Journal of Empirical Finance 16 (2009) 777–792

Contents lists available at ScienceDirect

Journal of Empirical Finance

journal homepage: www.elsevier.com/locate/jempfin

Intraday Value at Risk (IVaR) using tick-by-tick data with application to the Toronto Stock Exchange[☆]

Georges Dionne[a,b,*], Pierre Duchesne[b,c], Maria Pacurar[b,d]

[a] Canada Research Chair in Risk Management, HEC Montréal, 3000 Côte-Sainte-Catherine, Montreal, Canada H3T 2A7
[b] Centre for Research on E-Finance (CREF), Montreal, Canada
[c] Département de mathématiques et statistique, Université de Montréal, Montreal, Canada
[d] School of Business Administration, Dalhousie University, Halifax, Canada

ARTICLE INFO

Article history:
Received 25 November 2008
Accepted 23 May 2009
Available online 6 June 2009

JEL classification:
C22
C41
C53
G15

Keywords:
Intraday Value at Risk (IVaR)
Tick-by-tick data
ACD model
Intraday market risk
Market microstructure

ABSTRACT

This paper investigates the use of tick-by-tick data for intraday market risk measurement. We propose a method to compute an Intraday Value at Risk based on irregularly spaced high-frequency data and an intraday Monte Carlo simulation. A log-ACD–ARMA–EGARCH model is used to specify the joint density of the marked point process of durations and high-frequency returns. We apply our methodology to transaction data for three stocks actively traded on the Toronto Stock Exchange. Compared to traditional techniques applied to intraday data, our methodology has two main advantages. First, our risk measure has a higher informational content as it takes into account all observations. On the total risk measure, our method allows for distinguishing the effect of random trade durations from the effect of random returns, and for analyzing the interaction between these factors. Thus, we find that the information contained in the time between transactions is relevant to risk analysis, which is consistent with predictions from asymmetric-information models in the market microstructure literature. Second, once the model has been estimated, the IVaR can be computed by any trader for any time horizon based on the same information and with no need of sampling the data and estimating the model again when the horizon changes. Backtesting results show that our approach constitutes reliable means of measuring intraday risk for traders who are very active in the market.

1. Introduction

The traditional way of measuring and managing risk has been challenged by the current trading environment. Over the last several years, the speed of trading has been constantly increasing. Day-trading, once the exclusive territory of floor-traders, is now available to all investors. "High-frequency finance hedge funds" have emerged as a new and successful category of hedge funds. Consequently, risk management is now obliged to keep pace with the market. For day traders, market makers or other very active

☆ We are very grateful to Christian Gouriéroux and Denitsa Stefanova for valuable remarks and suggestions. The authors thank the two anonymous referees for constructive criticism. We have also benefited from remarks and comments through correspondence from Patrick Burns, Peter Christoffersen, Pierre Giot, Joachim Grammig, Erick Rengifo, and David Veredas. We also thank Iraj Fooladi, Geneviève Gauthier, Alain Guay, Khemais Hammami, Cristian Popescu, Jean-Guy Simonato, Joshua Slive, and participants at the 2006 Financial Management Association Annual Meeting, the 2006 Northern Finance Association Conference, the 2006 Econometric Society European Meeting, the 2006 European Financial Management Association Conference, the Bank of France Microstructure of Financial and Money Markets Conference, the 2005 International Conference on New Financial Market Structures, and seminar participants at Bank of Canada, Concordia University, Dalhousie University, HEC Montréal, in particular, Greg Bauer, Umberto Cherubini, Chris D'Souza, David Goldreich, Ingrid Lo, Albert Menkveld, Ludovic Mercier, Stephen Sapp, Kashi Nath Tiwari for helpful comments on a previous version of the paper. Jean-David Fournier, Francois Guertin, and Mohamed Jabir provided useful research assistance. Maria Pacurar acknowledges financial support from the *Institut de Finance Mathématique de Montréal (IFM2)*, the *Centre for Research on E-Finance (CREF)*, the *Canada Research Chair in Risk Management* and the *School of Business Administration*, Dalhousie University.
* Corresponding author. Canada Research Chair in Risk Management, HEC Montréal, 3000 Côte-Sainte-Catherine, Montreal, Canada H3T 2A7. Tel.: +1 514 340 6596; fax: +1 514 340 5019.
E-mail addresses: georges.dionne@hec.ca (G. Dionne), duchesne@dms.umontreal.ca (P. Duchesne), maria.pacurar@dal.ca (M. Pacurar).

0927-5398/$ – see front matter
doi:10.1016/j.jempfin.2009.05.005

agents on the market, risk should be evaluated on shorter-than-daily time intervals since the horizon of their investments is generally less than a day. For example, day traders liquidate any open positions at closing, in order to preempt any adverse overnight moves resulting in large gap openings. Brokers must also be able to calculate trading limits as fast as clients place their orders. Significant intraday variations in asset prices affect the margins a client has to deposit with a clearing firm, and this should be taken into account in the design of any margining engine. Moreover, as noted by Gouriéroux and Jasiak (1997), banks also use intraday risk analysis for internal control of their trading desk. For example, a trader could be asked at 11 a.m. to provide an estimation of his risk for the rest of the day or the next hour.

Much effort has been spent on developing increasingly sophisticated risk models of Value at Risk (VaR) type for daily data and/or longer horizons,[1] but the issue of intraday market risk measurement has been less explored. With increased access to intraday financial databases and advanced computing power, it became possible to address the question of how to define practical risk measures for investors or market makers operating on an intraday basis.

This paper proposes a method of computing an intraday VaR (IVaR) using tick-by-tick data and an approach that takes into account the irregular timing of trades. VaR refers to the maximum expected loss that will not be exceeded under normal market conditions over a predetermined period at a given confidence level (Jorion, 2000, p.xxii). Intraday VaR was initially discussed by Giot (2005) and it results from applying VaR techniques to intraday returns.

The irregular feature of intraday data complicates assessment of the results obtained from models based on such data. One possibility for computing intraday VaR consists of aggregating the irregularly time-spaced high-frequency data, in order to retrieve a regular sample to which traditional methods are applied. Intraday VaR models based on equidistantly time-spaced returns sampled from irregularly time-spaced data have been recently developed by Giot (2005) and Giot and Grammig (2006). Giot (2005) estimates a conditional parametric VaR in which intraday volatility is modeled using normal GARCH, Student GARCH, and RiskMetrics models on 15- and 30-minute returns of three stocks traded on the New York Stock Exchange. The Student GARCH model is found to perform best among all models. A different approach is followed by Giot and Grammig (2006) who introduce the notion of an Actual VaR that takes into account the potential price impact of liquidating a portfolio. Based on snapshots of the reconstructed real time order book from the Xetra platform, taken at 10- and 30-minute frequencies, Giot and Grammig compute returns that depend on the volume to be traded. Midquote returns are also used as input for computing a standard market VaR that is referred to as a frictionless VaR. The comparison of the two VaRs allows the authors to quantify the relative liquidity risk premium within a day. This approach requires a full knowledge of the order book and is suited mostly for automated auction markets able to provide accurate data.

The above-mentioned approaches not only require finding the optimal aggregating scheme, but also inevitably lead to the loss of important information contained in the time intervals between transactions. We propose an alternative approach which is more closely related to the asymmetric-information models from the market microstructure literature and which deals directly with irregularly spaced data by using models adapted to the particularities of such data. Except for Giot (2005), no other study has been published, to our knowledge, on the performance of intraday VaR-type risk measures based on models for irregularly time-spaced intraday data. Giot (2005) evaluates a conditional parametric intraday VaR in which the volatility is computed using log-ACD models applied to price durations.[2] This approach thus leads to irregularly time-spaced VaRs corresponding to subsequent price durations. These VaRs are then averaged for regularly time-spaced intervals for backtesting purposes. The method requires the user to fix exogenously the threshold for defining price durations prior to any estimation and to re-do the procedure every time this threshold changes. The backtesting results show that the performance of the log-ACD model applied to price durations is not satisfactory for forecasting the intraday VaR in a regularly time-spaced framework.

Our method for estimating and evaluating the IVaR is based on a GARCH model for tick-by-tick returns in the spirit of Engle (2000), coupled with the use of an intraday Monte Carlo simulation approach.[3] The advantage of this model is that it explicitly accounts for the irregular spacing of the data by considering the times between trades (defined as trade durations) when modeling returns. Our approach differs from Giot's (2005) on two points: (i) we use all available information as we focus on trade durations and not price durations; (ii) our use of the Monte Carlo method produces regularly spaced VaRs in a convenient way and avoids the need to use complicated time transformations to switch to the regularly time-spaced world.

Compared to traditional VaRs computed for intraday data, the IVaR has two advantages. First, our risk measure has a higher informational content as it is based on all available observations. The joint modeling of durations and returns allows for distinguishing the impact of random trade durations on the risk measure from the effect of random returns, and for analyzing the interaction between these factors. Our empirical study shows that ignoring the uncertainty about trade durations and the effect of the dependence between trade durations and returns may lead to underestimation of risk, especially in the first part of the trading day when the information asymmetry is greater. Second, once the model has been estimated, the IVaR can be computed for any time interval. This is different from standard methods based on regularly sampled data that involve a new dataset and estimation of a model when the risk horizon changes. As discussed later, this has important practical implications for traders. Backtesting results show that the model we adopt performs well out-of-sample for almost all the time horizons and the confidence levels considered, at least in our application.

[1] This is motivated by the fact that financial institutions generally produce their market VaR at the end of the business day to measure their total risk exposure over the next day. For regulated capital adequacy purposes, banks usually compute the market VaR daily and then re-scale it to a 10-day horizon.

