

How large is liquidity risk in an automated auction market?

Pierre Giot

Joachim Grammig*

September 21, 2002

ABSTRACT

We introduce a new empirical methodology that takes account of liquidity risk in a Value-at-Risk framework, and quantify liquidity risk premiums for portfolios and individual stocks traded on the automated auction market Xetra which operates at various European exchanges. When constructing liquidity risk measures we allow for the potential price impact incurred by the liquidation of a portfolio. We study the sensitivity of liquidity risk towards portfolio size and VaR time horizon, and interpret its diurnal variation in the light of market microstructure theory.

*Pierre Giot is from Department of Business Administration & CEREFIM at University of Namur, Rempart de la Vierge, 8, 5000 Namur, Belgium. Phone: +32 (0) 81 724887. Email: pierre.giot@fundp.ac.be. Joachim Grammig is from the Swiss Institute of Banking and Finance, University of St. Gallen, Rosenbergstr. 52, CH-9000 St. Gallen, Switzerland. Phone: +41 71 224 70 90. Fax: +41 71 224 7088. Email: joachim.grammig@unisg.ch. Both authors are research fellows at CORE, Université Catholique de Louvain, Belgium. We are grateful to Deutsche Boerse AG for providing access to the limit order data and to Kai-Oliver Maurer and Uwe Schweickert who provided invaluable expertise regarding the Xetra trading system, as well as Rico von Wyss and Michael Genser who offered helpful comments. We also thank Helena Beltran-Lopez for her cooperation in the preparation of the datasets and Bogdan Manescu for research assistance.

According to the 1988 Basel Accord the total (market risk) capital requirement for a financial institution is the sum of the requirements of positions in four different categories which are equities, interest rates, foreign exchange and gold and commodities. This sum is a major determinant of the eligible capital of the financial institution based on the 8% rule. The 1996 Amendment proposed an alternative approach for determining the market risk capital requirement, allowing the use of an internal model (subject to strong qualitative and quantitative requirements) in order to compute the maximum loss over 10 trading days at a 1% confidence level. This set the stage for Value-at-Risk models which can be broadly defined as quantitative tools designed to assess the possible loss that can be incurred by a financial institution over a given time period and for a given portfolio of assets.¹

In this paper we propose a new empirical methodology that explicitly accounts for liquidity risk when computing VaR measures. Using data from the automated auction system Xetra we investigate the dependence of liquidity risk on portfolio size and VaR time horizon, and interpret intra-day variations of liquidity risk premiums in the light of market microstructure theory.

In economics and finance, the notion of liquidity is generally conceived as the ability to trade quickly a large volume with minimal price impact. In an attempt to grasp the concept more precisely, Kyle (1985) identifies three dimensions of liquidity: tightness (reflected in the bid-ask spread), depth (the amount of one-sided volume that can be absorbed by the market without causing a revision of the bid-ask prices), and resiliency (the speed of return to equilibrium). Liquidity aspects enter the Value-at-Risk methodology quite naturally. The VaR approach is built on the hypothesis that “market prices represent achievable transaction prices” (Jorion, 2000). In other words, the prices used to compute market returns in the VaR models have to be representative of market conditions and traded volume. Consequently, the price impact of portfolio liquidation has to be taken into account. The debacle of the Long-Term Capital Management hedge

fund has shown that the price impact of the liquidation can be substantial and failing to account of liquidity risk might even stir economies as a whole.² However, empirical VaR analyses undertaken by academics and practitioners continue to use in almost all cases mid-quote prices as inputs, and disregard potential liquidity risk. Some recent contributions to the VaR literature have begun to address this issue. Subramanian and Jarrow (2001) characterize the liquidity discount (the difference between the market value of a trader's position and its value when liquidated) in a continuous time framework. Empirical models incorporating liquidity risk are developed in Jorion (2000), or Bangia, Diebold, Schuermann, and Stroughair (1999), but none of the methods does explicitly take into account the price impact incurred when liquidating a portfolio of assets. Instead, liquidity risk is approximated by and derived from the volatility of the inside spread (the difference between the best bid and ask price).

In this paper we will show that with suitable data at hand, the VaR can be adjusted for liquidity risk by explicitly modeling the price impact incurred by a trade of a given volume. In contrast to a standard (frictionless) VaR approach, in which one uses prices based on mid-quotes, the *Actual VaR* approach pursued in this paper uses as inputs volume-dependent transaction prices. This takes into account the fact that buyer (seller) initiated trades incur increasingly higher (lower) prices per unit share as the trade volume increases. The VaR liquidity risk component naturally originates from the volume dependent price impact incurred when the portfolio is liquidated. The Actual Var approach relies on the availability of intra-day bid and ask prices valid for the immediate trade of any volume of interest. Admittedly, procuring such data from traditional market maker systems would be an extremely tedious task. However, the advent of modern automated auction systems offers a new possibilities for empirical research. Using a unique database containing records of all relevant events occurring in an automated auction system, we construct real-time order book histories over a three-month period and compute time series of potential price impacts incurred by trading a given

portfolio of assets. Based on this data we estimate liquidity adjusted VaR measures and liquidity risk premiums for portfolios and single assets.

Our empirical results reveal a pronounced diurnal variation of liquidity risk which is consistent with predictions of microstructure information models. We show that when assuming a trader’s perspective, accounting for liquidity risk becomes a crucial factor: the traditional (frictionless) measures severely underestimate the true VaR. When the VaR time horizon is increased assuming the regulator’s perspective defined in the Basel Accord, liquidity risk is reduced compared to market risk, albeit remaining an economically significant factor as far as medium and large portfolios are concerned.

The remainder of the paper is organized as follows: In Section I, we provide background information about the organization of automated auction systems in general, and the Xetra system in particular. Section II describes our data set. The empirical method is developed in Section III. Results are reported in Section IV. Section V concludes and offers possible new research directions.

I. Market structure

In a dealership market, one or more dealers/market makers act as suppliers of liquidity. Market microstructure theory shows how inventory and asymmetric information effects account for the fact that the market maker’s quotes and depths – two of Kyle’s liquidity dimensions – become progressively less favorable as the traded volume increases (see O’Hara (1995) and Madhavan (2000) for reviews). Because no dedicated market makers are present in modern automated auction markets, liquidity supply solely depends on the state of the electronic order book which consists of previously entered, non-executed limit buy and sell orders. This set of standing orders determines the price-volume relationship that a trader who requires immediacy of execution is facing. If few limit buy or sell orders are present in the system or if many orders are present but are valid only for small trade

sizes, liquidity is low and trades may incur considerable price impacts. Because of the price and time priority rules implemented at automated auction markets, the price impact of a buy (sell) side trade is an increasing (decreasing) function of the trade size. Studying the Swedish stock index futures market Coppejans, Domowitz, and Madhavan (2001) consider as a key statistic for measuring liquidity the unit price obtained when selling v shares at time t :

$$b_t(v) = \frac{\sum_k b_{k,t} v_{k,t}}{v} \quad (1)$$

where v is the volume executed at k different unique bid prices $b_{k,t}$ with corresponding volumes $v_{k,t}$ standing in the limit order book at time t . This simple measure is able to meet Kyle's requirements for a liquidity measure by accounting simultaneously for tightness, depth and, by studying its time series dynamics, resiliency.³

