

EXPECTED AND UNEXPECTED COST OF TRADING IN THE XETRA AUTOMATED AUCTION MARKET

Helena Beltran Lopez¹, Pierre Giot^{1,2,4}, and Joachim Grammig^{1,3}

February 2002

Abstract

In this paper, we propose measures for characterizing the expected and unexpected cost of trading that can be applied to analyze automated electronic auction markets. Using a unique database which contains the full state of the order book for three stocks (Daimler-Chrysler, Deutsche Telekom and SAP) traded on the XETRA platform, we show that the unexpected cost of trading, but not the expected cost of trading, can depend on international linkages of stock markets. Both measures also depend on structural characteristics of the exchange (time of day patterns). More precisely, the expected cost of trading increases strongly with the traded volume and is much higher at the start of the day. The unexpected cost of trading increases moderately with the traded volume, is much higher at the start of the day and increases sharply when pre-trading starts in the US markets.

Keywords: XETRA, automated auction market, market liquidity, trading risk, intraday data

JEL classification:

¹CORE, Université Catholique de Louvain, Belgium.

²Department of Business Administration, University of Namur, Belgium.

³University of St-Gallen, Switzerland.

⁴Corresponding author: Pierre Giot, Department of Business Administration, University of Namur, Rempart de la Vierge 8, B-5000 Namur, Belgium. Email: pierre.giot@fundp.ac.be.

1 Introduction

The theoretical and empirical literature on continuous trading mechanisms traditionally classifies these as being either market maker driven or based on an order book system.¹ Market maker systems are used in the Foreign Exchange (FOREX) market, at the NASDAQ (although some order book features have been recently introduced) and in most over-the-counter trading systems. As to the order book systems, they are used (not an exhaustive list) at the Toronto Stock Exchange, the Paris Bourse and EuroNext, Lisbon and Madrid Stock Exchanges and at the XETRA trading system in Germany. Hybrid systems, which combine features from both market maker and order book mechanisms, are used for example at the New York Stock Exchange. The recent availability of order book data from the Paris Bourse and Madrid Stock Exchange has spurred empirical studies focusing on topics such as: the pattern of the spread during the trading day (Gouriéroux, Le Fol, and Meyer, 1998 and Gouriéroux, Le Fol, and Jasiak 1999), the price-volume relationship and the role of market depth (Gouriéroux, Le Fol, and Meyer, 1998 and Martinez, Tapia, and Rubio, 2000), the design of ACD or Log-ACD based duration models (Bisière and Kamionka, 2000 and Bauwens and Giot, 2000), the placement of limit orders and the competition for liquidity (Biais, Hillion, and Spatt, 1995), . . .

In this paper, we propose measures for characterizing the expected and unexpected cost of trading that can be applied to analyze automated electronic auction markets. In our framework, the expected cost of trading at time $t - 1$ for a subsequent (short) time interval is defined as the conditional expectation of the subsequent ask-bid return, or $E(r_t(v)|I_{t-1})$. Correspondingly, the unexpected cost of trading at time $t - 1$ is defined as the conditional variance of the return, or $V(r_t(v)|I_{t-1})$. This unexpected component can be viewed as the ‘trading risk’ of the return. Using a unique database which contains the full state of the order book for three stocks (Daimler-Chrysler, Deutsche Telekom and SAP) traded on the XETRA platform, we document that the unexpected cost of trading, but not the expected cost of trading, can depend on international linkages of stock markets. Both measures also depend on structural characteristics of the exchange (time of day patterns). Furthermore, we argue that time of day patterns of both expected and unexpected cost of trading for large traded volumes suggest the introduction of an ‘upstairs market’ for the German stock exchange’s electronic trading environment XETRA.

Taking into account previous empirical work for stocks traded on the Paris Bourse or on the Madrid Stock Exchange, we deal more specifically with two research agendas. First and in a more ‘descriptive’ approach, we investigate the price-volume relationship over the trading day. Indeed and going back to Kyle (1985) for the theoretical framework, it is well known that the price agreed upon for a trade very much depends on the traded volume. More specifically, the larger the traded volume, the larger the deviation of the price of the trade from an equilibrium price prior to the trade. This is the ‘depth’ part of market liquidity, which complements the inside spread as the ‘tightness’ measure of liquidity. Because our dataset displays the continuous recording of the full state of the order book, such a study can easily be undertaken.

In the second part of the paper, we put forward an econometric model for the returns (defined

¹See for example Domowitz (1993), O’Hara (1995), Madhavan (2000) or Bauwens and Giot (2001).

on equidistantly sampled prices), study the pattern of their mean level and conditional volatility, and characterize measures for computing the expected and unexpected cost of trading as defined above. More specifically we focus on ask-bid returns (defined on subsequent ask and bid prices) which are relevant for a trader performing a succession of buy-sell transactions on an intraday basis. One of the main feature of our analysis is that all returns are defined on ask and bid prices valid for the immediate trade of a given volume v . We thus go beyond the usual econometric analysis of intraday returns (see for example Andersen and Bollerslev, 1997, Giot, 2000 or Dacorogna, Gencay, Muller, Olsen, and Pictet, 2001) by a ‘volume dependent analysis’, where all returns depend on a given volume to be traded. We derive implications for market liquidity by characterizing the expected cost of trading and the unexpected cost of trading (i.e. variance and conditional variance) from our econometric model, both measures being dependent on the traded volume and the time-of-day.