[2] As first noted by Engle and Russell (1998), price durations (the minimum amount of time needed for the price of an asset to change by a certain amount) are closely linked to the instantaneous volatility of the price process.

[3] See Christoffersen (2003) and Burns (2002) for computing VaR using GARCH models and a Monte Carlo simulation with daily data.

The rest of the paper is organized as follows. In Section 2, we define the concept of IVaR. In Section 3, we outline the econometric model used to describe the dynamics of durations and tick-by-tick returns. In Section 4, we present the intraday Monte Carlo simulation approach used to estimate the IVaR. In Section 5, we apply our methodology to transaction data on three Canadian stocks traded on the Toronto Stock Exchange (TSX): the Royal Bank of Canada (RY) stock, the Placer Dome (PDG) stock and the Alcan (AL) stock. We investigate the model's performance by backtesting and we economically compare the IVaR to different intraday VaR measures. In Section 6, we discuss the gains from implementing the IVaR. Finally, Section 7 concludes the paper.

2. IVaR: definition

Asymmetric-information models from the market microstructure literature suggest that both time between trades and price dynamics are related to the existence of new information on the market and that time is not exogenous to the price process (see Diamond and Verrechia, 1987; Admati and Pfleiderer, 1988; Easley and O'Hara, 1992, among others). Predictions from theoretical models with regard to the relationship between intertrade durations and volatility are relatively ambiguous.[4] In the end, the common viewpoint is that durations and volatility are closely linked. Consequently, to estimate an IVaR that draws on insights from this microstructure literature, we employ a joint modeling of arrival times and price changes as described in the next section.

Consider two consecutive trades that happened at times t_{i-1} and t_i, at prices p_{i-1} and p_i, respectively. Let $x_i = t_i - t_{i-1}$ be the duration between these trades, and $r_i = \log(p_i/p_{i-1})$ be the corresponding continuously compounded return.

To define our IVaR, let $Y = \{y_k, k \in \mathbb{Z}\}$ be the process of the asset return sampled at regular time intervals equal to T units of time. We consider a realization of length n' of the process Y, $\{y_k; k = 1\ldots,n'\}$, with y_k obtained at times t'_k such that $t'_k - t'_{k-1} = T$. Thus, y_k denotes the T-period return which is simply the sum of correspondent tick-by-tick returns

$$y_k = \sum_{i=\tau(k-1)}^{\tau(k)-1} r_i, \tag{1}$$

where $\tau(k)$, $k = 1,\ldots n'$ is such that the cumulative duration exceeds T for the first time. This means that

$$\sum_{i=\tau(k-1)}^{\tau(k)-1} x_i \leq T \text{ and } \sum_{i=\tau(k-1)}^{\tau(k)} x_i > T.$$

By definition, $\tau(0) = 1$. Therefore, if we model specifically the durations x_i, we are able to determine every $\tau(k)$ and keep track of the time step.

The IVaR with confidence level $1-\alpha$, is formally defined as

$$Pr(y_k < -\text{IVaR}_k(\alpha) \,|\, \mathcal{G}_k) = \alpha. \tag{2}$$

where $\mathcal{G}_k$ denotes the information set that includes all the tick-by-tick data (durations and returns) up to time t'_{k-1}. Common confidence levels for market risk shall be considered in the empirical study. It follows from Eq. (2) that the IVaR is such that $-\text{IVaR}_k(\alpha) = Q_k(\alpha|\mathcal{G}_k)$ where $Q_k(\alpha|\mathcal{G}_k)$ is the quantile of the conditional distribution of y_k.

Various methods exist for computing a VaR: parametric or variance–covariance methods (e.g., RiskMetrics, GARCH), nonparametric (e.g., Historical Simulation), and semi-parametric (e.g., CAViar, Extreme Value Theory). We adopt a GARCH-type model to specify the time-varying volatility of the tick-by-tick returns r_i and a Monte Carlo simulation approach to simulate the future regularly spaced returns y_k from which the IVaR is computed.

3. A log-ACD–ARMA–EGARCH model

Since trade durations and tick-by-tick returns could be interdependent, as suggested by the microstructure theory, we adopt a model for their joint distribution. Following the literature (Engle, 2000; Russell and Engle, 2005), let $f(x_i, r_i | \breve{x}_{i-1}, \breve{r}_{i-1}; \theta)$ represent the joint density of durations and returns, where θ is a d-vector of parameters ($d < n$), finite and invariant over events, and $\breve{x}_{i-1}$ and $\breve{r}_{i-1}$ denote the past of the variables X and R, respectively, up to the $(i-1)$th transaction. The log-likelihood function for a sample of observations $\{x_i, r_i\}$ with $i = 1,\ldots,n$ can then be written as

$$\mathcal{L}(\theta_1, \theta_2) = \sum_{i=1}^{n} [\log g(x_i | \breve{x}_{i-1}, \breve{r}_{i-1}; \theta_1) + \log q(r_i | x_i, \breve{x}_{i-1}, \breve{r}_{i-1}; \theta_2)], \tag{3}$$

[4] In Easley and O'Hara (1992), informed traders trade only when there are information events that influence the asset price. Therefore, in this model, long intertrade durations are associated with no news and, consequently, with low volatility. Opposite relations between duration and volatility follow from the Admati-and-Pfleiderer model (1988) where frequent trading is associated with liquidity traders. Hence, low trading means a high proportion of informed traders on the market, which translates into higher volatility.

where $g(x_i|\breve{x}_{i-1},\breve{r}_{i-1};\theta_1)$ is the marginal density of the duration x_i with parameter θ_1, conditional on past durations and returns, and $q(r_i|x_i,\breve{x}_{i-1},\breve{r}_{i-1};\theta_2)$ is the conditional density of the return r_i with parameter θ_2, and conditional on past durations and returns as well as the contemporaneous duration x_i.

Several possibilities for g and q can be considered in the log-likelihood function. In the next subsections, we use the log-ACD model of Bauwens and Giot (2000) for describing the dynamics of trade durations and a general ARMA–EGARCH model for the price dynamics. The two parts of the log-likelihood function are estimated simultaneously to avoid any potential loss of efficiency.

3.1. Log-ACD models for durations

Let $\psi_i = E(x_i|F_{i-1})$ be the conditional duration, supposed to be non-negative and measurable with respect to F_{i-1}, the information set available at time t_{i-1}. Engle and Russell (1998) introduced the ACD model to describe the duration clustering observed in high-frequency financial data. The ACD model is based on the assumption that all the temporal dependence between durations is captured by the conditional duration mean function, such that

$$x_i / \psi_i = \varepsilon_i. \tag{4}$$

The process $\varepsilon = \{\varepsilon_i, i \in \mathbb{Z}\}$ is composed of independent and identically distributed (i.i.d.) random variables. We assume ε_i to be non-negative with probability one and admit a density $p(\varepsilon)$, such that $E(\varepsilon_i) = 1$.

By choosing different specifications for the conditional durations and the density $p(\varepsilon)$, several forms of the ACD model can be obtained (see Bauwens and Giot, 2001; Hautsch, 2004; Pacurar, 2008 for surveys). Here, in order to avoid the necessity of imposing non-negativity constraints on the parameters of the conditional mean function, we use the logarithmic version of the ACD model, the so-called log-ACD model, proposed originally by Bauwens and Giot (2000). This model specifies the autoregressive equation on the logarithm of ψ_i:

$$\psi_i = exp\left(\omega + \sum_{j=1}^{m} \alpha_j \varepsilon_{i-j} + \sum_{j=1}^{q} \beta_j ln\, \psi_{i-j}\right). \tag{5}$$

We use the generalized gamma distribution with unit expectation for the density $p(\varepsilon)$, i.e.

$$p(\varepsilon|\gamma_1,\gamma_2) = \begin{cases} \dfrac{\gamma_1 \varepsilon^{\gamma_1\gamma_2 - 1}}{\gamma_3^{\gamma_1\gamma_2}\Gamma(\gamma_2)} exp\left[-\left(\dfrac{\varepsilon}{\gamma_3}\right)^{\gamma_1}\right], & \varepsilon > 0, \\ 0, \text{ elsewhere.} \end{cases}$$

where $\gamma_1 > 0, \gamma_2 > 0$, $\Gamma(\cdot)$ denotes the gamma function and $\gamma_3 = \Gamma(\gamma_2)/\Gamma\left(\gamma_2 + \frac{1}{\gamma_1}\right)$. Note that the generalized gamma distribution nests the Weibull distribution when $\gamma_2 = 1$.