In our empirical analysis we will use data from the automated auction system Xetra which is employed at various European trading venues, like the Vienna Stock Exchange, the Irish Stock Exchange and the European Energy Exchange. Xetra was developed and is maintained by the German Stock Exchange and has operated since 1997 as the main trading platform for German blue chip stocks at the Frankfurt Stock Exchange (FSE). Whilst there still exist market maker systems operating parallel to Xetra - the largest of which being the Floor of the Frankfurt Stock Exchange- the importance of those venues has been greatly reduced, especially regarding liquid blue chip stocks. Similar to the Paris Bourse's CAC and the Toronto Stock Exchange's CATS trading system, a computerized trading protocol keeps track of entry, cancellation, revision, execution and expiration of market and limit orders. Until September 17, 1999, Xetra trading hours at the FSE extended from 8.30 a.m to 5.00 p.m. CET. Beginning with September 20, 1999 trading hours were shifted to 9.00 a.m to 5.30 p.m. CET. Between an opening and a closing call auction - and interrupted by a another mid-day call auction - trading is based on a continuous double auction mechanism with automatic matching of orders

based on clearly defined rules of price and time priority. Only round lot sized orders can be filled during continuous trading hours. Execution of odd-lot parts of an order (representing fractions of a round lot) is possible only in a call auction. During pre- and post-trading hours it is possible to enter, revise and cancel orders, but order executions are not conducted, even if possible.

According to a taxonomy introduced by Domowitz (1992) Xetra may be described as a “hit and take” system.⁴ Until October 2000, Xetra screens displayed not only best bid and ask prices, but the whole content of the order book to the market participants. This implies that liquidity supply and potential price impact of a market order (or marketable limit order) were exactly known to the trader. This was a great difference compared to e.g. Paris Bourse’s CAC system where “hidden” orders (or “iceberg” orders) may be present in the order book. As the name suggests, a hidden limit order is not 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. Iceberg orders were allowed in Xetra in October 2000, heeding the request of investors who were reluctant to see their (potentially large) limit orders, i.e. their investment decisions, revealed in the open order book.

The transparency of the Xetra order book does not extend to revealing the identity of the traders submitting market or limit orders. Instead, Xetra trading is completely anonymous and dual capacity trading, i.e. trading on behalf of customers and principal trading by the same institution is not forbidden.⁵ In contrast to a market maker system there are no dedicated providers of liquidity, like e.g. the NYSE specialists, at least not for blue chip stocks studied in this paper. For some small cap stocks listed in Xetra there may exist so-called Designated Sponsors - typically large banks - who are obliged, but not forced to, provide a minimum liquidity level by simultaneously submitting competing buy and sell limit orders.

II. Data

The German Stock Exchange granted access to a database containing complete information about Xetra open order book events (entries, cancellations, revisions, expirations, partial-fills and full-fills of market and limit orders) that occurred between August 2, 1999 and October 29, 1999.⁶ Due to the considerable amount of data and processing time, we had to restrict the number of assets we deal with in this study. Event histories were extracted for three blue chip stocks, DaimlerChrysler (DCX), Deutsche Telekom (DTE) and SAP. By combining these stocks we form small, medium and large portfolios as it could be argued that estimating the Value-at-Risk is interesting not so much at stock level, but on the level of (well-diversified) portfolios. At the end of the sample period the combined weight of DaimlerChrysler, SAP and Deutsche Telekom in the DAX - the value weighted index of the 30 largest German stocks - amounted to 30.4 percent (October 29, 1999). Hence, the liquidity risk associated with the three stock portfolios is quite representative of the liquidity risk that an investor faces when liquidating the market portfolio of German Stocks.

Based on the event histories we perform a real time reconstruction of the order book sequences. Starting from an initial state of the order book, we track each change in the order book implied by entry, partial or full fill, cancellation and expiration of market and limit orders. This is done by implementing the rules of the Xetra trading protocol outlined in Deutsche Börse AG (1999) in the reconstruction program.⁷ From the resulting real-time sequences of order books snapshots at 10 and 30-minute frequencies during the trading hours were taken. For each snapshot, the order book entries were sorted on the bid (ask) side in price descending (price ascending) order. Based on the sorted order book sequences we computed the unit price $b_t(v)$, as defined in Equation (1), implied by selling at time t volumes v of 1, 5,000, 20,000, and 40,000 shares, respectively. Mid-quote prices were computed as the average of best bid and ask prices prevailing at time t . Of course these are equivalent to $b_t(1)$ and $a_t(1)$, respectively. If the trade

volume v exceeds the depth at the prevailing best quote then $b_t(v)$ will be smaller than $b_t(1)$ (and $a_t(v) > a_t(1)$). By varying the trade volume v one can plot the slope of the instantaneous offer and demand curves.

III. Methodology

Bangia, Diebold, Schuermann, and Stroughair (1999) (henceforth referred to as BDSS) suggest a liquidity risk correction procedure for the Value-at-Risk framework. BDSS relate the liquidity risk component to the distribution of the inside half-spread. In the first step of the procedure, the VaR is computed as the α percent quantile of the mid-quote return distribution (assuming normality). This quantile is then increased by a factor based on the excess kurtosis of the returns. In a second step, liquidity cost is allowed for by taking as inputs the historical average half-inside-spread and its volatility. This adjusts the VaR for the fact that buy and sell orders are not executed at the quote mid-point, but that (extreme) variations in the spread may occur. BDSS assume a perfect correlation between the frictionless VaR and the exogenous cost of liquidity. This yields the total VaR being equal to the sum of the market VaR and liquidity cost. Switching from returns to price levels, BDSS express the VaR at level α (including liquidity costs) as:

$$P_t = \frac{a_t(1) + b_t(1)}{2} \left[(1 - e^{\mu + Z_\alpha \sigma}) + \frac{1}{2} (\mu_S + Z'_\alpha \sigma_S) \right] \quad (2)$$

where μ and σ are the mean and volatility of the market (mid-quote) returns, μ_S and σ_S are the mean and volatility of the relative spread, Z_α and Z'_α are the α percent quantiles of the distribution of market returns and spread respectively and P_t is the VaR at level α (expressed as a price) taking into account market risk and liquidity costs.

The BDSS procedure offers the possibility to allow for VaR liquidity risk when only best bid best ask prices are available. This is, for example, the case when using the

popular TAQ data supplied by NYSE. A volume dependent price impact is, quite deliberately, not taken into account as such information cannot be procured from standard databases. However, a more precise way to allow for liquidity risk becomes feasible with richer data at hand. The approach pursued in this paper relies on the availability of time series of intra-day bid and ask prices valid for the immediate trade of a given volume. In a market maker setting this requires a time series of quoted bid and ask prices for a given volume. In an automated auction market, unit bid and ask prices can be computed according to Equation (1) using open order book data. Obtaining such data for a market maker system will be almost impossible. As market makers are obliged to quote only best bid and ask prices with associated depths, quote driven exchanges can and will at best supply this limited information set for financial market research. As a matter of fact, this is the situation where the BDSS approach adds the greatest value in correcting VaR for liquidity risk. In a computerized auction market much richer data can be exploited. As the automated trading protocol keeps track of and records all events occurring in the system it is possible to reconstruct real time series of limit order books from which the required unit bid prices $b_t(v)$ can be straightforwardly computed.

