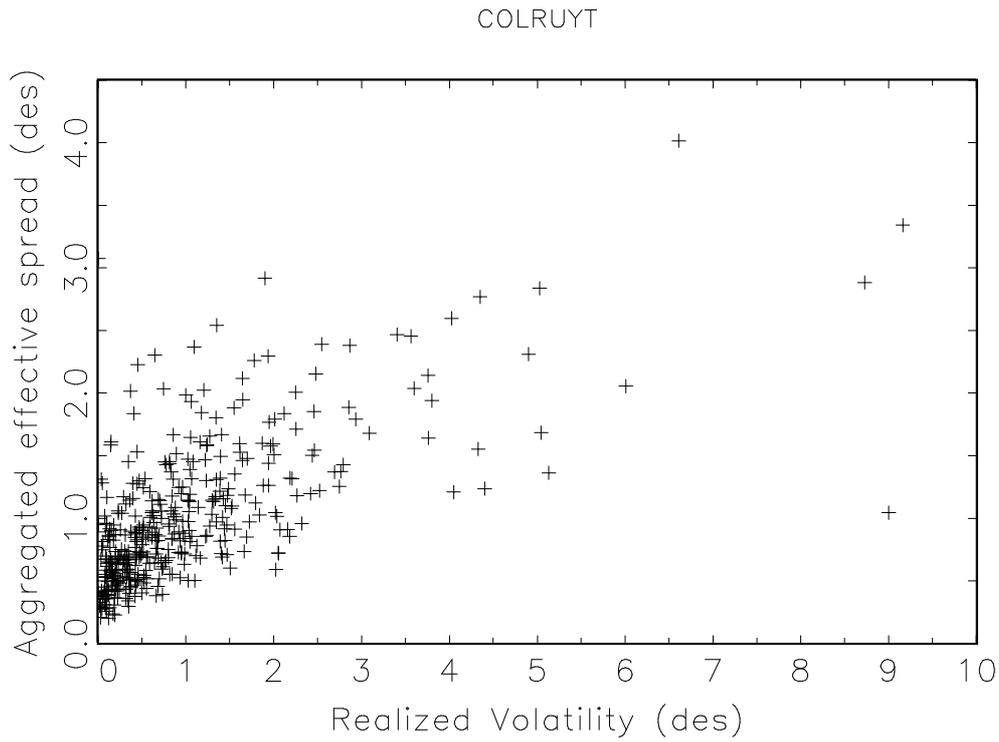
Figure 2. Fortis. From top left to bottom right: stock price (sampled at 15-minute intervals), time-of-day for the volatility, time-of-day for the traded volume, time-of-day for the ask inside depth, time-of-day for the current and effective