Our key results show that: (a) the slope of the ask and bid curves and the inside spread are much larger at the start of the trading day than at any other times, liquidity is thus quite low at the start of the trading day; (b) there is a slight drop in liquidity around 15h Frankfurt time, i.e. in the pre-trading period of the U.S. markets; (c) the inside spread is low and the depth is high at the end of the trading day, suggesting that liquidity does not dry out at the end of the day; (d) an AR(p)-GARCH(1,1) model with time-of-day seasonality for the level and volatility of the returns succeeds in modelling adequately the dynamics of the ask-bid returns; (e) the volatility of the returns is large at the start of the day and around 15h, the latter suggesting a high degree of international linkage between the US and European markets regarding the information flow and uncertainty of the open of trading in New York; (f) the time-of-day in the spread gives us a measure of the expected cost of trading; (g) the time-of-day in volatility combined with the volatility forecast of the GARCH model gives us a measure of the unexpected cost of trading, i.e. the variance of the return associated with a trade of v shares; (h) both measures given in (f) and (g) supplement the ‘direct’ liquidity information given in (a)-(c) and provide important complimentary information regarding the state of the market.

The rest of the paper is organized in the following way. In Section 2, we characterize the XETRA trading system and we present our datasets in Section 3. Section 4 gives the first part our empirical analysis, while Section 5 deals with the econometric models and volatility analysis. Market liquidity is analyzed in Sections 4 and 5.

2 XETRA trading system

Under normal circumstances, trading in the XETRA system is continuous during the opening hours and is based on the order book trading mechanism (also called continuous double auction). As in Paris (with the CAC trading system) or Toronto (with the CATS trading system), a computerized order book keeps track of all incoming market and limit orders. The matching of orders is automatically performed by the computer based on the usual rules of price and time priority ². Thus no market makers are involved in

²See Biais, Hillion, and Spatt (1995) or Bauwens and Giot (2001) for a complete description of an order book market.

the trading process.

Prior to the start of trading at 8h30 (or 9h in the second part of our dataset), an opening auction is held where the opening price is determined as the price which maximizes the volume that can be traded.³ While actual trading is not allowed before the opening hour, traders and market participants can enter limit orders in the order book to participate in the opening auction. Depending on prevailing market conditions, trading starts at 8h30 (or 9h) or slightly later if there are large order imbalances.

3 Constructing the order book and the datasets

We consider three stocks traded on the XETRA trading system of the Deutsche Borse. These stocks are Daimler-Chrysler (ticker DCX), SAP (SAP), and Deutsche Telekom (DTE); the period under study ranges from August, 2nd 1999 to October 29th, 1999. During the first part of this period (until September 19th), trading on the XETRA system was allowed from 8h30 to 17h; starting September 20th, trading is allowed from 9h to 17h30 as the opening and closing times were delayed by 30 minutes to match the opening and closing times of the other stock exchanges in Europe.

Our dataset contains *all* events (limit, market, fill or kill, cancellation, ...) entered in the order book (including auctions) during the three months and all matching outcomes (full or partial execution, ...). With this available information, we reconstruct the complete state of the order book in real time, i.e. we have at our disposal the full grid of bid and ask prices (with the corresponding volume) at any given time. In other words, the complete (aggregated) order book is available over the three month period, which corresponds to the information available to the market participant wishing to trade at a given moment. It should be stressed that the orders entered in the auction periods do not really matter for our analysis as we focus on the continuous trading periods and corresponding market liquidity; only the outcome of the auctions is taken into account as it yields the start-of-the-day state of the order book.

4 Descriptive analysis

4.1 Price-volume relationship and spread

In this subsection we deal with the price-volume relationship and the pattern of the spread during the trading day for the three stocks in our dataset. The price-volume relationship gives, at a specified time, the average price corresponding to an immediate trade of a specified volume.⁴ Indeed, it is a main feature of an order book market that the average bid price for a given (immediate) volume to be transacted

³For a description of the opening auction mechanism used in an order book market and corresponding trading strategies, see Biais, Hillion, and Spatt (1999).

⁴Thus, by immediate trade, one means a market order, i.e. an order to be filled at once, with no price limit. It should be emphasized that market orders are treated differently at XETRA than at the Paris stock exchange for example. Indeed, at the Paris stock exchange, market orders are only matched against the best bid and ask limit orders; the remaining unfilled part of the order (if any) is transformed into a limit order: there is thus no immediate walking up or down of the order book as at XETRA.

is a decreasing function of the submitted volume. Regarding the average ask price, it is an increasing function of the volume to be transacted.⁵ While the sign of the price-volume relationship is known beforehand, it is well known that the slope of this relationship is far from being constant and usually exhibits a time-of-day and market conditions dependence: in a tight market where few limit orders (or limit orders with a low volume) have been submitted to the trading system, the slope will be large as the requested (immediate) volume consumes the available liquidity quickly and thus leads to large changes in price. See for example Gouriéroux, Le Fol, and Meyer (1998) for an empirical investigation of such features at the Paris stock exchange. The slope of the price-volume relationship is thus a main indicator of the available liquidity at a given time in an order book system. It is also strongly related to the depth component of liquidity (i.e. amount of one sided volume that can be absorbed by the market without causing a revision of bid-ask prices) as given in Kyle (1985).