3.2. EGARCH models for returns

As explained in Section 5, we use transaction prices instead of mid-quotes for computing volatility of returns in an order-driven market. It is well known that transaction prices may be affected by various market microstructure effects, such as nonsynchronous trading and the bid–ask bounce (Tsay, 2002), that may induce serial correlation in high-frequency returns. Therefore, to capture such microstructure effects, we follow Grammig and Wellner (2002) and Ghysels and Jasiak (1998) and use ARMA(p,q) models on the tick-by-tick returns. In the next sections, the simple ARMA(1,1) specification is used:

$$r_i = c + \varphi_1 r_{i-1} + e_i + \theta_1 e_{i-1}. \tag{6}$$

The error term of Eq. (6) satisfies

$$e_i = z_i \sqrt{h_i}, \tag{7}$$

where $z = \{z_i, i \in \mathbb{Z}\}$ denotes a strong white noise, that is, the stochastic process z is composed of i.i.d. random variables, such that z_i has mean zero and a finite variance. The random variable z_i is such that $E(z_i^2) = 1$. The conditional variance per transaction is given by h_i, that is $h_i = V_{i-1}(r_i|x_i,\breve{x}_{i-1},\breve{r}_{i-1})$. We consider the following specification for the volatility of returns:

$$log(h_i) = \gamma log(x_i) + \tilde{\omega} + \sum_{j=1}^{P} \tilde{\beta}_j \left\{ log\left(h_{i-j}\right) - \gamma log\left(x_{i-j}\right) \right\} + \sum_{j=1}^{Q} \left[a_j \left\{ |z_{i-j}| - E\left(|z_{i-j}|\right) \right\} + \tilde{\alpha}_j z_{i-j} \right]. \tag{8}$$

The unknown parameters in the model composed of Eqs. (6), (7) and (8), are c, φ_1, θ_1, $\tilde{\omega}$, $\tilde{\beta}_j$, $j=1,\ldots,P$, a_j and $\tilde{\alpha}_j$, $j=1,\ldots,Q$, and finally, γ. The astute reader will notice the similarity between Eq. (8) and Engle's UHF-GARCH model (2000) that specifies a GARCH component for the volatility of returns per unit of time, that is h_i/x_i. Under Engle's framework, the conditional variance of returns from one transaction to the next equals the duration between the two consecutive trades times the variance per second (assuming durations are measured in seconds, as is often the case). Although it seems natural to model the variance as a function of time when using irregularly time-spaced data (see also Meddahi et al., 2006), this modeling for the unit of time might be quite restrictive for some empirical data. The conditional heteroskedasticity in the returns could depend on time in a more complicated way due to the fact that the impact of the trading frequency on volatility following news events depends on the trading behavior of different types of investors, some more sophisticated than others. Therefore, in our model, the parameter γ that specifies the duration weighting for the volatility of a particular stock has to be estimated for each stock.[5] When $\gamma=1$, the model is similar, though not equivalent, to the model studied in Engle (2000). Note that we specify the mean Eq. (6) on the returns r_i rather than on the returns per square root of time as in Engle (2000), in order to simplify the simulation of the future distribution of returns from which the IVaR could be extracted. When $\gamma=0$, Eq. (8) becomes a standard EGARCH applied to the high-frequency data by ignoring the irregular spacing of returns when modeling volatility. As shown later in the empirical analysis, the parameter γ is statistically significant, and statistically different from one, and has different values for different stocks, at least in our empirical application.

The use of an EGARCH(p,q) model in the empirical study is justified by the advantage of keeping the volatility component positive during simulations regardless of the variables added to the autoregressive equation. From a computational point of view, the absence of non-negativity constraints on the parameters also simplifies numerical optimization (see Nelson, 1991 for the properties of EGARCH models).

In the next section, we describe the intraday Monte Carlo approach that uses the combined log-ACD–ARMA–EGARCH model to simulate durations and returns from which the IVaR will be computed. While our method of computing the IVaR is not restricted to the specific models described in this section (for example, a user could specify an ACD model for the durations, and an ARMA–GARCH model for the returns), the empirical application provided in Section 5 suggests that the model presented in this section is suited for our stock data.

4. Intraday Monte Carlo simulation

In this section, we describe the Monte Carlo simulation approach used for computing the IVaR based on tick-by-tick data. As the time intervals between consecutive observations are irregular, the question of how to assess the results needs to be considered. Since our model fully specifies the distribution of returns in event time, one-step forecasts can be computed analytically. However, we need to run simulations in order to obtain forecasts of returns in any arbitrary length of calendar time, T. More specifically, the log-ACD model will define the time step of our simulations and the ARMA–EGARCH model introduced in the previous subsection will generate the corresponding tick-by-tick returns. We compute the forecasted returns for regular calendar-time intervals as the sum of irregular (tick-by-tick) intraday returns. Using the simulated distribution of returns over the chosen time interval T, we calculate the IVaR by extracting the desired percentile. The IVaR obtained is thus an IVaR for regular time intervals (therefore, conventionally comparable to regular real returns), but computed using tick-by-tick data and adapted to the non-regularity of time intervals.

More precisely, we proceed as follows:

(1) The original sample is divided into two parts, one for estimation and one for forecast/validation. The estimation sub-sample serves to calibrate a log-ACD–ARMA–EGARCH model for durations and tick-by-tick returns as presented in the previous section.
(2) We draw random numbers from the standard generalized gamma distribution and the standard normal distribution, respectively, to compute the innovations which will serve to generate scenarios for future durations and returns.
(3) In order to obtain the simulated durations within the first fixed interval, we forecast the durations in an iterative way using Eqs. (4) and (5) as long as the sum of forecasted tick-by-tick durations does not exceed the time interval of interest.
(4) Using Eqs. (6), (7), and (8), we forecast, conditional on the simulated durations, the tick-by-tick returns corresponding to the first regular interval. By summing up all these returns we obtain the (regular) return for the first interval.
(5) We continue the procedure until the desired number of fixed intervals or regular returns is obtained. This corresponds to the first path.
(6) We repeat steps (1) to (5) for the desired number of paths.
(7) For each time-fixed interval, the IVaR corresponds to the quantile of the simulated distribution of returns for that interval.

Fig. 1 summarizes the main steps of the simulation. As can be seen, the algorithm displays similarities with the traditional Monte Carlo simulation used for computing a multi-period VaR. The main difference is at step (3) where we make use of the ACD model to keep trace of the time step. While in the standard Monte Carlo approach the time unit is not very important because all observations are equally spaced, here we relax this constraint and are able to proceed tick-by-tick. It should be noted that we need to fix a time horizon for backtesting purposes. Fixing the time horizon will also allow us to compare the IVaR to intraday VaRs

[5] Note that, in our model, the parameter γ serves mainly for better fitting of our data as illustrated by backtesting tests. A detailed study of the relationship between the arrival of news and the trading intensity and volatility of a particular stock is beyond the scope of our paper and left for future research.

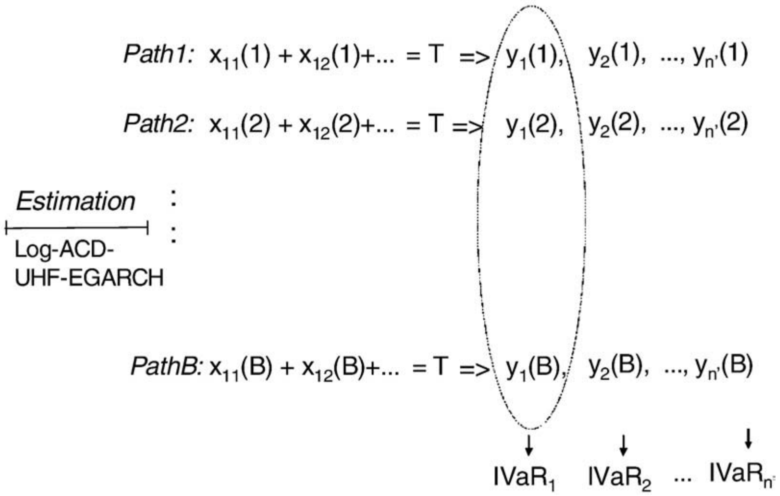

Fig. 1. Illustration of the intraday Monte Carlo simulation approach. B is the number of independently simulated paths. T is the length of the chosen time interval. n' is the number of regular intervals used for validation. $x_{ki}(p)$ is the i trade duration corresponding to the path p for the interval k. $y_k(p)$ is the regular return corresponding to the interval k for the path p. $p=1,...,B$, $k=1,...,n'$, $i=1$, ..., t such that $\Sigma_{i=1}^{t} x_{ki}(p) \leq T$ and $\Sigma_{i=1}^{t+1} x_{ki}(p) > T$. $IVaR_k$ represents the quantile of the simulated distribution of returns for interval k, $y_k(p)$ with $p=1,...,$ B. After each simulated interval k, the index i counting the simulated tick-by-tick durations and returns is reinitialized to 1.

computed using traditional methods on equidistantly spaced data. However, the IVaR may also be computed using the parametric method coupled with the log-ACD–ARMA–EGARCH model in the standard way. In this case, the IVaR obtained would be irregularly spaced and would describe the maximum potential loss at a given confidence level over the time until the next trade occurs. We illustrate our procedure using real data in the next section.

5. Empirical analysis

5.1. Data

In this section we compute the IVaR according to the methodology previously described using high-frequency data from the "Market Data Equity Trades and Quotes Files" CD-ROM of the TSX. Previous studies on irregularly spaced data used data from the NYSE (e.g., Engle and Russell, 1998; Giot, 2005), or from the Paris Bourse (e.g., Jasiak, 1998; Gouriéroux et al., 1999), the German Stock Exchange (e.g., Grammig and Wellner, 2002), and the Moscow Interbank Currency Exchange (Anatolyev and Shakin, 2007). TSX is essentially an electronic order-driven market with a structure similar to that of other exchanges. Therefore, our results should not be specific to TSX only. We focus on three actively traded stocks: Royal Bank of Canada stock (RY), Placer Dome stock (PDG) and Alcan stock (AL) for the period from April 1st to June 30, 2001 which contains 63 trading days. The Royal Bank is Canada's largest bank as measured by market capitalization (about C$ 63.8 billion in 2006) and assets. Placer Dome was Canada's second largest gold miner with a market capitalization of US$ 8.2 billion at the end of 2004. Alcan was one of the world's leading suppliers of bauxite, alumina and aluminium and had a market capitalization of around C$ 39 billion in 2007.[6]

The TSX operates as an automated, continuous auction market. Its structure and its trading process are discussed in Davies (2003) and Smith et al. (2006), among others. The TSX intraday database contains date-and-time stamped bid-and-ask quotes, transaction prices, and volume for all firms. Following the literature, we removed all non-valid trades and interdaily durations and returns and retained only those transactions made during regular trading hours. Open trades were also deleted in order to avoid effects induced by the opening auction. For simultaneous transactions, we considered a weighted average price and removed all remaining observations with this time stamp, thus considering these observations as split transactions.