spread and time-of-day for the bid price impacts. The time period is December 2, 2002 to April 30, 2003.

Figure 3. Aggregated effective spread vs realized volatility (Colruyt and Fortis). Relationship between the aggregated effective spread and the realized volatility. The aggregated effective spread is the average effective spread over the [9h:11h], [11h:13h], [13h:15h] and [15h:17h30] intervals; the realized volatility is defined over the same intervals. Both measures are deseasonalized by their respective time-of-day. The time period is December 2, 2002 to April 30, 2003.
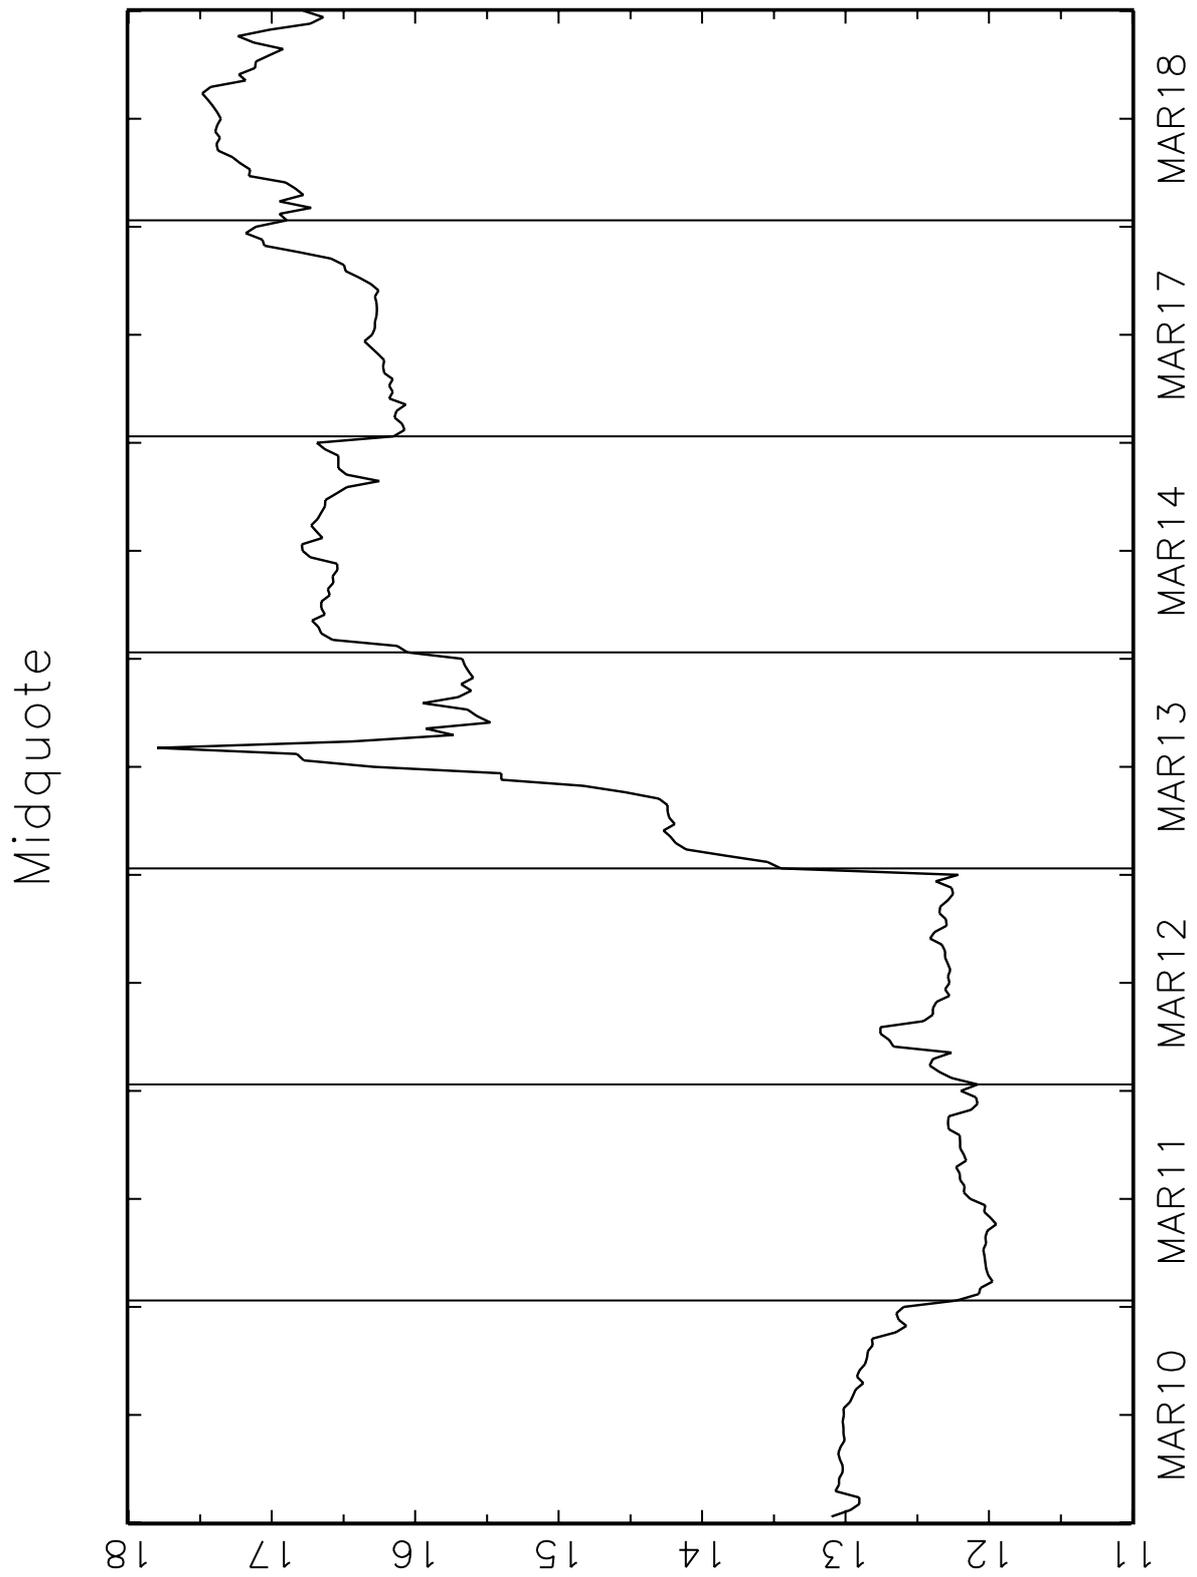
37

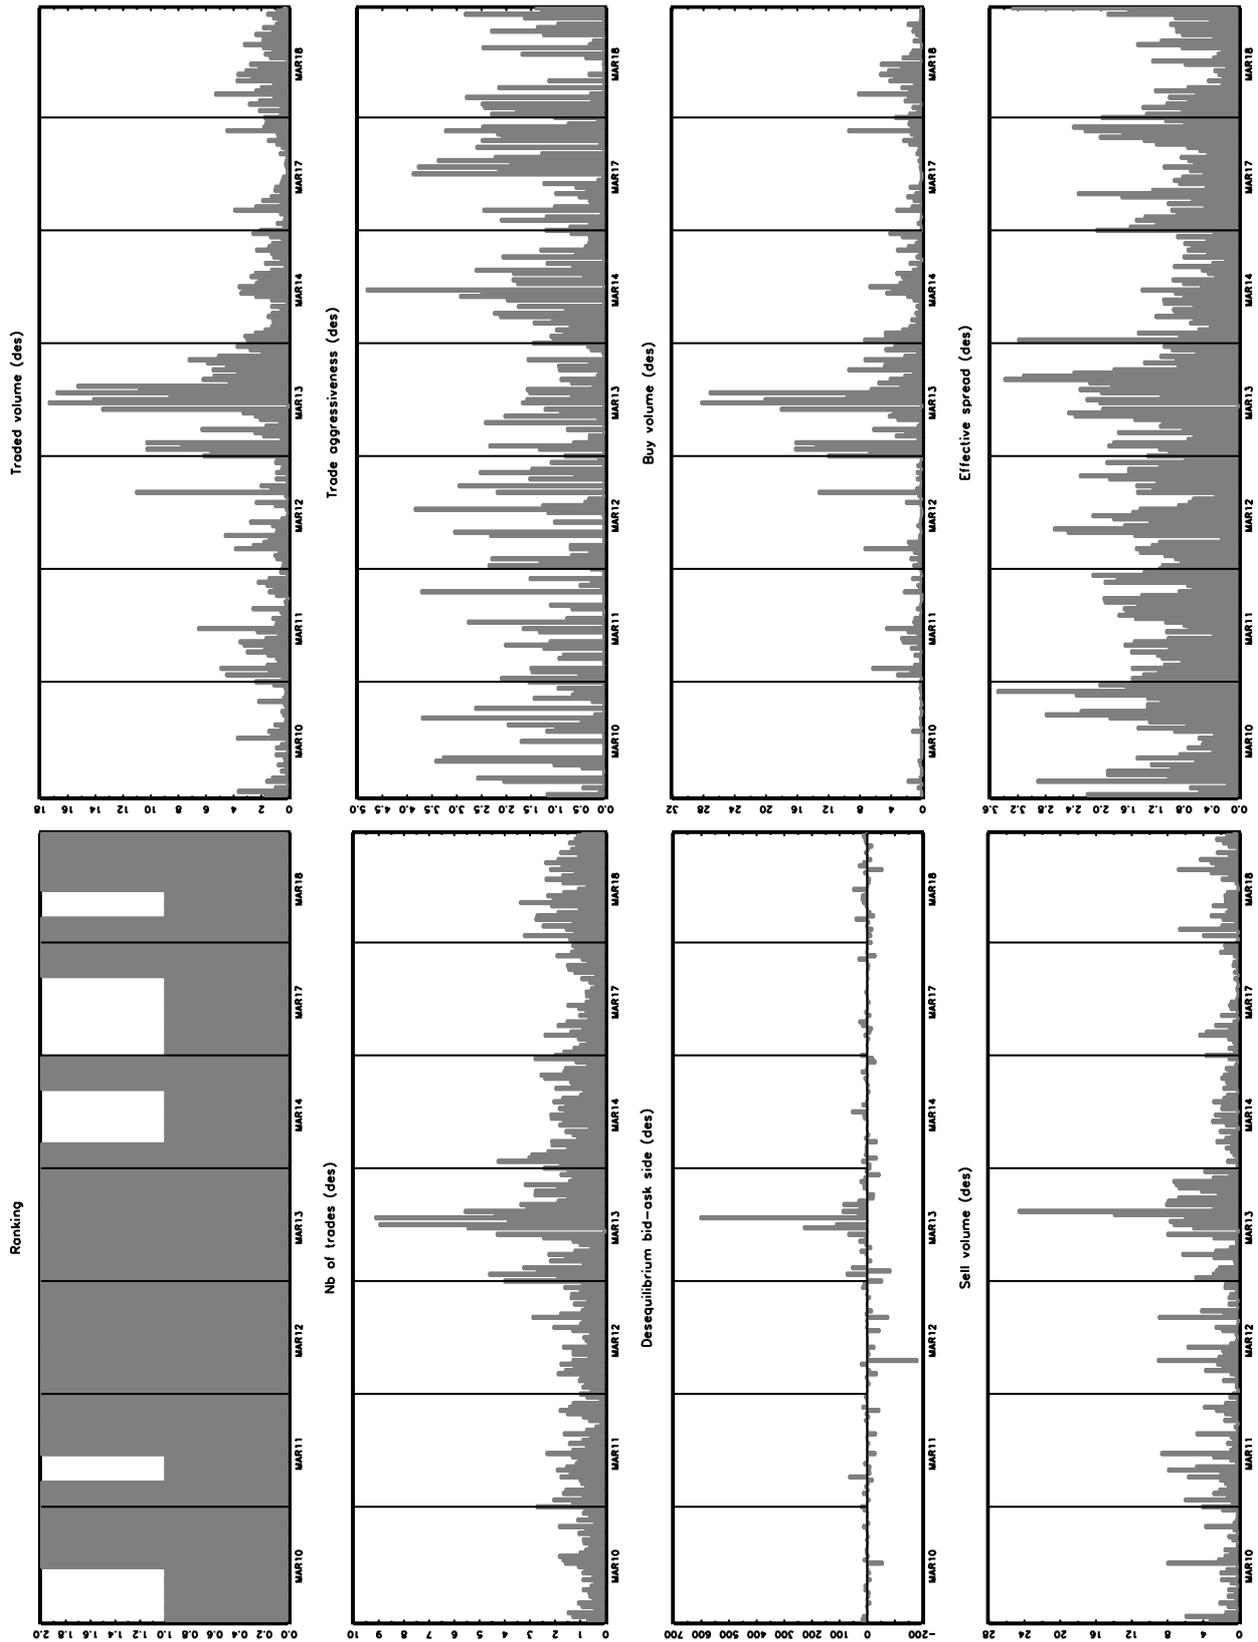**Figure 4. Delhaize (event study).** Price pattern for the five days centered on March 13, 2003.

Figure 5. Delhaize (event study). Trade dynamics and market