Regarding the spread, one defines the inside spread⁶ (i.e. difference between the best ask and bid prices) and the spread for a given volume (defined as the difference between the average ask and bid prices for the immediate trade of the specified volume). Due to the price-volume relationship, it is immediately seen that the spread is an increasing function of the traded volume. Furthermore the exact shape of this dependence is bound to depend on time-of-day and market conditions.

4.2 Empirical results

For our empirical analysis, we first define (by ‘relative’ we mean as percentage changes from the best bid and ask prices or spread):

- $a_{r,t}(v)$: average relative ask price at time t for the immediate trade (buy order) of a v shares order;
- $b_{r,t}(v)$: average relative bid price at time t for the immediate trade (sell order) of a v shares order;
- $s_{r,t}(0)$: inside relative spread at time t ;
- $s_{r,t}(v)$: relative spread at time t for a v shares immediate buy/sell trade.

Regarding the time dependence t , we compute the variables averaged over all available days. For example, $a_{r,t}(v)$ is to be understood as the average relative ask price at time t for the immediate trade (buy order) of a v shares order, averaged over all available days. Thus $a_{r,t}(v) = (a_{r,t,day1}(v) + \dots + a_{r,t,dayn}(v))/n$. This remark also applies to the other variables.

For the DCX, DTE and SAP stocks and for the first period of the sample, Figures 1, 2 and 3 plot the price-volume relationship for the average relative bid (bottom of figure) and ask (top of figure) prices as a function of time. Figure 4 gives the relative inside spread as a function of time for the three stocks.⁷ These figures indicate that:

⁵See for example Biais, Foucault, and Hillion (1997) or Bauwens and Giot (2001).

⁶In the framework of Kyle (1985), this corresponds to the first measure of liquidity which is the tightness.

⁷Because the general pattern of all these graphs is exactly the same in the second period of our sample than in the first period, we do not report the corresponding figures for the second period. They are however available on request.

- as expected, the bid (ask) curve is a decreasing (increasing) function of the immediate volume to be traded; correspondingly the spread is an increasing function of the immediate volume to be traded;
- a strong opening hour effect is at work, highlighting the time-of-day dependence of the variables. More specifically, the slope of $a_{r,t}(v)$ and $b_{r,t}(v)$ is much higher at the start of the trading day than at any other moments: liquidity is quite low at the start of the trading day. The same comment can be made for the relative inside spread, which is much higher at the start of the trading day than at any other times.
- liquidity appears to slightly dry up around 15h, i.e. prior to the open of the U.S. markets:
 - a) the slope of $a_{r,t}(v)$ and $b_{r,t}(v)$ is in some instances slightly larger at this time of the day than at other times, except of course close to the opening hour (not a pronounced effect however);
 - b) the relative inside spread increases a bit from 14h onwards, but decreases anew at the end of the day.
- interestingly the spread does not rise sharply at the end of the day, as it does at the NYSE for example (see Chung, Van Ness, and Van Ness, 1999 for example). Indeed, liquidity appears to be high at the end of the day as the slope of the price-volume relationship (both for bid and ask prices) is quite low at these times and the relative inside spread is at its lowest for the trading day.

To summarize, the empirical results presented in this section thus indicate that: (a) liquidity at XETRA features a strong time-of-day component, with both the inside spread and slope of the bid-ask curves much higher at the start of the trading day; (b) the inside spread is a limited indicator of the available liquidity in the order book, with the slope of the bid-ask curves providing important information regarding the price-volume relationship; (c) there is a slight drop in liquidity before the open of trading at the NYSE (i.e. around 15h Frankfurt time). Results (a) and (b) are not surprising as they are key features of a trading system based on an order book mechanism. See similar studies on Paris Bourse data by Gouriéroux, Le Fol, and Meyer (1998) and Gouriéroux, Le Fol, and Jasiak (1999), and on data for the Madrid Stock Exchange by Martinez, Tapia, and Rubio (2000).

5 Expected and unexpected cost of trading: an econometric analysis

In the second part of the paper, we perform an univariate econometric analysis of the intraday returns (defined on the equidistantly sampled prices) for the three stocks in our dataset which leads us to define the expected and unexpected cost of trading (see below). As mentioned in the introduction, one of the main feature of our analysis is that all returns are defined on ask and bid prices valid for the immediate trade of a given volume v . This type of analysis is made possible by our unique dataset where we can retrieve all ask and bid prices at a given time and valid for the immediate trade of a v share volume.

In this framework, we define ask-bid returns as returns defined as the log difference of consecutive ask and bid prices for the immediate trade of a volume v : $r_t(v) = \ln(b_t(v)) - \ln(a_{t-1}(v))$.⁸ Such ask-bid returns are relevant for a trader performing a succession of buy-sell transactions for a v share volume on an intraday basis.