In order to compute the liquidity risk measures to be introduced below, econometric specifications for two return processes are required. First, for mid-quote returns (referred to as *frictionless* returns) which are defined as the log ratio of consecutive mid-quotes:⁸

$$r_{mm,t} = \ln \frac{a_t(1) + b_t(1)}{a_{t-1}(1) + b_{t-1}(1)}.$$

Second, for *actual* returns which are defined as the log ratio of mid-quote and consecutive unit bid price valid for selling a volume v shares at time t :

$$r_{mb,t}(v) = \ln \frac{b_t(v)}{0.5(a_{t-1}(1) + b_{t-1}(1))}.$$

For the analysis of liquidity risk associated with a portfolio consisting of $i = 1, \dots, N$ assets with volumes v^i , actual returns are obtained by computing the log ratio of the market value when selling the portfolio at time t , $\sum_{i=1}^N b_t(v^i)v^i$, and the value of the portfolio evaluated at time $t - 1$ mid-quote prices. To compute frictionless portfolio returns, the portfolio is evaluated at mid-quote prices both at t and $t - 1$.

For both types of returns the VaR is estimated in the standard way, namely as the one-step ahead forecast of the α percent return quantile. We refer to the VaR computed on the $\{r_{mb,t}(v)\}_{t=1}^T$ returns sequence as the *Actual VaR*. Our econometric specifications of the return processes build on previous results on the statistical properties of intra-day spreads and return volatility. Two prominent features of intra-day return and spread data have to be accounted for. First, spreads feature considerable diurnal variation (see e.g. Chung, Van Ness, and Van Ness (1999)). Microstructure theory suggests that inventory and asymmetric information effects play a crucial role in procuring these variations. Information models predict that liquidity suppliers (market makers, limit order traders) widen the spread in order to protect themselves against potentially superiorly informed trades around alleged information events, such as the open. Second, as shown by e.g. by Andersen and Bollerslev (1997), conditional heteroskedasticity and diurnal variation of return volatility have to be taken into account. When specifying the conditional mean of the actual return processes we therefore allow for diurnal variations in actual returns, since these contain, by definition, the half-spread. We adopt the specification of Andersen and Bollerslev (1997) to allow for volatility diurnality and conditional heteroskedasticity in the actual return process. Furthermore, diurnal variations in mean returns and return volatility are assumed to depend on the trade volume, as suggested by Gouriéroux, Le Fol, and Jasiak (1999). For convenience of notation we suppress the volume dependence of actual returns and write the model as:

$$r_{mb,t} = \psi_t + \delta_0 + \sum_{i=1}^r \delta_i r_{mb,t-i} + u_t \quad (3)$$

where

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

$$h_t = \omega + \sum_{i=1}^q \alpha_i u_{t-i}^2 + \sum_{i=1}^p \beta_i h_{t-i}. \quad (5)$$

The innovations ϵ_t are assumed to be independently identically Student distributed with ν_1 degrees of freedom. The functions ψ_t and ϕ_t account for diurnal variation in the level of actual returns and return volatility, respectively. We have suppressed the volume dependence of actual returns only for brevity of notation. The discerning reader will recognize that in a more extensive notation all greek letters would have to be written with a volume index v . Accordingly, the model parameters are estimated for each volume dependent actual return process. We employ a four-step procedure that is described as follows.

First, the diurnal component ψ_t is estimated by a non-parametric regression approach. Given returns available at s -minute sampling frequency, we sub-divide the trading day into s -minute bins, compute the average actual return (over all days in the sample) by bin and smooth the resulting time series using the Nadaraya-Watson estimator.⁹ In the second step, a time series of diurnally adjusted returns is obtained by subtracting the estimate $\hat{\psi}_t$ from the actual return $r_{mb,t}$. The resulting time series is used to estimate the AR-parameters process by OLS. The sequence of AR residuals provides the input for modelling actual return volatility. In the third step, the diurnal volatility function ϕ_t is estimated non-parametrically by applying the Nadaraya-Watson estimator to the estimated squared AR residuals, $\{\hat{u}_t^2\}$ which are sorted in s -minute bins. In step four, the squared AR residuals are divided by the estimates $\hat{\phi}_t$. The resulting series is finally used to estimate the GARCH parameters by conditional Maximum Likelihood.¹⁰

This specification implies that the conditional standard deviation of the actual return at time t , $\sigma_t(r_{mb,t}(v))$, evolves as:

$$\sigma_t(r_{mb,t}) = \sqrt{\phi_t h_t}. \quad (6)$$

The Actual VaR at time $t - 1$ for the actual return at time t given confidence level α is then given by:

$$VaR_{mb,t} = \mu_{mb,t} + t_{\alpha,\nu_1} \sigma_t(r_{mb,t}) \quad (7)$$

where

$$\mu_{mb,t} = \psi_t + \delta_0 + \sum_{i=1}^r \delta_i (r_{mb,t-i} - \psi_{t-i})$$

and t_{α,ν_1} is the α -percent quantile of the Student distribution with ν_1 degrees of freedom.¹¹ We refer to $\mu_{mb,t}$ as the mean component and to $t_{\alpha,\nu_1} \sigma_t(r_{mb,t})$ as the volatility component of the Actual VaR.

The computation of the liquidity risk premium measures which will be discussed below also requires a VaR estimate based on mid-quote returns (referred to as frictionless VaR). The econometric specification corresponds to the Actual VaR with the exception that there is no need to account for a diurnal variation in mean returns.¹² For notational convenience let us use the same greeks as for the actual return specification:

$$\begin{aligned} r_{mm,t} &= \mu_{mm,t} + u_t \\ \mu_{mm,t} &= \delta_0 + \sum_{i=1}^r \delta_i r_{mm,t-i} \\ u_t &= \sqrt{\phi(t) h_t} \epsilon_t \\ h_t &= \omega + \sum_{i=1}^q \alpha_i u_{t-i}^2 + \sum_{i=1}^p \beta_i h_{t-i}. \end{aligned} \quad (8)$$

The innovations ϵ_t are assumed to be independently identically Student distributed with ν_2 degrees of freedom. ϕ_t accounts for diurnal variation in frictionless return volatility. Parameter estimation is performed along the same lines as outlined above. The frictionless VaR at α percent confidence level is:

$$VaR_{mm,t} = \mu_{mm,t} + t_{\alpha,\nu_2}\sigma_t(r_{mm,t}). \quad (9)$$

where $\sigma_t(r_{mm,t}) = \sqrt{\phi(t)h_t}$. As above we refer to the two terms on right hand side of Equation (9) as mean and volatility component of the frictionless VaR.

We propose to quantify liquidity risk by comparing frictionless and Actual VaR. Two liquidity risk premium measures are used, one based on the difference, the other on the ratio of frictionless and Actual VaR:

$$\Lambda_t = VaR_{mm,t} - VaR_{mb,t} \quad (10)$$

$$\lambda_t = \frac{\Lambda_t}{VaR_{mm,t}}. \quad (11)$$

Omitting the economically negligible mean component of the frictionless VaR, we can rearrange Equation (11) and write the relative liquidity risk premium as:

$$\lambda_t = \left[\frac{\mu_{mb,t}}{t_{\alpha,\nu_2}\sigma_t(r_{mm,t})} \right] + \left[\frac{t_{\alpha,\nu_1}\sigma_t(r_{mb,t})}{t_{\alpha,\nu_2}\sigma_t(r_{mm,t})} - 1 \right]. \quad (12)$$

Equation (12) shows that λ_t can be conceived as the sum of two terms naturally referred to as mean and volatility component of the (relative) liquidity risk premium. Note that in a more extensive notation, one would write the dependence of both liquidity premiums on portfolio size v and confidence level α .