Previous studies dealing with high-frequency data commonly used mid-quotes instead of transaction prices (e.g., Engle, 2000; Engle and Lange, 2001; Manganelli, 2005). While this may be appropriate for the NYSE which is a quote-driven market, working with transaction prices appears a better choice in our case since we are interested in VaR estimation and want to forecast returns for real transactions. To liquidate positions in an order-driven market, one has to transact either on the ask or the bid and, therefore, using the midquote-change quantiles may understate the true VaR.[7]

We also noticed what appeared to be an anomaly on 4 May before 3 p.m. when trade durations for the RY and the AL stocks were excessively long (37 and 38 min, respectively). As documented in the press,[8] the TSX actually had a technical problem with its data-

[6] Placer Dome was acquired by Barrick Gold in 2006. Following its acquisition by Rio Tinto, Alcan shares have been delisted from the TSX starting on 1 November 2007.

[7] We are grateful to Joachim Grammig for shedding light on this issue.

[8] Futures World News Financial via Comtex, May 4, 2001.

Table 1
Descriptive statistics of raw data for Royal Bank (RY), Placer Dome (PDG), and Alcan (AL) stocks.

	Mean	Std. dev	Skew	Kurt	Max	Min
Durations						
RY	28.44	35.19	2.98	17.54	526	1
PDG	52.59	82.11	5.04	59.87	2248	1
AL	33.39	50.95	3.79	25.73	776	1
Returns						
RY	0.000	0.001	−0.021	7.798	0.008	−0.008
PDG	0.000	0.002	0.010	7.822	0.017	−0.016
AL	0.000	0.001	−0.389	21.124	0.015	−0.022

The sample period runs from April 1st to June 30, 2001. It consists of 51,661 observations for the RY stock, 27,965 observations for the PDG stock and 43,971 observations for the AL stock. *Mean* is the sample mean, *Std. dev* is the sample standard deviation, *Skew* is the sample skewness, *Kurt* is the sample kurtosis, *Max* is the sample maximum, *Min* is the sample minimum.

feed updates on that day and, consequently, halted trading for about 50 min. We removed these data, since this event had a strong impact on the descriptive statistics for the RY and AL stocks. This leaves us with a total of 51,661 observations for the RY stock, 27,965 observations for the PDG stock, and 43,971 observations for the AL stock.

The price increment for the three stocks RY, PDG and AL for the period analyzed was one penny (C$0.01).[9] We observe a disproportionately large number of zero returns (almost 50%) which seems rather typical for high-frequency data, especially for single stocks. For the IBM dataset used by Engle and Russell (1998), Tsay (2002, p. 182) reported that about two-thirds of the intraday transactions were without price change. Gorski et al. (2002) report the same phenomenon for DAX[10] returns and call it the *zero return enhancement*. Bertram (2005) also finds a large number of zero price changes for equity data from the Australian Stock Exchange and argues that zero returns follow from the absence of significant new information on the market. According to the efficient market hypothesis, a price changes when new information arrives on the market and, consequently, traders simply continue to trade at the previous price when the amount of information is insufficient to move the price. So zero returns yield information.[11]

Information about the raw data is given in Table 1. Of the three stocks, RY, with an average duration of almost 29 s, is traded almost twice as frequently as PDG, while AL is traded on average every 33 s. We also note overdispersion of durations, i.e., the standard deviation is higher than the mean. This is typically found in the literature for the trade duration process and it may suggest that exponential distribution cannot properly describe the durations. The transaction returns have a sample mean equal to zero for all stocks and a standard deviation equal to 0.001, 0.002, and 0.001, respectively. RY and AL returns exhibit negative sample skewness while PDG returns have a positive skewness. All stocks' returns display a kurtosis higher than that of a normal distribution.

5.2. Seasonal adjustment

As noted by Engle and Russell (1998), high-frequency data exhibits a strong intraday seasonality explained by the fact that markets tend to be more active at the beginning and towards the end of the day. While most of the studies on high-frequency data ignore interday variations in variables, Anatolyev and Shakin (2007) found that durations and return volatilities of the considered Russian stocks fluctuate throughout different trading days. To prevent distortion of the results, these interday and intraday seasonalities must be taken out prior to estimating any model. We inspected our data for such evidence. We noted, for example, that durations are higher on Mondays and Fridays than during the rest of the week. A Wald test in a regression of average RY durations on five day-of-the-week dummies rejected the null hypothesis of equality of all coefficients, thus providing evidence of interday seasonal effects. Interday effects were also identified for PDG durations. When found, we removed interday seasonality under a multiplicative form, following the approach taken by Anatolyev and Shakin (2007):

$$x_{t,\text{inter}} = \frac{x_t}{\bar{x}_s}, \tag{9}$$

$$r_{t,\text{inter}} = \frac{r_t}{\sqrt{\overline{r^2}_s}}, \tag{10}$$

where $\bar{x}_s$ corresponds to the average duration for day s if observation t belongs to day s, and $\overline{r_s^2}$ is the average of squared returns for day s.

[9] Beginning January 29, 2001 the TSX introduced the penny tick size for stocks selling at over $0.50.
[10] Deutsche Aktienindex (DAX) represents the index for the 30 largest German companies quoted on the Frankfurt Stock Exchange.
[11] We keep the zero returns in our sample. See Darolles et al. (2000) and Russell and Engle (2005) for models that specifically accommodate the high proportion of zero returns.

To take out the time-of-day effect, Engle and Russell (1998) suggest computing "diurnally adjusted" durations by dividing the raw durations by a seasonal deterministic factor related to the time at which the duration was recorded. We obtained intraday seasonally adjusted (isa) durations and returns in the following way:

$$x_{t,\mathrm{intra}} = \frac{x_{t,\mathrm{inter}}}{E\left(x_{t,\mathrm{inter}} \,|\, \mathcal{F}_{t-1}\right)},$$

$$r_{t,\mathrm{intra}} = \frac{r_{t,\mathrm{inter}}}{\sqrt{E\left(r^2_{t,inter} \,|\, \mathcal{F}_{t-1}\right)}},$$

where the expectation was computed by averaging the variables over thirty-minute intervals for each day of the week and then using cubic splines on the thirty-minute intervals to smooth the seasonal factor. Thus, intraday patterns are different for different days of the week. Fig. 2A and B shows, respectively, the estimated intraday factor for durations and squared returns for the RY stock. Similar seasonal-factor patterns are found for the other stocks and, therefore, are not reported. The patterns are analogous to what has been found in previous studies. Durations are shorter at the beginning and at the end of the day and longer in the middle. The

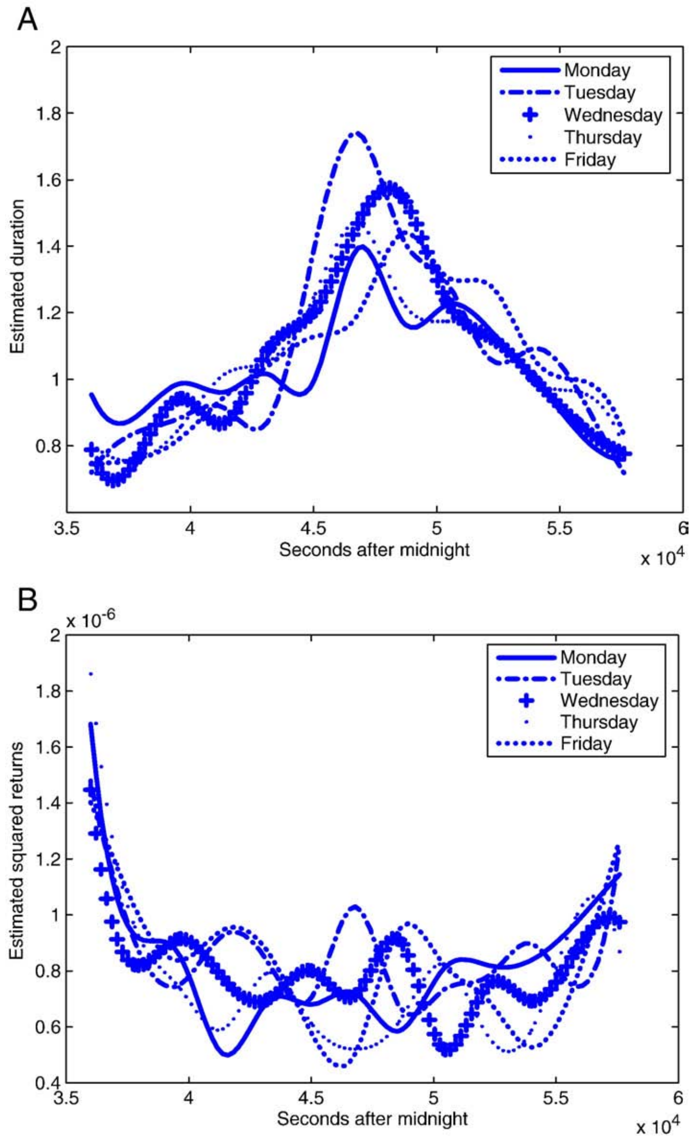

Fig. 2. A) Estimated intraday factors for RY durations. B) Estimated intraday factors for RY squared returns.

Table 2
Descriptive statistics of deseasonalized data for Royal Bank (RY), Placer Dome (PDG), and Alcan (AL) stocks.