5.1 An univariate model for the ask-bid returns

To keep notations short, we describe the econometric model using r_t instead of $r_t(v)$. The dynamics of the level of the returns is characterized by an AR(p) process. However, it is necessary to introduce a time-of-day function for the level of these returns as they include (by definition) the spread, which usually features a time-of-day pattern. See for example Chung, Van Ness, and Van Ness (1999) for NYSE data or Gouriéroux, Le Fol, and Jasiak (1999) for Paris Bourse data. Moreover, this time-of-day pattern should be strongly dependent on the volume to be transacted as large volumes should be associated with pronounced intraday patterns. The full specification for the dynamics of the ask-bid returns is thus:

$$r_t = \psi(t) + r'_t \quad (1)$$

where $\psi(t)$ is the time-of-day function for the return (as for $\phi(t)$ above, $\psi(t)$ depends on the traded volume, i.e. we have $\psi(t, v)$) and r'_t is modelled as

$$r'_t = \delta_0 + \delta_1 r'_{t-1} + \dots + \delta_p r'_{t-p} + u_t \quad (2)$$

i.e. an AR(p) specification for the deseasonalized ask-bid returns. As pointed out in several papers (see for example Andersen and Bollerslev, 1997 or Giot, 2000), intraday returns usually feature conditional heteroskedasticity combined with a strong time-of-day effect for the variance. Hence, we model the error term as:

$$u_t = e_t \sqrt{\phi(t)} \quad (3)$$

where $\phi(t)$ is the time-of-day function for the volatility⁹,

$$e_t = \sqrt{h_t} \epsilon_t \quad (4)$$

with ϵ_t IID $t(0, 1, \nu)$, and

$$h_t = \omega + \alpha e_{t-1}^2 + \beta h_{t-1} \quad (5)$$

which is the GARCH(1,1) specification of Bollerslev (1986) combined with a Student distribution for the error term. Another possibility is:

⁸Note that the time interval between times t and $t - 1$ is equal to the length of the sampling interval.

⁹Note that in our framework there exists a different time-of-day function for each possible traded volume, we thus have a dependence of $\phi(t)$ on v , i.e. $\phi(t, v)$.

$$h_t = \omega + \alpha e_{t-1}^2 + \gamma e_{t-1}^2 I_{e_{t-1} < 0} + \beta h_{t-1} \quad (6)$$

which is the GJR-GARCH specification of Glosten, Jagannathan, and Runkle (1993) ($I_{e_{t-1} < 0}$ is an indicator function which is positive if the return at time $t - 1$ is negative). Thus a positive γ indicates that there is an asymmetrical relationship between the conditional variance and the lagged squared error term: past negative shocks have a deeper impact on current conditional volatility than past positive shocks (see Black, 1976; French, Schwert, and Stambaugh, 1987 or Pagan and Schwert, 1990).

5.2 Empirical results

To estimate the univariate model presented above, we pool the data from both periods so that we work with the same econometrical specification across the whole sample. However we allow for a period-specific time-of-day pattern for the level and volatility of the returns. Indeed, as the opening and closing times are different in both periods, a common time-of-day pattern for the level or volatility of the returns would not be meaningful. When the AR(p) specification is estimated, the dataset is slightly transformed to ensure that intraday returns from a given day are only regressed on past intraday returns from the same day.

Our estimation sample for the stocks DCX, DTE and SAP ranges from August, 2nd 1999 to October 29th, 1999 and we consider six possible volumes: 1; 5,000; 10,000; 20,000; 30,000 and 40,000 shares. Bid and ask prices for these six volumes are sampled each s minute, with s equal to 5 and 10 minutes. All models are estimated using the CML library in GAUSS. Because the quantity of available results is quite large, we only report key results and graphs for some selected volumes, times and stocks. Unless indicated otherwise, it is to be understood that the given results are valid for the three stocks, all times and all volumes.

5.2.1 Mean specification

A strong time-of-day component for the returns is present (see Figure 5 for the three stocks in period 1 and 5 minute returns) and this effect is much more pronounced for large volumes. Of course, this time-of-day component closely tracks the spread whose level features a strong time-of-day effect and whose magnitude increases with the traded volume. The spread is highest at the start of the trading day, increases around 14h30 prior to the open of the US markets and increases with the volume to be traded. Correspondingly (see the full results in Table 1), the unconditional mean increases steadily as the traded volume increases: on average it is 5.5 times more costly to trade 40,000 shares than 1 share (and 3.6 times more costly to trade 40,000 shares than 5,000 shares) for the DCX stock, 5-minute returns. Ask-bid returns for the DTE stock have more or less the same characteristics although the unconditional mean at $v = 1$ share is almost 50% larger than for the DCX stock: the inside spread is thus larger for DTE than for DCX. While the unconditional mean at $v = 1$ share for the SAP returns is slightly larger than the one for DTE, it increases strongly with the traded volume: at $v = 20,000$ shares, the unconditional mean for the SAP returns is more than 20 times larger than the unconditional mean for the DCX or

Table 1: Ask-bid returns: unconditional mean vs volume

	5-minute returns			10-minute returns		
	DCX	DTE	SAP	DCX	DTE	SAP
$v = 1$ share	-0.1065	-0.1623	-0.1791	-0.1064	-0.1594	-0.1741
$v = 5,000$ shares	-0.1634	-0.2141	-0.7510	-0.1621	-0.2116	-0.7473
$v = 10,000$ shares	-0.2224	-0.2724	-1.7206	-0.2212	-0.2702	-1.7139
$v = 20,000$ shares	-0.3415	-0.3880	-7.5033	-0.3399	-0.3861	-7.4765
$v = 30,000$ shares	-0.4637	-0.5051	-	-0.4620	-0.5034	-
$v = 40,000$ shares	-0.5903	-0.6253	-	-0.5888	-0.6238	-

Unconditional mean (in %) for the ask-bid returns as a function of volume and sampling time (pooled results from both periods).