We want to stress two points before applying the liquidity measures Λ_t and λ_t . First, λ_t is, by definition, a relative, conditional measure. If the VaR horizon is short and the volatilities of both actual and frictionless returns are relatively small and of

comparable size then the mean component of the Actual VaR will be the most important determinant of λ_t . Ceteris paribus, the importance of the actual return mean component will increase with trade volume and so will both liquidity risk premium measures. At longer VaR horizons, both actual and frictionless return volatility will naturally increase due to non-liquidity related market risk. This reduces the relative importance of the λ -mean component as the denominator of the first term in Equation (12) grows. If both actual and frictionless return volatility increase by the same factor then the relative liquidity risk premium is expected to approach zero at longer VaR horizons. Second, when studying intra-day variations of the liquidity risk premium the difference measure Λ_t is more appropriate. Both spreads and volatility of intra-day returns are expected to exhibit diurnal variation. By construction, small changes in the intra-day frictionless return volatility may exert a considerable impact on the diurnal variation of λ_t , whilst Λ_t is robust against such fluctuations.

IV. Empirical results

A. Parameter estimates

The Nadaraya-Watson, OLS and Maximum likelihood estimates are obtained using GAUSS procedures written by the authors. Table I reports the estimates of the parametric model part (AR-GARCH parameters) based on 10-minute returns of equal volume stock (EVS) portfolios and individual stocks. Table III (deferred to the appendix) contains the half-hour frequency results.¹³ An EVS portfolio contains the same number of shares for each stock in the portfolio. We will henceforth generally use the notion *portfolio* both for EVS portfolios and single stocks, conceiving the latter as a portfolio containing only a single stock. We report results on small volume ($v = 5,000$), medium volume ($v = 20,000$) and big volume ($v = 40,000$) portfolios. Based on the Schwarz-Bayes-Criteria and Ljung-Box statistics, AR(1)-GARCH(1,1) (half-hour frequency) and

AR(3)-GARCH(1,1) (10-minute frequency) specifications were selected. Whilst the AR parameter estimates and the Ljung-Box statistics computed on raw returns indicate only small autocorrelations of frictionless returns (as expected in an at least weakly efficient market) , the serial dependence of the actual returns on the medium and big portfolios is more pronounced. This results holds true especially at the higher frequency and indicates that persistence in spreads increase with trade volume. After accounting for mean diurnality and serial dependence in actual returns the AR residuals do not display significant autocorrelations. Comparing the Ljung-Box statistic before and after accounting for conditional heteroskedasticity and volatility diurnality reveals that the model does a good job in reducing serial dependence in squared returns. The GARCH parameter estimates and degrees of freedom are quite stable across portfolio sizes and their order of magnitude is comparable to what is found when estimating GARCH models on intra-day returns (see Andersen and Bollerslev (1997)). Based on the estimation results we compute frictionless and Actual VaR at $\alpha = 0.05$ as well as the sequence of relative and difference liquidity measures Λ_t and λ_t .

Insert Table I about here.

B. The diurnal variation of liquidity risk

For the purpose of studying intra-day variation of liquidity risk we take sample averages of the difference measure Λ_t by time of day t , smooth the resulting series by applying the Nadaraya-Watson estimator, and investigate its diurnal variation during trading hours. Figure 1 displays the considerable diurnal variation of the liquidity risk premium especially for big portfolios. Liquidity risk is highest at the start of the trading day and sharply declines during the next two hours whilst remaining at a constant level throughout the remainder of the trading day. This pattern is stable across both sample sub-periods with different trading hours. A trader who plans to sell large volumes at

the start of the trading day is expected to incur a significant price impact, i.e. has to be ready to pay a considerable liquidity risk premium. The diversification effect smoothes the intra-day variation, but the pronounced liquidity risk premium during the first half hours after the open cannot be diversified away.

Insert Figure 1 about here.

Figure 2 details this finding by displaying the intra-day variation of mean and volatility component of the Actual VaR for the big EVS portfolio ($v = 40,000$).¹⁴ Both mean and volatility component contribute to the diurnal variation of the Actual VaR and hence to the time-of-day pattern of the liquidity risk premium. In the afternoon, NYSE pre-trading exerts an effect on volatility component of both frictionless and Actual VaR, but as both VaR measures are affected by the same order of magnitude, the liquidity risk premium is not affected.

Insert Figure 2 about here.

The intra-day pattern of the liquidity risk premium and Actual VaR provides additional empirical support for the information models developed by Madhavan (1992) and Foster and Viswanathan (1994). Madhavan (1992) considers a model in which information asymmetry is gradually resolved throughout the trading day implying higher spreads at the opening. In the Foster and Viswanathan (1994) model, competition between informed traders leads to high return volatility and spreads at the start of trading. Analyzing NYSE intra-day liquidity patterns using the inside spread, Chung, Van Ness, and Van Ness (1999) have argued that the high level of the spread at the NYSE opening and its subsequent decrease provides evidence for the information models à la Madhavan and Foster/Viswanathan. Accordingly, the diurnal variation of liquidity risk is consistent with the predictions implied by those models. Due to alleged information asymmetries, liquidity suppliers are initially cautious, i.e. the liquidity risk premium is

large. As the information becomes gradually incorporated during the trading process, the liquidity risk premium decreases with increasing liquidity supply.

C. Unconditional liquidity risk premiums from traders' and regulators' perspective

The aftermath of the LTCM debacle showed that disregard of liquidity risk associated with intra-day trading of large volumes can lead to devastating results even from a macroeconomic perspective. Let us assess the importance of short term liquidity risk in the present sample. The relative liquidity risk measure λ as well as the difference measure Λ are defined as *conditional* measures given information at time $t - 1$. One can estimate the *unconditional* liquidity risk premium $\lambda = \mathbb{E}(\lambda_t)$ by taking sample averages

$$\bar{\lambda} = T^{-1} \sum_{t=1}^T \lambda_t = T^{-1} \sum_{t=1}^T \frac{\mu_{mb,t}}{t_{\alpha,\nu_2} \sigma_t(r_{mm,t})} + T^{-1} \sum_{t=1}^T \frac{t_{\alpha,\nu_1} \sigma_t(r_{mb,t})}{t_{\alpha,\nu_2} \sigma_t(r_{mm,t})} - 1. \quad (13)$$

and study the dependence of the unconditional liquidity risk premium on the size of the portfolio to be liquidated. Equation 13 shows that the decomposition of the relative liquidity risk premium λ_t into mean and volatility component remains valid for the unconditional liquidity risk premium, too.

Insert Table II about here.

Table II reports the estimated unconditional liquidity risk premium $\bar{\lambda}$. The decomposition into mean and volatility component is contained in Table IV. The results show that taking account of liquidity risk at the intra-day level is quite crucial. Even at medium portfolio size, the liquidity risk premium is considerable. At half-hour horizon the underestimation of the VaR of the medium EVS portfolio amounts to 34 percent. For the big EVS portfolio the VaR is underestimated at half-hour horizon by 61 percent.