	Mean	Std. dev	Skew	Kurt	Max	Min	$Q(15)$	$Q_2(15)$
Durations								
RY	0.99	1.14	2.55	12.73	15.97	0.01	1585	480
PDG	1.03	1.59	9.22	289.14	81.47	0.01	2547	68
AL	1.00	1.48	3.59	22.81	21.39	0.01	5628	1477
Returns								
RY	0.001	0.983	−0.001	7.310	6.889	−8.391	8155	14,690
PDG	0.002	0.991	0.025	6.752	7.860	−6.907	2425	3112
AL	0.007	0.984	−0.200	11.521	9.195	−19.795	4323	3432

Mean is the sample mean, *Std. dev* is the sample standard deviation, *Skew* is the sample skewness, *Kurt* is the sample kurtosis, *Max* is the sample maximum, *Min* is the sample minimum. $Q(15)$ is the Ljung–Box test statistics with 15 lags, $Q_2(15)$ is the Ljung–Box test statistics applied to squared variables using 15 lags. The associated 95% critical value is 24.99.

return volatility is lower in the middle of the day than at the beginning and at the end. This reflects the behavior of traders who are very active at the beginning of the trading session and adjust their positions to incorporate the overnight change in information. Towards the end of the day, traders are changing their positions in anticipation of the close and to preempt the risk posed by any information that could arrive during the night. We also notice a difference between the patterns for different days.

Descriptive statistics for deseasonalized data are given in Table 2. We note the salient features of high-frequency data, such as overdispersion in trade durations and high autocorrelations. The Ljung–Box test statistics for fifteen lags $Q(15)$ tests the null hypothesis that the first 15 autocorrelations are zero. The large values of these statistics greatly exceed the 5% critical value of 24.99, indicating strong autocorrelation of both durations and returns. There is still excess kurtosis for all stocks even if at a lesser degree compared to the raw data. We chose the normal distribution when estimating the GARCH model and we obtained satisfactory simulation results.[12] It is well known that the conditional normality assumption in ARCH models generates some degree of unconditional excess kurtosis (see, e.g., Bollerslev et al., 1992).

5.3. Estimation results

In this section we apply the model presented in Section 3 to the deseasonalized data for the three stocks. Observations of the first month are used for estimation and those of the last two months serve for forecast and validation. The likelihood function is maximized using Matlab v. 7 with the Optimization toolbox v. 3.0 and numerical derivatives are used for computing the standard errors of the estimates.

We first tested our durations for the clustering phenomenon using the Ljung–Box test statistic with 15 lags, $Q(15)$. The high coefficients (reported in Table 2) suggest the presence of ACD effects in our durations data at any reasonable significance level. The high positive serial dependence of the squared returns greatly exceeding the critical values, as illustrated by the Ljung–Box test statistics, noted $Q_2(15)$, represents evidence of volatility clustering and justifies the application of a GARCH-type model.

Estimation results are presented in Table 3 together with the Ljung–Box test statistics applied to the standardized residuals. We tried to find the best fit for our data and reestimated the model chosen without the variables that were statistically insignificant. We tested the adequacy of the log-ACD model using the Ljung–Box test statistic applied to the standardized residuals.[13] We judged the quality of the GARCH fit using the Ljung–Box test statistics for standardized residuals and their squares.

We retained a log-ACD(2,2)–ARMA(1,1)–EGARCH(1,1) model for the RY and AL stocks and a log-ACD(1,1)–ARMA(1,1)–EGARCH (1,1) model for the PDG stock. We found that the log-ACD specification used is successful in removing the autocorrelation in durations. For all stocks the Ljung–Box test statistics with 15 lags are inferior to the critical value. The parameters of the generalized gamma distribution are all significant. For each stock, the sum of the autoregressive parameters $\beta_1+\beta_2$ is close to one, revealing high persistence in durations.

With regard to the high-frequency returns, the ARMA(1,1)–EGARCH(1,1) specification accounts satisfactorily for the dependence of both the returns and squared returns of the PDG stock, as evidenced by the values of the order-15 Ljung–Box test statistics that are all less than the associated 95% critical value of 24.99. For the RY and the AL stocks, the Ljung–Box test statistics indicate that some serial dependence is still present in the data, which is inconsistent with the model's adequacy. However, the dependence is dramatically reduced compared to the original data. For example, the Ljung–Box test statistic with 15 lags $Q(15)$ for autocorrelation of RY returns was reduced from 8155 to 57. Similarly, the Ljung–Box test statistic $Q_2(15)$ for autocorrelation of RY squared returns was reduced from 14,690 to 41. Similar results have been found by Engle (2000) using Ljung–Box test statistics. Passing standard tests of model adequacy seems to remain an issue when using irregular high-frequency data due to the extremely high number of observations. We have tried higher-order models both for the mean and the variance

[12] We also estimated the model with the Student distribution. However, the results were not better than those obtained by presuming a conditionally normal distribution. In our applications, it seems that the tails are described just as well whether measured by distributions which are either conditionally Student or conditionally normal.

[13] We have also applied the test statistics for ACD adequacy introduced by Duchesne and Pacurar (2008) using the Bartlett kernel. The results are similar and therefore not reported.

Table 3
Estimates of log-ACD-ARMA-EGARCH models.

Parameter	RY		PDG		AL	
	Estimate	*t*-Student	Estimate	*t*-Student	Estimate	*t*-Student
β_0	−0.012	−30.51	−0.074	−14.13	−0.021	−46.91
α_1	0.078	189.57	0.073	14.66	0.097	213.73
α_2	−0.065	−159.89			−0.078	−172.99
β_1	1.556	464.59	0.923	44.33	1.472	1771.32
β_2	−0.567	−169.09			−0.479	−575.49
γ_1	0.423	91.13	0.378	56.33	0.220	110.79
γ_2	4.475	48.18	3.995	29.76	14	55.91
c			0.026	1.69	0.020	3.46
ϕ_1	0.077	4.38	0.060	2.03	0.031	1.99
θ_1	−0.609	−46.83	−0.350	−13.07	−0.410	−30.13
ω			0.072	6.34	0.016	5.56
a_1	0.248	21.49	0.257	9.38	0.213	23.54
$\widetilde{\beta_1}$	0.912	139.90	0.783	21.03	0.932	155.28
γ	0.072	3.94	0.224	8.99	0.223	15.03
$Q_d(15)$	22.52		15.18		19.95	
$Q(15)$	56.81		24.64		34.37	
$Q_2(15)$	40.90		22.25		68.78	

This table contains the parameter estimates of the log-ACD-ARMA-EGARCH models for RY, PDG and AL. *t-Student* are the *t*-statistics associated with the parameters estimates. $Q_d(15)$ represents the value of the Ljung–Box test statistic computed with 15 lags applied to the log-ACD standardized residuals. $Q(15)$ and $Q_2(15)$ represent, respectively, the values of the Ljung–Box test statistic computed with 15 lags for serial correlation in the standardized ARMA-EGARCH residuals and in their squares.

equations but we have not been able to gain any considerable improvement. We kept the ARMA(1,1)–EGARCH(1,1) specification of the model for its parsimony and because we are rather interested in assessing its forecasting ability under a risk management framework. As in Engle (2000), the MA(1) coefficient represented by θ_1 is negative and highly significant for all stocks. The AR(1) term represented by ϕ_1 is positive and significant. This can be explained by the fact that traders split large orders into smaller orders to obtain a better price overall and therefore make prices move in the same direction and thus induce a positive autocorrelation of returns (Engle and Russell, 2005). The positive autocorrelation can also be related to negative feedback trading (Sentana and Wadhwani, 1992). Engle (2000) has also found evidence of positive autocorrelation in high-frequency returns. The autoregressive parameter $\widetilde{\beta_1}$ for the EGARCH(1,1) model is close to one for the RY and the AL stocks, indicating a higher persistence in volatility than for the PDG stock. The parameter $\widetilde{\alpha_1}$ is statistically insignificant (not reported in the table) for all stocks, so no leverage effect is supported. The parameter γ is statistically different from zero and it is approximately equal to 0.224 for the PDG stock, to 0.223 for the AL stocks and to 0.072 for the RY stock.[14] Therefore, assuming an exogenous γ appears to be very restrictive.

Now that we have calibrated the model to our data, we may proceed to the simulation of future durations and returns.

5.4. IVaR backtesting

In this section we simulate tick-by-tick durations and returns using the estimated coefficients of our model and the observations of the last day as starting values. We sum up the irregular simulated returns over a fixed-time interval. Six different interval lengths are used: 15, 25, 35, 45, 55 and 90. Since all models are applied on deseasonalized data, the time unit of the interval does not represent a calendar unit (e.g., seconds, minutes). An approximate correspondence can nevertheless be established for each stock given the number of resampled intervals generated, depending on the trading intensity of that stock for the 44 days of the validation period. For example, a length = 15 for the RY stock results in 2462 intervals for 44 trading days, which is equivalent to an average interval of 7 min. Assessing the results in a regularly spaced framework allows us to evaluate the performance of the model in a traditional way, thus providing an appropriate benchmark. We have generated 5000 independent paths and we have extracted the IVaR as a percentile from the simulated distribution of returns.

We first analyze the performance of the model by using Kupiec's test (1995) for the percentage of failures, which is also embedded in the regulatory requirements on the backtesting of VaR models. We compute the empirical failure rate ($\hat{\alpha}$) as the percentage of times actual returns (y_k) are greater than the estimated IVaR. If the IVaR estimates are accurate, the failure rate should equal α. Under the null hypothesis that the model is correct we have $\hat{\alpha}=\alpha$ and Kupiec's likelihood ratio statistic takes the form:

$$\mathrm{LR} = 2\left[ln\left(\widehat{\alpha}^{\,m}(1-\widehat{\alpha})^{n'-m}\right) - ln\left(\alpha^m(1-\alpha)^{n'-m}\right)\right]$$

[14] We have also estimated our model on the original IBM data set used by Engle (2000) and the parameter γ had the value 0.387. Using a different sample period and an ACD/UHF-GARCH framework in which volatility is also a power function of duration, Engle and Sun (2007) estimated a value of 0.325 for a similar parameter.