DTE stocks. Thus the inside spread for DTE and SAP (and even DCX) are not very different but the market depth for DTE and DCX is much larger than for SAP.¹⁰

Regarding the dynamics of the returns once the time-of-day component has been removed, an AR(3) specification is needed to take into account their autocorrelation prior to the GARCH modelling.

5.2.2 Volatility specification

There is a strong time-of-day component for the volatility of all ask-bid returns, with the time-of-day effect being stronger for returns defined on large volumes (see Figure 6 for the three stocks in period 1 and 5 minute returns): the volatility risk is larger for large volume trades and is not constant during the day. Volatility is strong at the open and around 15h, i.e. as the pre-market trading activity starts in New York, and it picks up again around 16h. Thus there is a need for volatility deseasonalisation before estimation by GARCH models.¹¹ The GARCH(1,1) model performs adequately and there is no conditional asymmetry in the GJR-GARCH specification as γ is not significant in almost all cases. The unconditional variance (from the GARCH(1,1) model) increases steadily as the traded volume increases, but this effect is much less pronounced than for the unconditional mean; for example, the unconditional variance increases by only 10% as the traded volume goes from 1 share to 40,000 shares (DCX stock, 5 minute returns). However, one should add the effect from the time-of-day pattern in the volatility (which also increases with the traded volume), see next sub-section.

5.2.3 International linkage between the US and German markets

While the high level of volatility at the start of the trading day is not surprising (similar features are observed at the Paris Bourse, the NYSE or the NASDAQ for example), the increase in volatility around

¹⁰Again, these results indicate that the inside spread is a misleading indicator for the available liquidity, obviously in an order book trading mechanism and frequently in a market maker trading system.

¹¹As pointed out in Andersen and Bollerslev (1997), by not including the time-of-day function for the volatility in the GARCH(1,1) specification, one would get incorrect coefficients for the conditional variance model.

15h indicates the high degree of international linkage between the US markets and the European markets. Indeed, as the pre-trading starts in New York around 14h30-15h, the information flow and uncertainty regarding the opening of the US markets both increase, leading to a higher volatility regime in European trading.

5.2.4 Expected and unexpected cost of trading

As mentioned in the introduction of the paper, we consider two aspects of market liquidity for the ask-bid returns: the expected cost of trading, which is defined as $E(r_t(v)|I_{t-1})$ and the unexpected cost of trading, which is defined as $V(r_t(v)|I_{t-1})$. Both are parameterized by v , the amount of shares to be traded. In our framework, the expected cost of trading is directly related to the level of the spread at time t . Indeed, taking the conditional expectation of Equation (1) gives us:

$$E(r_t(v)|I_{t-1}) = \psi(t, v) + \delta_0 + \delta_1 r'_{t-1} + \delta_2 r'_{t-2} + \delta_3 r'_{t-3}. \quad (7)$$

However, as the conditional mean of the AR(3) is very small, one has that $E(r_t(v)|I_{t-1})$ is approximately equal to $\psi(t, v)$. This is plotted in Figure 5 for the three stocks. As indicated previously, this type of market liquidity depends on the time-of-day and decreases strongly with the volume to be traded.

Regarding the unexpected cost of trading, one has that:

$$V(r_t(v)|I_{t-1}) = \phi(t, v) \hat{h}_t. \quad (8)$$

Thus, the unexpected cost of trading is equal to the product of the conditional variance as forecasted by the GARCH model and the prevailing time-of-day function for the volatility at that time. Because there is no asymmetry in the conditional variance for the intraday returns (coefficient γ almost never significant in the GJR-GARCH specification), a plain GARCH(1,1) is adequate for modelling \hat{h}_t and thus the unexpected cost of trading can be re-written as:

$$V(r_t(v)|I_{t-1}) = \phi(t, v) \left(\hat{\omega} + \hat{\alpha} e_{t-1}^2 + \hat{\beta} \frac{V(r_{t-1}(v)|I_{t-2})}{\phi(t-1, v)} \right) \quad (9)$$

as $h_{t-1} = \frac{V(r_{t-1}(v)|I_{t-2})}{\phi(t-1, v)}$, with e_{t-1}^2 given by Equations (2) and (3). As an example, we plot $V(r_t(v)|I_{t-1})$ as a function of the first 240 observations in Figure 7 for the DCX and DTE stocks. Quite surprisingly, while the $V(r_t(v)|I_{t-1})$ are different when v changes, the differences are rather small, even when comparing the unexpected cost of trading for $v = 1$ and $v = 40,000$ shares.