At the shorter horizon the underestimation becomes even more severe with the medium EVS portfolio's VaR being underestimated by 68 percent. The decomposition exercise shows that when increasing portfolio size the volatility component of the liquidity risk premium remains small relative to the mean component since the price impact incurred by trading a large portfolio becomes the dominating factor. The relative liquidity risk premium decreases, *ceteris paribus*, at the longer VaR horizon. The reason is that both the frictionless and the Actual VaR's volatility component obey the square root of time rule, and increase (in absolute terms) by about the same order of magnitude. This reduces the significance of the mean component while the order of magnitude of the volatility component remains approximately constant. In other words, market risk compared to liquidity risk becomes more important at longer horizons.

Because of the nature of portfolio liquidation in the banking industry, the time horizon for VaR models considered in the Basel Accord has been a longer one. From a regulatory point of view one could therefore be interested in assessing the economic significance of the liquidity risk factor at the one and ten-day VaR time horizon. Referring to the square root of time rule one could use a high frequency estimate (e.g. 10-minute frequency) of frictionless and Actual VaR and multiply the volatility components by a factor \sqrt{f} that scales short horizon into long horizon volatility.¹⁵ For example, when half-hour frequency estimates are available, and one wants to assess the liquidity risk premium at the one-hour horizon then $f = 2$. Taking, as before, sample averages one can approximate the unconditional relative liquidity premium at the desired frequency by

$$\bar{\lambda} = T^{-1} \sum_{t=1}^T \frac{\mu_{mb,t}}{t_{\alpha,\nu_2} \sqrt{f} \sigma_t(r_{mm,t})} + T^{-1} \sum_{t=1}^T \frac{t_{\alpha,\nu_1} \sqrt{f} \sigma_t(r_{mb,t})}{t_{\alpha,\nu_2} \sqrt{f} \sigma_t(r_{mm,t})} - 1. \quad (14)$$

If VaR horizon and sampling frequency coincide then $f = 1$. Hence, we did not change the notation and continue to write $\bar{\lambda}$ to denote unconditional liquidity risk premiums. We apply the approximation formula to the one and ten-day horizons using the half hour Actual VaR estimates as a basis and scale to daily ($f = 17$) and ten day horizon ($f =$

170). The results reported in Table IV indicate that the relative importance of liquidity risk premium is reduced at longer horizons, since market risk dominates liquidity risk. Yet, liquidity risk remains economically significant for larger portfolios. Disregarding liquidity risk, the one-day VaR estimate for the big EVS portfolio is underestimated by 12 percent, a non-negligible amount from the regulatory perspective. At the 10-day horizon, the underestimation is further reduced to 2 percent.¹⁶

V. Conclusion and outlook

This paper quantified liquidity risk in an automated auction market by employing a new empirical technique which extends the classical, frictionless VaR methodology. The notion of an Actual VaR measure was introduced which takes account of the potential price impact of liquidating a portfolio. Liquidity risk measures were proposed that are defined on the difference and ratio of Actual and standard (frictionless) VaR. It was argued that automated auction markets, in which a computer manages order entry and matching, provide suitable and accurate data for the task of estimating the Actual VaR. In order to measure liquidity risk using the methodology pursued in this paper one has to provide a real time limit order book reconstruction from which price impacts of trading a given volume can be computed. Using data from the automated auction system Xetra liquidity risk was quantified both for portfolios and for individual stocks. The dependence of liquidity risk premiums on time-of-day, trade volume and VaR time horizon was outlined.

The analysis revealed a pronounced diurnal variation of the liquidity risk premium. The peak of the liquidity risk premium at the open and its subsequent decrease is consistent with the predictions of the microstructure information models considered by Madhavan (1992) and Foster and Viswanathan (1994). The results show that when assuming a trader's perspective, accounting for liquidity risk becomes a crucial factor:

the traditional (frictionless) measures severely underestimate the Value-at-Risk. This result is the more pronounced the bigger the portfolio size and the shorter the time horizon. When the VaR time horizon is increased to the regulator's perspective given in the Basel Accord, the relative importance of liquidity risk is reduced compared to market risk, albeit remaining an economically significant factor as far as medium and large portfolios are concerned.

Avenues for further research stretch in several directions. In the classification of Dowd (1998), this paper has focused on the normal liquidity risk in contrast to crisis liquidity risk. The latter refers to situations where "a market can be very liquid most of the time, but lose its liquidity in a major crisis" (Dowd, 1998). The methodology applied in this paper could be readily used to evaluate crisis liquidity risk using intra-day data specific to such crisis periods where liquidity dried up for a few days/weeks (for example during the Asian crisis of the summer months of 1997 or close to the LTCM debacle in September 1998). Second, in a cross section study for a larger collection of stocks one could relate liquidity risk premiums to firm-specific characteristics like e.g. the equity-debt ratio. This would facilitate investigating the question whether corporate financing decisions or ownership structure affect firm specific liquidity risk.

References

- Andersen, T.G., and T. Bollerslev, 1997, Intraday periodicity and volatility persistence in financial markets, *Journal of Empirical Finance* 4, 115–158.
- Bangia, A., F.X. Diebold, T. Schuermann, and J.D. Stroughair, 1999, Modeling liquidity risk, with implications for traditional market risk measurement and management, The Wharton Financial Institutions Center WP 99-06.
- Bauwens, L., and P. Giot, 2001, *Econometric modelling of stock market intraday activity*. (Kluwer Academic Publishers).

- Beltran, H., P. Giot, and J. Grammig, 2002, Expected and unexpected cost of trading in the XETRA automated auction market, Mimeo, CORE.
- Biais, B., P. Hillion, and C. Spatt, 1999, Price discovery and learning during the pre-opening period in the Paris Bourse, *Journal of Political Economy* 107, 1218–1248.
- Buhl, C., and R. von Wyss, 2002, Exogenous liquidity in the Value at Risk concept, Mimeo, Universities of Basel and St. Gallen.
- Campbell, J. Y., A. W. Lo, and A. C. MacKinlay, 1997, *The Econometrics of Financial Markets*. (Princeton University Press Princeton).
- Chung, K.H., B.F. Van Ness, and R.A. Van Ness, 1999, Limit orders and the bid-ask spread, *Journal of Financial Economics* 53, 255–287.
- Coppejans, M., I. Domowitz, and A. Madhavan, 2001, Liquidity in an automated auction, Mimeo.
- Deutsche Börse AG, 1999, Xetra Market Model Release 3 Stock Trading, Technical Report.
- Domowitz, I., 1992, Automating the price discovery process: some international comparisons and regulatory implications, *Journal of Financial Services Research* 6, 305–326.
- Dowd, K., 1998, *Beyond Value-at-Risk*. (Wiley).
- Drost, F. C., and T. E. Nijman, 1993, Temporal aggregation of GARCH processes, *Econometrica* 61, 907–927.
- Drost, F. C., and B. J. M. Werker, 1996, Closing the GARCH gap: Continuous time GARCH modeling, *Journal of Econometrics* 74, 31–57.
- Engle, R.F., and J. Russell, 1998, Autoregressive conditional duration; a new model for irregularly spaced transaction data, *Econometrica* 66, 1127–1162.
- Engle, R. F., 2000, The econometrics of ultra high frequency data, *Econometrica* 68, 1–22.