Table 4
IVaR backtesting results.

VaR Level T	Kupiec test				DQ test				
	5%	2.5%	1%	0.5%	5%	2.5%	1%	0.5%	N
RY									
15	**0.000**	**0.003**	**0.010**	0.321	**0.000**	**0.004**	**0.030**	0.929	2462
25	**0.006**	0.177	0.535	0.661	0.163	0.375	0.868	0.986	1427
35	0.087	0.383	0.241	**0.022**	0.321	0.322	0.055	**0.004**	1008
45	0.086	0.184	0.673	0.593	0.267	0.672	0.992	1.000	780
55	0.143	0.772	0.796	0.923	0.080	0.252	0.998	0.999	634
90	0.202	0.843	0.935	0.463	0.679	0.356	0.987	0.965	384
PDG									
15	**0.000**	**0.001**	0.281	0.703	**0.000**	**0.000**	**0.000**	0.999	1611
25	0.133	0.349	0.654	0.752	0.108	0.549	0.099	0.999	933
35	0.148	0.381	0.827	0.177	0.300	0.802	0.999	0.542	655
45	**0.045**	0.446	0.405	0.388	**0.039**	**0.029**	**0.000**	**0.000**	503
55	0.195	0.808	0.086	0.226	0.833	0.594	**0.028**	0.833	409
90	0.470	0.483	0.158	0.181	0.833	0.672	0.785	0.689	248
AL									
15	**0.000**	**0.002**	**0.001**	**0.035**	**0.002**	**0.045**	**0.008**	0.825	2261
25	**0.041**	0.184	0.593	0.130	0.393	0.821	0.995	0.961	1286
35	0.643	0.746	0.742	0.816	**0.021**	0.457	0.997	0.993	900
45	0.052	0.410	0.973	0.774	0.485	0.678	0.999	0.941	691
55	0.065	0.077	0.184	0.906	**0.025**	0.222	0.716	0.999	560
90	0.452	0.396	0.408	0.816	0.500	**0.000**	0.980	0.986	338

This table contains the p-values for Kupiec's (1995) test and the DQ test of Engle and Manganelli (2004) using 5 lags and the current IVaR as explicative variable. T is the length of the interval used for simulation. N represents the number of intervals used for validation. Bold entries denote a failure of the model at the 95% confidence level since the p-values are inferior to 0.05.

where m is the number of exceptions and n' is the sample size. Under general conditions, this likelihood ratio is asymptotically distributed as a $\chi^2(1)$ under the null hypothesis. The left panel of Table 4 shows the p-values for the Kupiec test for the three stocks with a 5%, 2.5%, 1% and 0.5% IVaR level. Bold entries denote a failure of the model at the 95% confidence level, since the p-value is inferior to 0.05. The results show that the model performs well for all stocks. For almost all the intervals and the tails considered the p-values are superior to 0.05. The only exception is for smallest interval ($T=15$) and higher VaR levels ($\alpha=5\%,2.5\%$) which could be explained by the fact that the interval length is too small for obtaining non-zero returns when sampling the tick-by-tick real returns on regularly spaced intervals and, therefore, the theoretical number of IVaR violations cannot be achieved. For an interval equal to 25 and a 5% IVaR level, the model is rejected at a 95% confidence level but not at a 97.5% confidence level.

We also apply the dynamic quantile (DQ) test of Engle and Manganelli (2004) to check for another property a VaR measure should display: the hits (IVaR violations) should not be serially correlated. This can be tested by defining a sequence:

$$\mathrm{Hit}_k \equiv I(y_k < -\mathrm{IVaR}_k) - \alpha.$$

The expected value of Hit_k is 0 and the DQ test is computed using the regression of the variable Hit_k on its past, on current IVaR, and any other variables:

$$\mathrm{Hit}_k = XB + \epsilon_k$$

Then, $DQ = \hat{B}'X'\hat{B}/(\alpha(1-\alpha)) \sim \chi^2(l)$, where l is the number of explanatory variables and $\hat{B}$ is the OLS estimate of B. We perform the test using a matrix of explanatory variables, X composed of 5 lags of the hits and the current IVaR. Results are given in the right panel of Table 4 which reports the p-values of the DQ test. The results are satisfactory for the RY stock and coherent with the results from Kupiec's test. In most cases the p-values are larger than 0.05. For the PDG and the AL stocks, some estimates still show some predictability, especially for the smallest interval. Overall, it seems that the model performs best for the three stocks when a highest confidence level is considered.[15]

One might argue that the Monte Carlo simulation is very time-consuming. While this may be true for backtesting purposes that generally require a sufficiently large number of validation intervals, computing the next forecasted IVaR (which is normally required for practical purposes) is reasonably fast. For example, producing the next IVaR with a Pentium 4, one CPU 3.06 GHz and 5000 simulated paths for an interval equal to 90, takes us approximately 3.5 min for each of the three stocks considered. For our samples, an interval equal to 90 corresponds on average to 44 min for RY, 68 min for PDG and 50 min for AL. The computing time can be reduced to a few seconds by using more than one CPU and paralleled programming.

[15] We also did simulations with different means for durations and the model performed well in all cases (results are available upon request). It seems that the market's being more or less active has no impact on the model's performance.

Table 5
Comparison of three intraday VaR measures.

T (min)	Mean			Standard deviation		
	IVaR	VaR-GARCH	VaR-HS	IVaR	VaR-GARCH	VaR-HS
RY						
7	0.0033	0.0037	0.0034	0.0008	0.0015	0+
12	0.0039	0.0046	0.0040	0.0010	0.0022	0+
17	0.0044	0.0047	0.0045	0.0010	0.0019	0+
22	0.0049	0.0051	0.0053	0.0010	0.0019	0+
AL						
7	0.0044	0.0034	0.0037	0.0020	0.0023	0+
13	0.0054	0.0048	0.0048	0.0026	0.0044	0+
19	0.0062	0.0056	0.0056	0.0031	0.0036	0+
24	0.0067	0.0066	0.0060	0.0035	0.0035	0+
PDG						
10	0.0079	0.0080	0.0062	0.0025	0.0032	0+
18	0.0098	0.0099	0.0088	0.0035	0.0026	0.0001
25	0.0110	0.0122	0.0101	0.0041	0.0034	0+
33	0.0125	0.0150	0.0120	0.0047	0.0084	0.0001

T (min) indicates the intraday horizon of the IVaR expressed in minutes. *Mean* and *Standard deviation* represent, respectively, the mean and standard deviation for one day (May 1st). *IVaR* denotes the intraday VaR computed based on irregularly spaced data according to the method we propose. *VaR-GARCH* denotes a parametric VaR whose volatility is computed based on a GARCH model applied to the regularly spaced data. *VaR-HS* denotes a VaR computed by historical simulation, using regularly spaced data. 0+ denotes a number smaller than 10^{-4}.

5.5. Comparison of IVaR with regularly spaced intraday VaRs

While backtesting results show an adequate overall performance of the model, one could legitimately wonder, economically, just how close our IVaR estimates are to intraday VaR estimates based on regular sampling and computed over the same time interval. We compared the IVaR to an intraday VaR computed according to the historical simulation method (VaR-HS) and to a parametric intraday VaR computed using a GARCH model (VaR-GARCH) for volatility (as studied in Giot, 2005). To facilitate the comparison, seasonality has been re-introduced in the IVaR estimates.[16] We analyzed the values of the intraday VaRs for the first day of the validation interval (i.e., May 1st) for each of the three stocks. In computing the VaRs, we considered several confidence levels and the four shortest intraday time horizons used in backtesting, in order to obtain enough regularly time-spaced observations for estimation of all models. Note that these validation intervals can change from one stock to another depending on the liquidity of each stock. In other words, a T equal to 15 on the deseasonalized scale will correspond to 7 min for RY but to 10 min for PDG, because the higher trading frequency of RY increases the number of validation intervals obtained on the validation sample. Table 5 reports the 95%-level results for the three stocks.

As shown by the means of the three VaRs for 1 May 2001, the risk measure based on tick-by-tick data is, on average, reasonably close to the estimates based on regular sampling. The IVaR mean is a little higher than the VaR-GARCH mean for the AL stock and the VaR-HS mean for the AL and PDG stocks. The IVaR mean is slightly lower than the VaR-GARCH mean for the two other stocks, with a higher difference for the PDG stock for large T.

Looking at the standard deviations of the three VaRs, we notice more differences. The standard deviation of the intraday VaR-HS is close to zero. It seems that for our sample no wide variations in market conditions (that would have changed the intraday quantiles significantly) exist. We note that the intraday VaR-GARCH is usually more variable than the IVaR. This is important for risk-averse investors who prefer to have more information, but with less fluctuations during the day. Judging from our experience, it would appear that ignoring some information leads to higher and more variable risk measures.

Fig. 3A in Appendix A shows how the IVaR, VaR-GARCH, and VaR-HS evolve within the day (May 1st) for the RY stock, at a given T for a 95% level of confidence. Figures for other Ts and other stocks are similar and are also available upon request. These figures confirm that the three VaRs are, in general, reasonably close. However, we do note that, within the day, there can be large differences. As previously indicated, wide variations in market conditions are needed to produce any significant changes in quantiles for short intraday periods. Consequently, the intraday VaR-HS barely moves in comparison to the other VaRs.[17] In

[16] The seasonal factors were computed during the initial deseasonalization of the data and they depend on the time of transaction and the day of the week. When these coefficients are known and we then want to reintroduce seasonality, it is just a matter of incorporating the time of the next transaction (as predicted by the ACD model on the deseasonalized scale), in order to find the time of the next transaction in units corresponding to calendar time — that is, in order to predict the next raw duration. The same holds for raw returns which are also predicted by the model.