liquidity (all variables are deseason-alized) for the five days centered on March 13, 2003.
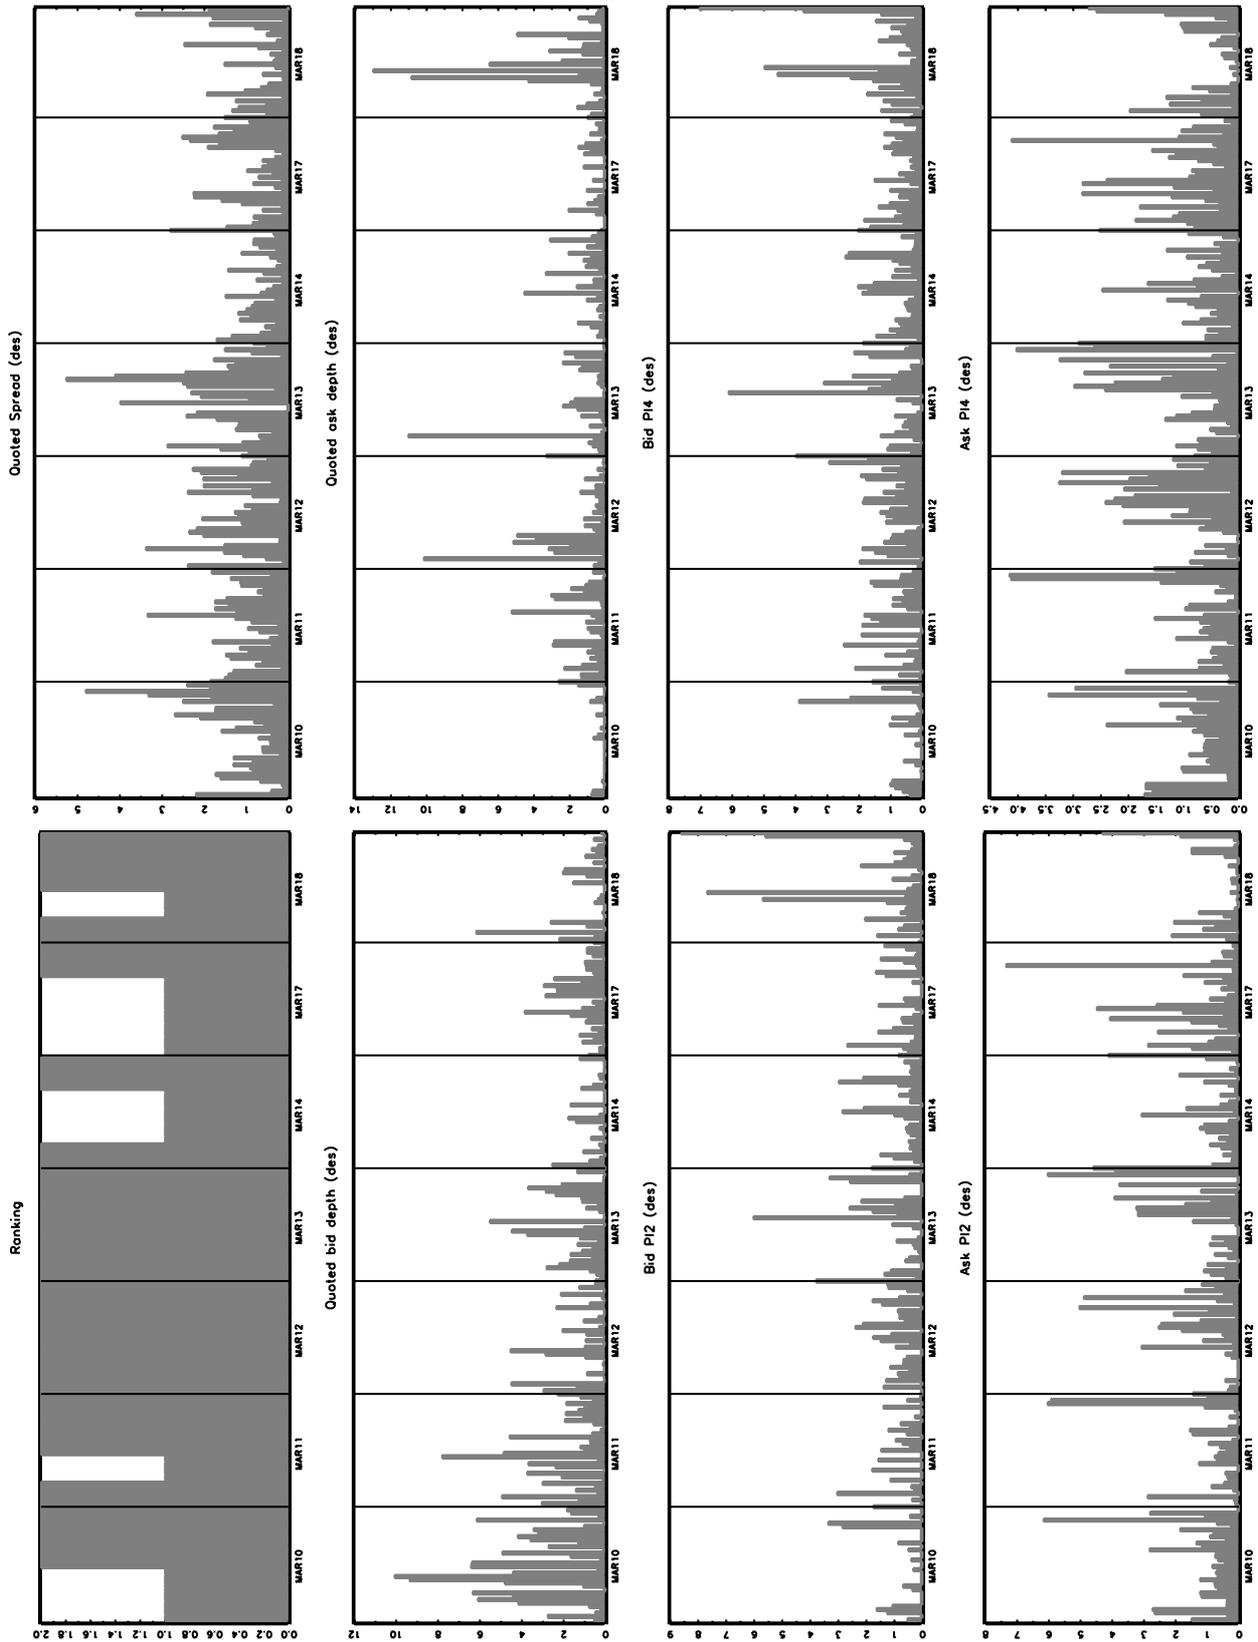
Figure 6. Delhaize (event study). Book dynamics and market liquidity (all variables are deseason-alized) for the five days centered on March 13, 2003.