- Foster, F., and S. Viswanathan, 1994, Strategic trading with asymmetrically informed investors and long-lived information, *Journal of Financial and Quantitative Analysis* 29, 499–518.
- Gomber, P. and U. Schweickert (Deutsche Börse AG), 2002, Der Market Impact: Liquiditätsmass im elektronischen Handel, *Die Bank* 7, 185–189.
- 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.
- Grammig, J., D. Schiereck, and E. Theissen, 2001, Knowing me, knowing you: Trader anonymity and informed trading in parallel markets, *Journal of Financial Markets* 4, 385–412.
- Jorion, P., 2000, *Value-at-Risk*. (McGraw-Hill).
- Kyle, A.S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.
- Lowenstein, R., 2001, *When genius failed: the rise and fall of Long-Term Capital Management*. (Fourth Estate, HarperCollinsPublishers).
- Madhavan, A., 1992, Trading mechanisms in securities markets, *Journal of Finance* 47, 607–642.
- Madhavan, A., 2000, Market microstructure: A survey, *Journal of Financial Markets* 3, 205–258.
- Martens, M., Y. Chang, and S.J. Taylor, 2002, Intraday volatility forecasts using different seasonal adjustment methods, *Journal of Financial Research*.
- O’Hara, M., 1995, *Market microstructure theory*. (Basil Blackwell Oxford).
- Saunders, A., 2000, *Financial institutions management*. (McGraw-Hill).
- Subramanian, A., and R. A. Jarrow, 2001, The liquidity discount, *Mathematical Finance* 11, 447–474.

Appendix A. Additional tables

Insert Table III about here.

Insert Table IV about here.

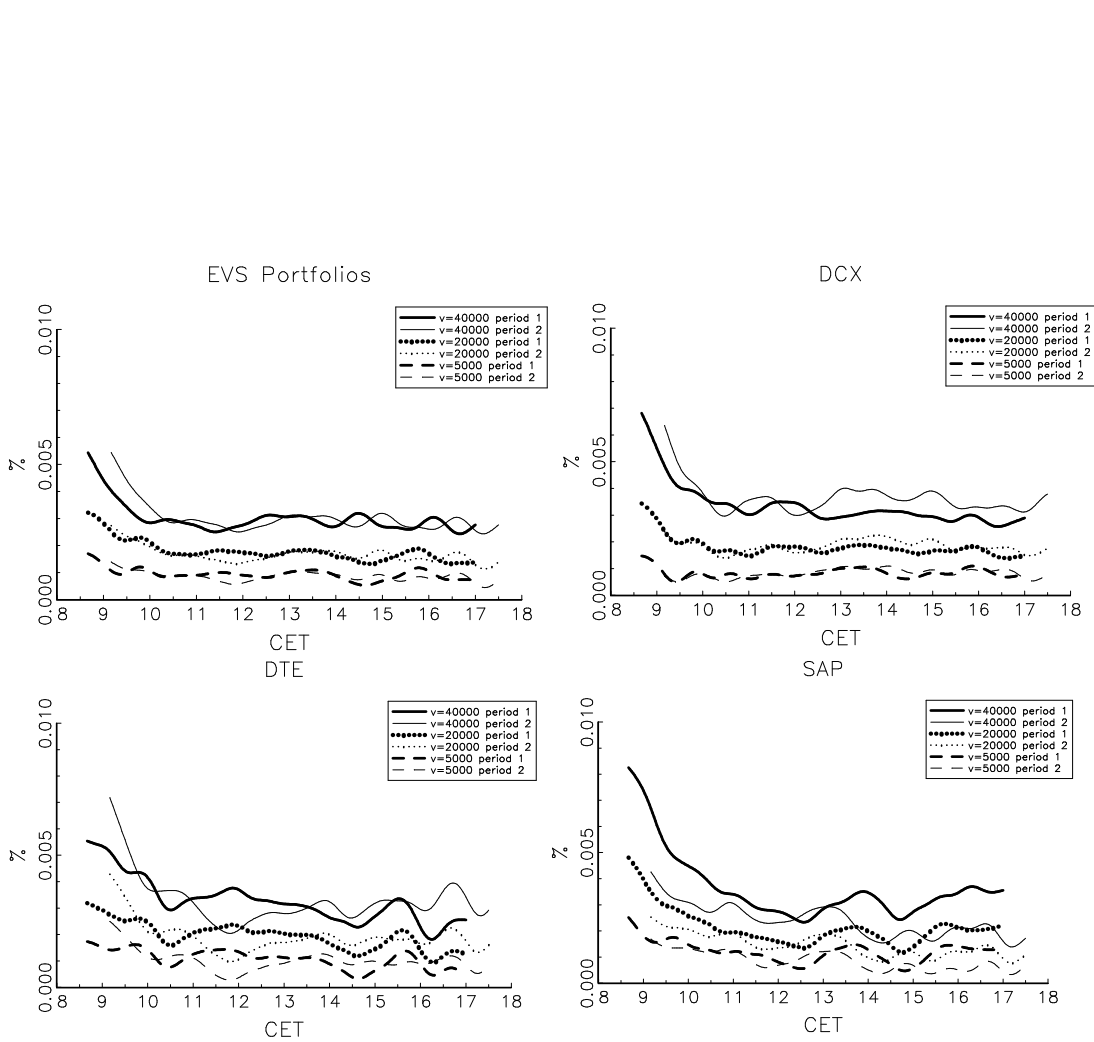


Figure 1. Diurnal variation of the liquidity risk premium Λ_t for equal volume stock portfolios and individual stocks Period 1 refers to the first half of the sample period, August, 2, 1999 - September 17, 1999 when the Xetra continuous trading hours extended from 8.30 a.m. CET - 5.00 p.m. CET. Period 2 refers to the second half of the sample period, September 20, 1999 - October 29, 1999 when Xetra continuous trading hours extended from 9.00 a.m. CET - 5.30 p.m. CET. The numbers have been multiplied by 100 to represent percentages. $\alpha = 0.05$.




Figure 2. Decomposition of liquidity premium $\Lambda_t = VaR_{mm,t} - VaR_{mb,t}$. The figure displays for the big EVS portfolio ($v = 40,000$) and for the first half of the sample period (August, 2, 1999 - September 17, 1999) the diurnal variation of the components of Λ_t . The Actual VaR ($VaR_{mb,t}$) is the sum of the mean component, $\mu_{mb,t}$, and the volatility component, $t_{\alpha, \nu_1} \sigma_t(r_{mb,t})$ (see equation 7). The frictionless VaR ($VaR_{mm,t}$) is defined in equation 9. The numbers are multiplied by 100 to represent percentages. $\alpha = 0.05$

Table I
Estimation results 10-minute frequency.