[17] The VaR-HS is the only VaR subjected to rolling computation (otherwise, it would remain constant).

Table 6
Decomposition of duration and return effects on total risk measure.

T (min)	Mean		Standard deviation	
	IVaR	IVaR_r	IVaR	IVaR_r
RY				
7	0.0033	0.0031	0.0008	0.0008
12	0.0039	0.0037	0.0010	0.0008
17	0.0044	0.0042	0.0010	0.0009
22	0.0049	0.0047	0.0010	0.0010
AL				
7	0.0044	0.0040	0.0020	0.0016
13	0.0054	0.0051	0.0026	0.0021
19	0.0062	0.0060	0.0031	0.0024
24	0.0067	0.0067	0.0035	0.0028
PDG				
10	0.0079	0.0071	0.0025	0.0016
18	0.0098	0.0091	0.0035	0.0023
25	0.0110	0.0105	0.0041	0.0026
33	0.0125	0.0119	0.0047	0.0033

T (min) indicates the intraday horizon of the IVaR expressed in minutes. *Mean* and *Standard deviation* represent respectively the mean and standard deviation for one day (May 1st). *IVaR* denotes the intraday VaR computed based on irregularly spaced data according to the method we propose. IVaR_r denotes the IVaR computed according to our method and making all the durations equal to the empirical mean.

contrast, the IVaR and the intraday VaR-GARCH evolve with daily market conditions. They are generally higher at the opening of the trading session, as predicted in the literature on market microstructure based on explanations drawn from information asymmetry models. In the Madhavan model (1992), information asymmetry is gradually resolved during the day, whereas Foster and Viswanathan (1994) develop a model in which the competition among informed investors leads to greater volatility at the start of the session. The IVaR fluctuates more at the beginning of the day. It is interesting to note that, seemingly, right around lunch time (when durations are longer) VaRs pass each others and that VaRs are more closely aligned at the end of the day than at its beginning. Finally, the VaR-GARCH does seem to fluctuate more than the IVaR, which can be explained as noted by the fact that its construction leaves out more information, particularly for large Ts. The smaller the T, the more comparable the VaRs appear to be, which confirms our intuition, since in this case the three techniques use almost the same data set. When we compare Fig. 3A and B in Appendix A that shows the three VaRs for RY for a small and a large T respectively, we note that the IVaR stays almost the same. The VaR-GARCH on the contrary fluctuates more when T increases which suggests that using an arbitrary period can have a different impact for traders.

6. Gains of implementing the IVaR

Compared to intraday VaRs based on regularly spaced data, the IVaR has two main advantages. First, it can be computed for any time horizon, once the model has been estimated, whereas in traditional methods the data set will usually change with the VaR's horizon (unless the volatility is scaled over time using the square-root-of-time rule).[18] In our case, estimation of the parameters of the log-ACD–ARMA–EGARCH model for computing the IVaR is independent of the horizon chosen. This has important practical consequences, since the practitioner who wants to change the VaR's horizon during the day will not have to estimate a new model. Moreover, two individuals or traders with different time horizons can very easily use the same model to compute their IVaR, whereas, with traditional methods, data processing and estimation procedures will be extra steps and the VaRs will not be based on the same data. If market conditions change during the day (for example, a very heavy increase in trading), an analyst or a trader might want to shorten the time horizon of his VaR. With the IVaR there would be no need to sample the data on the new horizon and estimate a new model.

Second, and perhaps more importantly, the IVaR has a higher informational content than traditional risk measures as it takes into account all of the observations available (instead of only one per time interval) and it explicitly incorporates the information contained in the times between transactions when modeling returns. Using a model that jointly describes the dynamics of duration and returns makes it possible to differentiate the impact of random trade durations on the risk measure from that of random returns.

[18] For example, computing a VaR for the next 5 min with traditional methods implies sub-sampling the data every 5 min while computing the VaR for the next 10 min will imply a new sub-sampling procedure which corresponds to this new horizon.

In fact, the IVaR can be broken down into three components: $IVaR = IVaR_r + IVaR_d + IVaR_{rd}$.[19] The $IVaR_r$ is computed by making all the durations equal to their empirical mean when computing the IVaR by Monte Carlo simulation. The $IVaR_r$ consequently measures the pure effect of the returns in the IVaR or the uncertainty about price movements. The $IVaR_d$ is obtained by making all the returns equal to their empirical mean when computing the IVaR. The $IVaR_d$ thus represents the quantile of a degenerated distribution of the returns and it is akin to a residual. We have in effect found it to be very close to 0 for the three stocks analyzed in this article. Finally, the $IVaR_{rd}$ takes into account the interaction between returns and durations. To be more precise, it measures the dependence effect between durations and returns. In practice, our model can use the gap between the IVaR and the $IVaR_r$ as a means of measuring the impact of irregularly spaced durations on risk, via the potential dependence between durations and returns.

Table 6 reports the means and standard deviations of the IVaR and the $IVaR_r$ for the three stocks for one day, using different *T*s. The IVaR is, on average, larger than the $IVaR_r$, which suggests that the component linked to the dependence between durations and returns is important. Since the IVaR accounts for several sources of risk, it is generally more variable that the $IVaR_r$.

Fig. 4 in Appendix A shows how the IVaR and the $IVaR_r$ evolve over one day for the RY stock for a given *T* and level of confidence. We have done the same exercise for the other stocks and several *T*s and other degrees of confidence, and the results are available. We see that the IVaR is higher than the $IVaR_r$ in the morning when the latter seems to underestimate the effect of duration on risk. The gaps between the two IVaRs narrow as the day advances. It thus seems to be worse to ignore the effect of durations at the beginning of the session than at the end. The fact that there is stronger dependence between durations and returns (as measured by the difference between the IVaR and the $IVaR_r$) at the beginning of the session, when information asymmetry is at its height, might be explained by a practice commonly used by informed investors: this practice consists of dividing up large orders to make them less noticeable and to minimize the impact on the price of their transaction (see e.g., Kyle, 1985; Foster and Viswanathan, 1990; Manganelli, 2005).

7. Conclusion

In this paper, we have proposed a method of computing an IVaR using tick-by-tick data. While the literature is full of efforts to develop sophisticated VaR models for daily data, here we investigate the use of irregularly spaced data for intraday risk management. This is particularly useful for agents who are very active in the market. We adopted an intraday Monte Carlo simulation approach which enabled us to forecast high-frequency returns for any arbitrary interval length, thus avoiding complicated time manipulations in order to return to a convenient regularly spaced framework. In our setup, the log-ACD model yields the consecutive steps in time while the ARMA–EGARCH model allows us to simulate the corresponding conditional tick-by-tick returns. Regularly spaced intraday returns are simply the sum of tick-by-tick returns simulated conditional on the forecasted duration.

We applied our methodology to three actively traded stocks on the TSX (RY, PDG and AL). Backtesting results indicate that our methodology for computing VaR with tick-by-tick data constitutes a reliable approach for measuring intraday risk. We show that there are two main advantages of the IVaR measure compared to intraday VaRs based on regularly spaced data. First, the IVaR has a higher informational content as it accounts for information contained in durations and in their interaction with returns. Using tick-by-tick data allows the user to differentiate on the total risk measure the impact of random durations from the impact of random returns. We found that the time between transactions is relevant to risk analysis, which is consistent with predictions from the market microstructure literature. Thus, the IVaR risk measure offers a more faithful reflection of traders' total risk. Second, the IVaR can be computed for any time horizon once the model has been estimated without requiring a new sampling and estimation every time the horizon changes, as is generally the case with traditional VaR techniques applied to intraday data.

Potential users of our approach would be traders on various types of markets who need intraday measures of risk; brokers and clearing firms looking for more accurate computation of margins; or any other entity interested in computing the VaR during a trading day in order to improve risk control. To compute the next IVaR using the method we propose, one has to monitor the time using the starting time (e.g., 11 a.m.) and the simulated durations for each path and, consequently, to apply the appropriate seasonal factors to reintroduce both interday and intraday seasonality in returns. The IVaR extracted from the simulated raw returns then has the usual interpretation.

One may argue that the true VaR is underestimated by using transaction data since transactions are observed only if the spread is favorable to the trader. In this respect, since transaction events are certainly timed, liquidity risk is not taken into account.[20] Like any standard VaR measure, our IVaR does not take into account the effect of liquidating a position. However, if data allowing full reconstruction of the real time order book were available, one could use our model in the spirit of Giot and Grammig (2006) to compute a liquidity-adjusted IVaR pertinent to impatient traders. Based on snapshots of the reconstructed order book taken for each trade, one could compute not only the price per share obtained when selling a given number of shares but also the duration associated with liquidating this position. These new series of prices and durations would then serve as input to our model and the resulted IVaR measure would thus take into account the uncertainty about both prices and durations from liquidating a given volume of shares. We plan to pursue this issue in future work.

[19] We are grateful to Christian Gouriéroux for sharing his insights into this issue.
[20] We thank Joachim Grammig for pointing it out.