Once a parametric model for the ask-bid returns is specified (a Student model in our case) and $V(r_t(v)|I_{t-1})$ has been estimated along with the degrees of freedom ν of the Student distribution, it is straightforward to compute a measure of the intraday Value-at-Risk for the returns at our very short time horizon. As shown in Giot (2002) who motivates the use of Value-at-Risk measures for very short time horizons, the VaR at time $t - 1$ for the intraday ask-bid returns up to time t (for a traded volume of v shares and a confidence level of 100α percent in our case) is equal to:

$$VaR_{t-1}(v) = t_{\alpha,\nu} \sqrt{V(r_t(v)|I_{t-1})} \quad (10)$$

where $t_{\alpha,\nu}$ is the corresponding quantile of the Student distribution with ν degrees of freedom and $V(r_t(v)|I_{t-1})$ is the forecast of the conditional variance by the Student GARCH(1,1) model. This additional risk measure takes into account the ‘fat-tail’ feature of the distribution of returns and thus complements the unexpected trading risk measure.

To summarize, our empirical analysis indicates that the expected cost of trading strongly increases when the traded volume increases, while the unexpected cost of trading increases only slightly. However the unexpected cost of trading displays a more pronounced time-of-day effect in the afternoon: while $E(r_t(v)|I_{t-1})$ increases slightly around 15h, there is a sharp increase in $\phi(t, v)$ around 15h and this leads to a sharp increase in the conditional variance at that time, whatever the forecasted variance by the GARCH(1,1) model.


6 Conclusion

References

- ANDERSEN, T., AND T. BOLLERSLEV (1997): “Intraday periodicity and volatility persistence in financial markets,” *Journal of Empirical Finance*, 4, 115–158.
- BAUWENS, L., AND P. GIOT (2000): “The Logarithmic ACD model: an application to the bid-ask quote process of three NYSE stocks,” *Annales d’Economie et Statistique*, 60, 117–149.
- (2001): *Econometric modelling of stock market intraday activity*. Kluwer Academic Publishers.
- BIAIS, B., T. FOUCAULT, AND P. HILLION (1997): *Microstructure des marchés financiers: institutions, modèles et tests empiriques*. Presses Universitaires de France.
- 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.
- (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.
- BLACK, F. (1976): “Studies of Stock Market Volatility Changes,” *Proceedings of the American Statistical Association, Business and Economic Statistics Section*, pp. 177–181.
- BOLLERSLEV, T. (1986): “Generalized autoregressive conditional heteroskedasticity,” *Journal of Econometrics*, 31, 307–327.

- CHUNG, K., B. VAN NESS, AND R. VAN NESS (1999): “Limit orders and the bid-ask spread,” *Journal of Financial Economics*, 53, 255–287.
- DACOROGNA, M., R. GENÇAY, U. MULLER, R. OLSEN, AND O. PICTET (2001): *An introduction to high-frequency finance*. Academic Press.
- DOMOWITZ, I. (1993): “A taxonomy of automated trading systems,” *Journal of International Money and Finance*, 12, 607–631.
- FRENCH, K., G. SCHWERT, AND R. STAMBAUGH (1987): “Expected Stock Returns and Volatility,” *Journal of Financial Economics*, 19, 3–29.
- GIOT, P. (2000): “Time transformations, intraday data and volatility models,” *Journal of Computational Finance*, 4, 31–62.
- (2002): “Market risk models for intraday data,” University of Namur Discussion Paper.
- GLOSTEN, L., R. JAGANNATHAN, AND D. RUNKLE (1993): “On the relation between expected return on stocks,” *Journal of Finance*, 48, 1779–1801.
- GOURIÉROUX, C., G. LE FOL, AND J. JASIAK (1999): “Intraday market activity,” *Journal of Financial Markets*, 2, 193–226.
- GOURIÉROUX, C., G. LE FOL, AND B. MEYER (1998): “Analyse du carnet d’ordres,” *Banque et Marchés*, 36, 5–20.
- KYLE, A. (1985): “Continuous auctions and insider trading,” *Econometrica*, 53, 1315–1335.
- MADHAVAN, A. (2000): “Market microstructure: a survey,” *Journal of Financial Markets*, 3, 205–258.
- MARTINEZ, M., M. TAPIA, AND G. RUBIO (2000): “Understanding liquidity: a closer look at the limit order book,” Working Paper 00-52. Department of Economics, Universidad Carlos III Madrid.
- O’HARA, M. (1995): *Market microstructure theory*. Basil Blackwell, Oxford.
- PAGAN, A., AND G. SCHWERT (1990): “Alternative Models for Conditional Stock Volatility,” *Journal of Econometrics*, 45, 267–290.