Section III provides details of the estimation procedure. The first column gives estimation results using frictionless returns. The other columns report estimation results based on actual returns. Parameter standard errors are given in parentheses. The Q -rows report the Ljung-Box Q -statistic computed on the AR residuals. For the computation of the Ljung-Box statistics the number of autocorrelations included is equal to 2, and observations of different trading days were excluded from the estimation of autocovariances. In brackets the Ljung-Box statistic of the raw return data is reported. The Q^2 rows reports the Ljung-Box statistic of the GARCH residuals. The figures in brackets are Ljung-Box statistic of squared raw returns.

	frictionless	$v = 5,000$	$v = 20,000$	$v = 40,000$
EVS portfolios				
δ_0	3.9E-05 (3.1E-05)	0.0E+00 (3.1E-05)	0.0E+00 (3.0E-05)	1.0E-06 (3.2E-05)
δ_1	0.015 (0.018)	0.020 (0.018)	0.041 (0.018)	0.108 (0.018)
δ_2	0.011 (0.018)	0.013 (0.018)	0.020 (0.018)	0.068 (0.018)
δ_3	0.015 (0.018)	0.040 (0.018)	0.055 (0.018)	0.090 (0.018)
ω	0.054 (0.027)	0.061 (0.032)	0.072 (0.032)	0.104 (0.038)
α_1	0.051 (0.015)	0.049 (0.015)	0.050 (0.014)	0.055 (0.014)
β_1	0.868 (0.050)	0.863 (0.055)	0.850 (0.052)	0.807 (0.057)
ν	9.360 (1.437)	9.996 (1.625)	10.567 (1.784)	11.346 (1.995)
Q	0.06 [1.26]	0.06 [3.35]	0.18 [17.76]	1.05 [154.91]
Q^2	4.60 [23.81]	2.49 [14.32]	0.69 [24.58]	0.17 [188.48]
DCX				
δ_0	3.0E-06 (3.7E-05)	-2.0E-06 (3.6E-05)	-1.0E-06 (3.6E-05)	-2.0E-06 (4.0E-05)
δ_1	0.025 (0.018)	0.033 (0.018)	0.082 (0.018)	0.189 (0.018)
δ_2	0.001 (0.018)	0.004 (0.018)	0.020 (0.018)	0.069 (0.018)
δ_3	0.011 (0.018)	0.024 (0.018)	0.036 (0.018)	0.079 (0.018)
ω	0.114 (0.036)	0.118 (0.051)	0.164 (0.049)	0.126 (0.048)
α_1	0.105 (0.021)	0.089 (0.023)	0.098 (0.020)	0.081 (0.019)
β_1	0.667 (0.079)	0.694 (0.104)	0.628 (0.089)	0.718 (0.085)
ν	5.249 (0.507)	5.666 (0.577)	6.312 (0.683)	7.230 (0.861)
Q	0.02 [2.96]	0.06 [5.65]	0.03 [34.57]	0.34 [258.85]
Q^2	1.56 [58.49]	1.25 [42.07]	3.49 [47.48]	2.78 [228.73]
DTE				
δ_0	7.6E-05 (4.9E-05)	2.0E-06 (4.9E-05)	2.0E-06 (4.9E-05)	2.0E-06 (5.2E-05)
δ_1	0.013 (0.018)	0.033 (0.018)	0.062 (0.018)	0.121 (0.018)
δ_2	0.030 (0.018)	0.029 (0.018)	0.034 (0.018)	0.067 (0.018)
δ_3	0.005 (0.018)	0.015 (0.018)	0.038 (0.018)	0.070 (0.018)
ω	0.052 (0.021)	0.040 (0.021)	0.056 (0.022)	0.082 (0.027)
α_1	0.055 (0.013)	0.041 (0.013)	0.057 (0.014)	0.075 (0.015)
β_1	0.841 (0.046)	0.883 (0.045)	0.843 (0.046)	0.789 (0.051)
ν	5.657 (0.578)	6.009 (0.657)	6.704 (0.784)	7.497 (0.917)
Q	0.19 [4.36]	0.26 [11.94]	0.46 [32.41]	0.27 [132.27]
Q^2	12.42 [82.20]	12.55 [67.59]	6.52 [74.93]	4.55 [176.03]
SAP				
δ_0	7.3E-05 (4.7E-05)	1.0E-06 (4.7E-05)	2.0E-06 (4.8E-05)	4.0E-06 (5.1E-05)
δ_1	0.017 (0.018)	0.050 (0.018)	0.109 (0.018)	0.204 (0.018)
δ_2	-0.001 (0.018)	0.009 (0.018)	0.028 (0.018)	0.074 (0.018)
δ_3	-0.011 (0.018)	0.008 (0.018)	0.038 (0.018)	0.071 (0.018)
ω	0.099 (0.028)	0.086 (0.029)	0.073 (0.027)	0.125 (0.047)
α_1	0.089 (0.016)	0.067 (0.014)	0.065 (0.014)	0.093 (0.020)
β_1	0.673 (0.067)	0.752 (0.063)	0.792 (0.057)	0.685 (0.089)
ν	4.123 (0.338)	4.748 (0.421)	5.486 (0.543)	6.029 (0.634)
Q	0.14 [0.10]	0.15 [8.44]	0.07 [62.64]	2.72 [310.22]
Q^2	3.80 [73.24]	7.10 [66.27]	6.26 [129.33]	4.66 [505.71]

Table II
Unconditional relative liquidity risk premium $\bar{\lambda}$ estimates at different VaR
time horizons.

The one-day and ten-day horizon estimates were obtained by scaling the half-hour estimates to the respective horizons. Table IV (appendix) contains the decomposition $\bar{\lambda}$ into mean and volatility components.

VaR horizon	$v = 5,000$	$v = 20,000$	$v = 40,000$
EVS portfolios			
10-min.	0.35	0.68	1.20
half-hour	0.17	0.34	0.61
one-day	0.03	0.06	0.12
ten-day	0.00	-0.01	0.02
DCX			
10-min.	0.30	0.64	1.21
half-hour	0.15	0.32	0.60
one-day	0.03	0.06	0.15
ten-day	0.00	0.00	0.05
DTE			
10-min.	0.28	0.51	0.86
half-hour	0.11	0.22	0.39
one-day	0.02	0.04	0.08
ten-day	-0.01	0.00	0.02
SAP			
10-min.	0.29	0.50	0.83
half-hour	0.12	0.23	0.38
one-day	0.02	0.05	0.11
ten-day	0.00	0.01	0.05

Table III
Estimation results 30-minute frequency.

See table I for explanations.