Appendix A

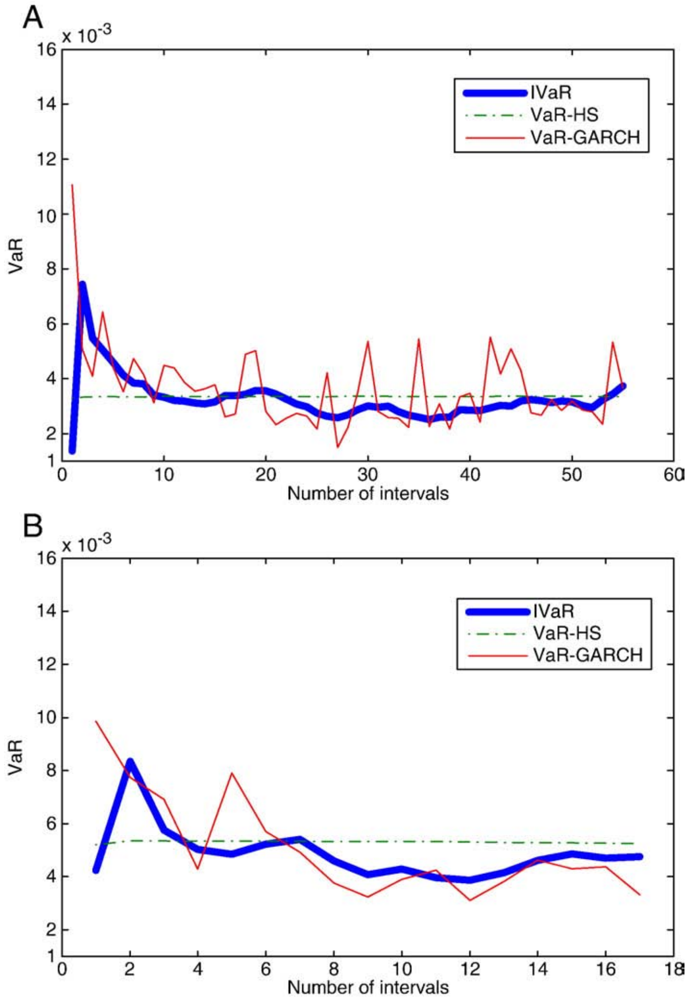

Fig. 3. A) RY estimated intraday VaR at 95% and $T=7$ min for May 1st, 2001. B) RY estimated intraday VaR at 95% and $T=22$ min for May 1st, 2001.

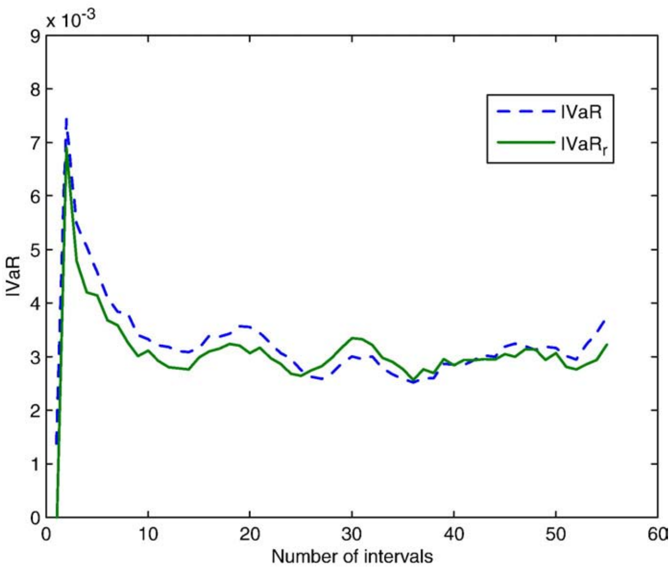

Fig. 4. RY estimated IVaR vs $IVaR_r$ at 95% and $T=7$ min for May 1st, 2001.

References

Admati, A.R., Pfleiderer, P., 1988. A theory of intraday patterns: volume and price variability. Review of Financial Studies 1, 3–40.

Anatolyev, S., Shakin, D., 2007. Trade intensity in the Russian stock market: dynamics, distribution and determinants. Applied Financial Economics 17, 87–104.

Bauwens, L., Giot, P., 2000. The logarithmic ACD model: an application to the bid–ask quote process of three NYSE stocks. Annales d'Économie et de Statistique 60, 117–149.

Bauwens, L., Giot, P., 2001. Econometric Modelling of Stock Market Intraday Activity. Advanced Studies in Theoretical and Applied Econometrics. Kluwer Academic Publishers, Dordrecht. 196 pages.

Bertram, W.K., 2005. A threshold model for Australian stock exchange equities. Physica A 346, 561–576.

Bollerslev, T., Chou, R.Y., Kroner, K.F., 1992. ARCH modeling in finance — a review of the theory and empirical evidence. Journal of Econometrics 52, 5–59.

Burns, P., (2002), "The quality of value at risk via univariate GARCH," Burns Statistics working paper, http://www.burns-stat.com.

Christoffersen, P.F., 2003. Elements of Financial Risk Management. Academic Press, San Diego. 214 pages.

Darolles, S., Gouriéroux, C., Le Fol, G., 2000. Intraday transaction price dynamics. Annales d'Économie et de Statistique 60, 207–238.

Davies, R.J., 2003. The Toronto Stock Exchange preopening session. Journal of Financial Markets 6, 491–516.

Diamond, D.W., Verrechia, R.E., 1987. Constraints on short-selling and asset price adjustments to private information. Journal of Financial Economics 18, 277–311.

Duchesne, P., Pacurar, M., 2008. Evaluating financial time series models for irregularly spaced data: a spectral density approach. Computers & Operations Research, Special Issue: Applications of OR in Finance, vol 35, pp. 130–155.

Easley, D., O'Hara, M., 1992. Time and the process of security price adjustment. Journal of Finance 47, 577–606.

Engle, R.F., 2000. The econometrics of ultra-high frequency data. Econometrica 68, 1–22.

Engle, R.F., Lange, J., 2001. Measuring, forecasting and explaining time varying liquidity in the stock market. Journal of Financial Markets 4 (2), 113–142.

Engle, R.F., Manganelli, S., 2004. CAViaR: conditional autoregressive Value at Risk by regression quantiles. Journal of Business and Economic Statistics 22, 367–381.

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

Engle, R.F. and Russell, J.R., (2005), "Analysis of high frequency data," forthcoming in Handbook of Financial Econometrics, ed. by Y. Ait-Sahalia and L. P. Hansen, Elsevier Science: North-Holland.

Engle, R.F. and Sun, Z., (2007), "When is noise not noise — a microstructure estimate of realized volatility," working paper, Stern School of Business, NY University.

Foster, F.D., Viswanathan, S., 1990. A theory of the interday variations in volume, variance, and trading costs in securities markets. Review of Financial Studies 3, 593–624.

Foster, F.D., Viswanathan, S., 1994. Strategic trading with asymmetrically informed traders and long-lived information. Journal of Financial and Quantitative Analysis 29, 499–518.

Ghysels, E., Jasiak, J., 1998. GARCH for irregularly spaced financial data: the ACD-GARCH model. Studies in Nonlinear Dynamics & Econometrics 2 (4), 133–149.

Giot, P., 2005. Market risk models for intraday data. European Journal of Finance 11, 309–324.

Giot, P., Grammig, J., 2006. How large is liquidity risk in an automated auction market? Empirical Economics 30, 867–887.

Gorski, A.Z., Drozdz, S., Speth, J., 2002. Financial multifractality and its subleties: an example of DAX. Physica A 316, 496–510.

Gouriéroux, C. and Jasiak, J., 1997. Local likelihood density estimation and Value at Risk, working paper, York University.

Gouriéroux, C., Jasiak, J., Le Fol, G., 1999. Intra-day market activity. Journal of Financial Markets 2, 193–226.

Grammig, J., Wellner, M., 2002. Modeling the interdependence of volatility and inter-transaction duration processes. Journal of Econometrics 106, 369–400.

Hautsch, N., 2004. Modelling Irregularly Spaced Financial Data — Theory and Practice of Dynamic Duration Models, Lecture Notes in Economics and Mathematical Systems, Vol 539. Springer, Berlin, 291 pages.

Jasiak, J., 1998. Persistence in intertrade durations. Finance 19, 166–195.

Jorion, P., 2000. Value at Risk: The New Benchmark for Managing Financial Risk. McGraw-Hill. 544 pages.

Kupiec, P., 1995. Techniques for verifying the accuracy of risk measurement models. Journal of Derivatives 2, 73–84.

Kyle, A.S., 1985. Continuous auctions and insider trading. Econometrica 53, 1315–1336.

Madhavan, A., 1992. Trading mechanisms in securities markets. Journal of Finance 47, 607–641.

Manganelli, S., 2005. Duration, volume and volatility impact of trades. Journal of Financial Markets 8, 377–399.

Meddahi, N., Renault, E., Werker, B., 2006. GARCH and irregularly spaced data. Economics Letters 90, 200–204.

Nelson, D.B., 1991. Conditional heteroskedasticity in asset returns: a new approach. Econometrica 59, 347–370.

Pacurar, M., 2008. Autoregressive conditional duration (ACD) models in finance: a survey of the theoretical and empirical literature. Journal of Economic Surveys 22, 711–751.

Russell, J., Engle, R.F., 2005. A discrete-state continuous-time model of financial transactions prices and times: the autoregressive conditional multinomial-autoregressive conditional duration model. Journal of Business & Economic Statistics 23, 166–180.

Sentana, E., Wadhwani, S., 1992. Feedback traders and stock return autocorrelations: evidence from a century of daily data. Economic Journal 102, 415–425.

Smith, B.F., Turnbull, D.A., White, R.W., 2006. The impact of pennies on the market quality of the Toronto Stock Exchange. Financial Review 41, 273–288.

Tsay, R.S., 2002. Analysis of Financial Time Series. Wiley, New York. 472 pages.