Average ask price vs volume and vs time



Average bid price vs volume and vs time


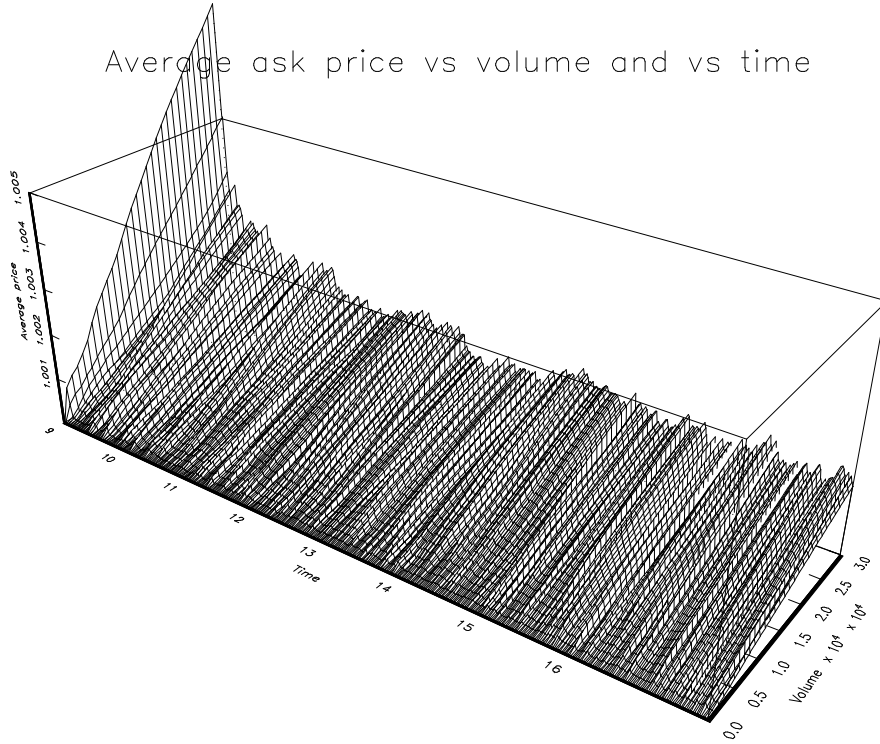


Figure 1: Average relative bid and ask price vs volume and vs time (Daimler-Chrysler, 2 August to 19 September 1999)

Average ask price vs volume and vs time



Average bid price vs volume and vs time





Figure 2: Average relative bid and ask price vs volume and vs time (Deutsche Telekom, 2 August to 19 September 1999)

Average ask price vs volume and vs time



Average bid price vs volume and vs time




Figure 3: Average relative bid and ask price vs volume and vs time (SAP, 2 August to 19 September 1999)

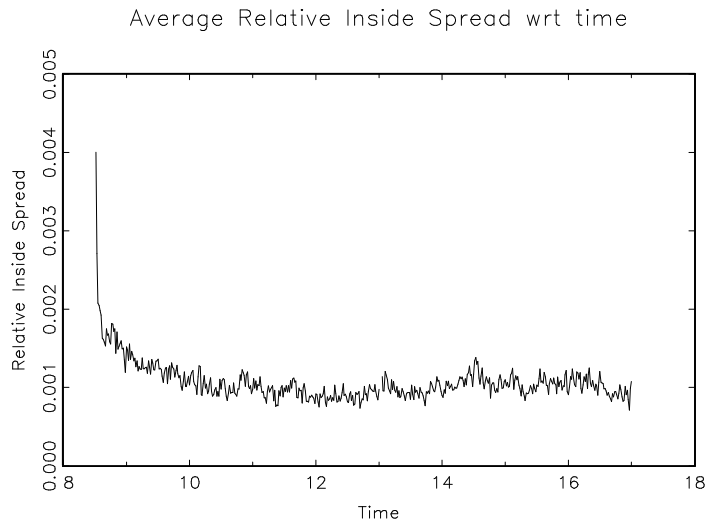

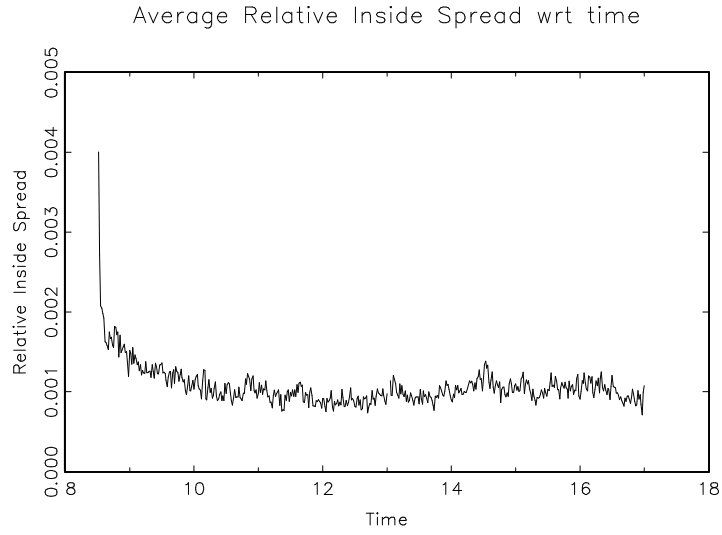


Figure 4: Average relative inside spread for Daimler-Chrysler (top), Deutsche Telekom (middle) and SAP (bottom), 2 August to 19 September 1999