	frictionless	$v = 5,000$	$v = 20,000$	$v = 40,000$
EVS portfolios				
δ_0	1.2E-04 (9.3E-05)	-3.0E-06 (9.2E-05)	-3.0E-06 (9.0E-05)	-2.0E-06 (9.0E-05)
δ_1	0.052 (0.031)	0.065 (0.031)	0.067 (0.031)	0.081 (0.031)
ω	0.029 (0.019)	0.038 (0.024)	0.035 (0.021)	0.045 (0.026)
α_1	0.029 (0.013)	0.031 (0.014)	0.030 (0.013)	0.027 (0.013)
β_1	0.935 (0.031)	0.923 (0.037)	0.926 (0.033)	0.916 (0.039)
ν	9.386 (10.048)	20.167 (10.548)	16.230 (7.030)	13.687 (5.048)
Q	0.20 [3.46]	0.20 [4.45]	0.19 [4.87]	0.32 [7.69]
Q^2	0.26 [4.99]	0.66 [1.13]	0.66 [1.62]	0.81 [2.34]
DCX				
δ_0	2.9E-02 (1.9E-02)	-7.0E-06 (1.1E-04)	-6.0E-06 (1.1E-04)	-7.0E-06 (1.1E-04)
δ_1	0.029 (0.013)	0.012 (0.031)	0.021 (0.031)	0.061 (0.031)
ω	0.122 (0.071)	0.129 (0.073)	0.194 (0.141)	0.274 (0.185)
α_1	0.060 (0.024)	0.058 (0.023)	0.069 (0.031)	0.087 (0.037)
β_1	0.744 (0.123)	0.744 (0.121)	0.647 (0.220)	0.535 (0.271)
ν	6.590 (1.396)	7.058 (1.585)	7.536 (1.756)	8.666 (2.159)
Q	2.94 [3.14]	3.03 [3.13]	2.40 [1.86]	1.96 [3.70]
Q^2	0.60 [16.69]	0.09 [13.30]	0.12 [10.11]	0.14 [6.93]
DTE				
δ_0	2.3E-04 (1.5E-04)	-2.0E-06 (1.5E-04)	-1.0E-06 (1.5E-04)	3.0E-06 (1.5E-04)
δ_1	0.036 (0.031)	0.042 (0.031)	0.054 (0.031)	0.070 (0.031)
ω	0.051 (0.023)	0.053 (0.024)	0.054 (0.025)	0.058 (0.027)
α_1	0.029 (0.012)	0.030 (0.012)	0.029 (0.012)	0.028 (0.013)
β_1	0.881 (0.042)	0.877 (0.042)	0.876 (0.045)	0.872 (0.049)
ν	5.889 (1.156)	5.944 (1.171)	5.631 (1.072)	5.766 (1.107)
Q	0.09 [1.52]	0.04 [2.05]	0.23 [3.53]	1.36 [6.93]
Q^2	5.38 [5.42]	3.14 [2.30]	3.54 [1.66]	4.47 [1.57]
SAP				
δ_0	2.2E-04 (1.4E-04)	4.0E-06 (1.4E-04)	4.0E-06 (1.4E-04)	5.0E-06 (1.4E-04)
δ_1	0.056 (0.031)	0.058 (0.031)	0.081 (0.031)	0.107 (0.031)
ω	0.233 (0.133)	0.233 (0.123)	0.252 (0.116)	0.304 (0.163)
α_1	0.075 (0.034)	0.057 (0.028)	0.065 (0.029)	0.061 (0.031)
β_1	0.592 (0.200)	0.626 (0.175)	0.595 (0.161)	0.528 (0.223)
ν	7.677 (1.771)	8.524 (2.157)	8.826 (2.324)	8.448 (2.134)
Q	1.99 [3.39]	1.96 [5.32]	1.84 [10.47]	1.67 [20.75]
Q^2	2.16 [11.23]	0.97 [8.94]	0.39 [10.17]	0.46 [15.58]

Table IV
Decomposition of the unconditional relative liquidity risk premium $\bar{\lambda}$ at different VaR time horizons.

The table reports the three parts of $\bar{\lambda} = T^{-1} \sum_{t=1}^T \frac{\mu_{mb,t}}{t_{\alpha,\nu_2} \sigma_t(r_{mm,t})} + T^{-1} \sum_{t=1}^T \frac{t_{\alpha,\nu_1} \sigma_t(r_{mb,t})}{t_{\alpha,\nu_2} \sigma_t(r_{mm,t})} - 1$. Because of rounding errors, the sum of mean and volatility components might differ from the figures reported in Table II.

Table V			
VaR horizon	$v = 5,000$	$v = 20,000$	$v = 40,000$
EVS portfolios			
10-min.	0.36+0.99-1	0.70+0.98-1	1.16+1.04-1
half-hour	0.19+0.98-1	0.38+0.96-1	0.64+0.97-1
DCX			
10-min.	0.29+1.01-1	0.62+1.03-1	1.08+1.13-1
half-hour	0.17+0.99-1	0.35+0.98-1	0.60+1.00-1
DTE			
10-min.	0.27+1.01-1	0.49+1.02-1	0.77+1.08-1
half-hour	0.12+0.98-1	0.25+0.98-1	0.40+0.99-1
SAP			
10-min.	0.28+1.02-1	0.47+1.03-1	0.73+1.10-1
half-hour	0.13+0.99-1	0.23+1.00-1	0.37+1.02-1

Notes

¹Further general information about VaR techniques and regulation issues are available in Dowd (1998), Jorion (2000), Saunders (2000) or at the Bank of International Settlement website <http://www.bis.org>.

²A lively account of the rise and fall of Long-Term Capital Management is given in Lowenstein (2001).

³The unit price $a_t(v)$ of a buy of size v at time t can be computed analogously. Empirical studies focusing on the characterization of the price-volume relationship have been conducted by Beltran, Giot, and Grammig (2002) (stocks) and Coppejans, Domowitz, and Madhavan (2001) (stock index futures). The German Stock Exchange recently adopted the price impact as defined in Equation (1) as the key liquidity indicator for the automated auction system Xetra (see Gomber, P. and U. Schweickert (Deutsche Börse AG) (2002)).

⁴Bauwens and Giot (2001) provide a complete description of an order book market and Biais, Hillion, and Spatt (1999) describe the opening auction mechanism used in an order book market and corresponding trading strategies. Further information about the organization of the Xetra trading process is provided in Deutsche Börse AG (1999).

⁵Grammig, Schiereck, and Theissen (2001) have recently shown that the anonymity feature of automated auction systems can severely aggravate adverse selection effects.

⁶Note that during this period hidden orders were not allowed, and that trading hours shifted in the midst of the sample period.

⁷GAUSS programs for order book reconstruction are available from the authors upon request.

⁸Frictionless returns are, somewhat misleadingly, referred to as trading returns in the BDSS framework.

⁹See for example Gouriéroux, Le Fol, and Meyer (1998).

¹⁰As the trading hours shifted during the sample period, the diurnal functions ψ_t and ϕ_t are estimated separately for each sub-sample. When estimating the parameters of the autoregressive model components, we have to account for non-trading periods (overnight, weekends) and prevent that end-of-day observations shape the dynamics of the start-of-day returns. For this purpose, we adopt a procedure proposed by Engle and Russell (1998) and re-initialize the AR process at the start of each day. Sample average returns are used as initial values. We do not consider joint estimation which, although feasible, would impose a considerable computational burden. Recent research on related models of intra-day price processes has shown that the joint estimation results are quite similar to multi-step procedures like the one outlined above (see Engle and Russell (1998) or Martens, Chang, and Taylor (2002)).

¹¹For notational simplicity we suppress the dependence of the VaR measures on α as we do not vary the significance level in the empirical analysis.

¹²In a weakly efficient market, frictionless returns are serially uncorrelated. However, for ultra-high frequency time horizons, some small statistically significant autocorrelations can be expected (see Campbell, Lo, and MacKinlay (1997) or Engle (2000)).

¹³Diurnal variation in actual returns and volatility diurnality is taken into account employing nonparametric regression techniques, hence no estimation results are presented regarding the diurnal functions ψ_t and ϕ_t . Our GAUSS procedures library is available upon request. The MAXLIK or CML module is required.

¹⁴As above we take sample averages by time of day for each component and apply the Nadaraya-Watson smoother.

¹⁵This is, admittedly, a crude approximation. In a more precise strategy one could try derive temporal aggregation of GARCH processes, extending the results in Drost and Nijman (1993) and Drost and Werker (1996). However, GARCH aggregation properties have been derived for a specific class of GARCH processes (weak GARCH) and the extension to GARCH with diurnal volatility variation is beyond the scope of the present paper.

¹⁶Note that this analysis should be interpreted with caution as we rely on two hypotheses that are appealing but imperfect in practice: (1) the square root of time rule is perfectly valid; (2) volatility occurs when markets are open, i.e. we do not consider the volatility of the overnight returns. By increasing or decreasing f , it is nevertheless possible to mitigate these effects and get a range of possible $\bar{\lambda}$ (akin to a scenario analysis).