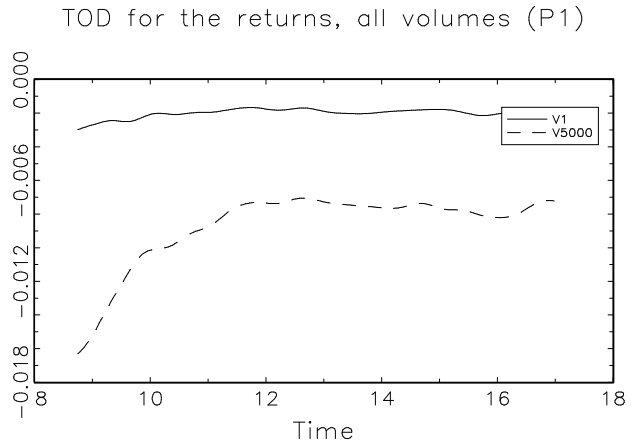

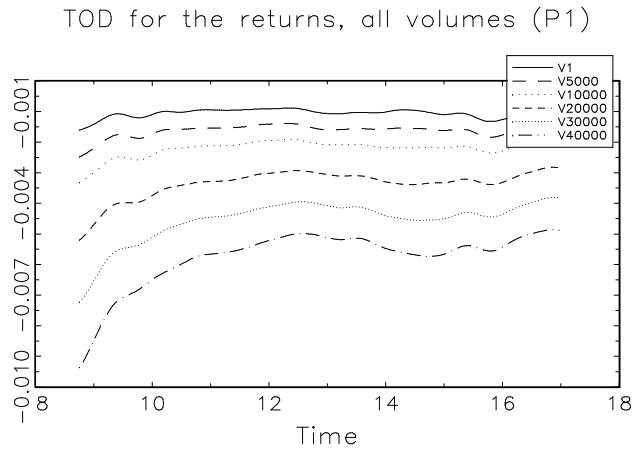


Figure 5: Time-of-day function for the mean of the ask-bid return, Daimler-Chrysler (top), Deutsche Telekom (middle) and SAP (bottom), 2 August to 19 September 1999

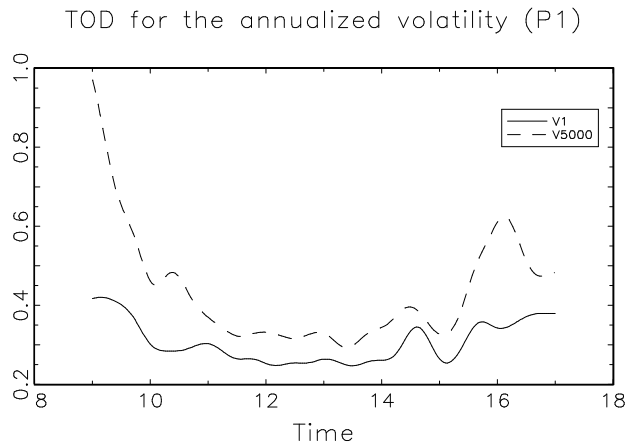

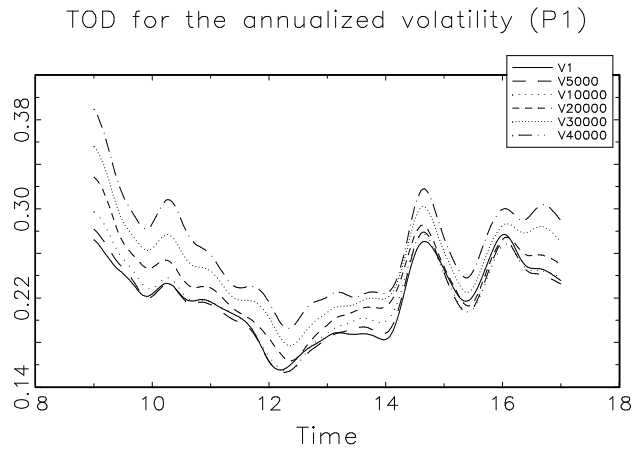


Figure 6: Time-of-day function for the annualized volatility of the ask-bid return, Daimler-Chrysler (top), Deutsche Telekom (middle) and SAP (bottom), 2 August to 19 September 1999

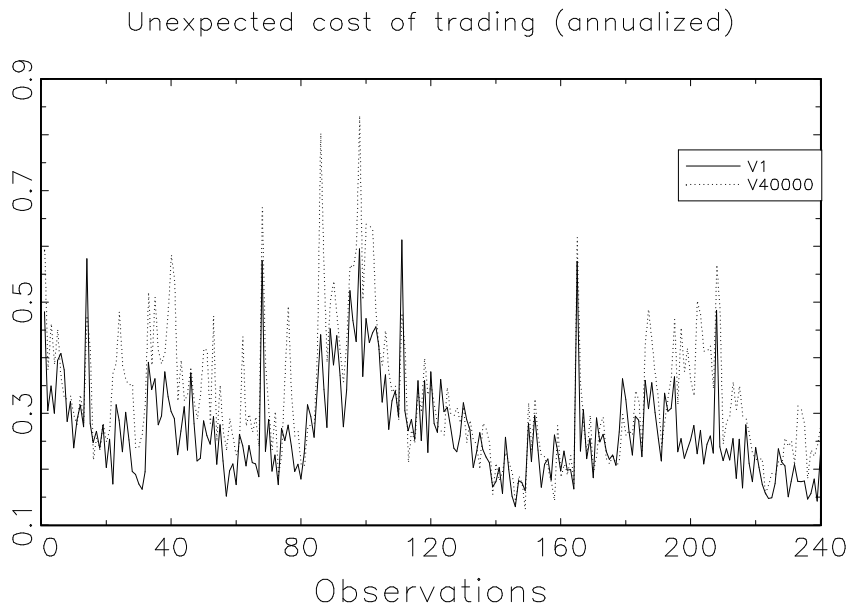



Figure 7: Trading risk for DCX (top) and DTE (bottom), ask-bid 5 minute returns