High-Frequency Risk Measures*

Denisa Banulescu[†], Gilbert Colletaz[‡], Christophe Hurlin[§] and Sessi Tokpavi[¶]

September 27, 2013

Abstract

This paper proposes an intraday high-frequency risk (HFR) measure specifically designed for HFR management and high-frequency trading (HFT). The HFR measure is a conditional joint measure of market risk and liquidity risk for irregularly spaced high-frequency data. It combines two well-known risk measures, *i.e.*, value at risk (VaR) and time at risk (TaR). We propose a forecasting procedure for both measures, which complies with HFR management requirements, particularly in terms of the information set. We also differentiate between three concepts of intraday VaR: total, marginal (or per time unit) and instantaneous VaR. Finally, we propose a backtesting procedure specifically designed to assess the validity of the VaR and TaR forecasts for each trade or other market microstructure event. The performance of the HFR measure is illustrated in an empirical application to two stocks (Bank of America and Microsoft) and an exchange-traded fund (ETF) based on Standard and Poor's (the S&P) 500 index. We show that the intraday VaR and TaR forecasts accurately capture the volatility and duration dynamics for these three assets.
Keywords: High Frequency Risk Measure, High Frequency Trading, Value at Risk, Time at Risk, Backtesting.

J.E.L Classification : C22, C52, G28

1 Introduction

This paper proposes a risk measure specially designed for high-frequency trading (HFT). In recent years, HFT has received extensive public attention and represents a significant share of financial market activity. As recently as 2010, HFT accounted for 56% of the equity trades in the US and 38% in Europe (Grant, 2010), and its importance is increasing. Consequently, the risk management of this activity becomes essential to financial market stability. However, the high-frequency risk (HFR) management is fundamentally different from the risk management imposed by the Basel regulations or the internal risk management practices generally adopted by financial institutions. This difference is related to the following unique characteristics of HFT (Gomber et al., 2013): (*i*) rapid order cancellation, (*ii*) very short holding positions, (*iii*) extracting very low margins per trade and (*iiv*) no significant position at the end of the day (flat position). In particular, end-of-day

*The authors thank Elena Dumitrescu for helpful comments on the paper as well as the participants of the the 7th international workshop "Methods in International Finance Network" (Namur, 2013). Usual disclaimers apply.
[†]Maastricht University, Department of Economics and University or Orléans, LEO, UMR CNRS 7322.
[‡]University of Orléans, LEO, UMR CNRS 7322.
[§]University of Orléans, LEO, UMR CNRS 7322. Corresponding author : christophe.hurlin@univ-orleans.fr.
[¶]University Paris-Ouest Nanterre, Economix, UMR CNRS 7235.

or intraday regularly spaced risk measures are not relevant in the HFT context. For instance, the risk caused by a statistical arbitrage algorithm or a short-term momentum strategy may be high during a short time period during the day, whereas no risk is reported at the end of the trading day.

Risk management for HFT requires new HFR measures for at least three reasons. First, to prevent or decrease the impact of a potential flash crash, regulators and academics advocate the establishment of circuit breakers throughout the market (Gomber et al., 2013).[1] A circuit breaker enforces a trade stop in stocks if their prices vary too quickly. In addition, these circuit breakers can be internalized in the HFT algorithms, in which case they are based on an expected change in price or an extreme risk measure, such as an intraday value at risk (VaR). Second, regulators worldwide are currently discussing whether HFT should require regulatory intervention. HFT can be regulated by imposing capital requirements, as it is performed for other types of risky financial activity (Biais, 2011). Such capital requirements would be based on specific intraday risk measures for the reasons previously mentioned. Finally, risk measurement for HFT should refer to both the market risk and the liquidity risk, even if the trading strategies are concentrated on highly liquid instruments.

To our knowledge, only two papers derive intraday market risk models using tick-by-tick data. Giot (2005) relies on two approaches.[2] First, he proposes to re-sample the data along a pre-specified time grid, which yields temporally equidistant observations (10 or 30 minutes). Then, standard time-series models (RiskMetrics- or GARCH-type models) are used to forecast the conditional volatility and the VaR for equidistantly time-spaced returns. However, this approach neglects the irregular timing of trades. Second, Giot derives the returns volatility (and thus the VaR) from the conditional intensity process associated with the price duration.[3] Following Engle and Russell (1998), he considers an autoregressive conditional duration (ACD)-type model to describe the dynamics of duration and compute the irregularly spaced VaR for price events. Then, this VaR is rescaled at fixed-time intervals for backtesting purposes. More recently, Dionne, Duchesne and Pacurar (2009) propose an intraday VaR (IVaR), which is based on a UHF-GARCH-type model (Engle, 2000) and a Monte Carlo simulation approach, to infer VaR at any fixed-time horizon. However, both studies consider only the market risk, through the VaR, and this risk measure is always rescaled at fixed time intervals because of the way it is constructed or for backtesting

[1] The most hazardous event on the stock market in which HFT has been involved occurred on May 6, 2010. In the course of 30 minutes, the Dow Jones Industrial Average dropped approximately 9%, followed by an equally rapid rebound. This event is now known as the Flash Crash and typically cited by HFT critics as an example of the substantial risk and volatility generated by this activity (Ahlstedt and Villysson, 2012).

[2] We do not mention here the papers that intraday VaRs for regularly spaced returns (Coreneo and Veredas, 2012,.etc.). Note also that Kozhan and Tham (2013) have recently proposed a measure for the execution risk in high-frequency arbitrage.

[3] Price durations are defined as the time necessary for a price of an asset to change by a given amount.

purposes.

This paper proposes an original intraday HFR measure defined as a joint measure of market risk and liquidity risk for irregularly spaced high-frequency data. The HFR measure combines two well-known risk measures: VaR and Time at Risk (TaR). The irregularly spaced intraday VaR corresponds to the maximum expected loss that will not be exceeded (for a given confidence level) during the time horizon until the next trade. The TaR, which was initially introduced by Ghysels, Gouriéroux and Jasiak (2004), is defined as the minimum duration prior to the next trade with a given confidence level. The HFR is measured at each trade and considers the irregular timing of transactions. One advantage of our measure is that it can be extended to any type of market microstructure event by considering a subset of the trades with specific characteristics or marks. We can define an HFR measure for the transactions associated with significant price changes (*i.e.,* price events), for example, or with a minimum volume (*i.e.,* volume events). This feature differentiates our measure from the intraday VaR proposed by Giot (2005) that is only valid for price events.[4]

The VaR and TaR are defined as quantiles of two conditional distributions: the distribution of the intraday returns and the distribution of durations. These variables are linked (Diamond and Verrecchia, 1987; Easley and O'Hara, 1992), and many approaches can be used to model their joint dynamic (Engle, 2000; Engle and Russell, 1997, 1998; Ghysels and Jasiak, 1998; Gerhard and Hautsch, 2002; Meddahi, Renault and Werker, 2005). In this context, the definition of the conditioning information set and the exogeneity assumptions made for both processes are important. When forecasting the volatility (and thus the VaR) for the next trade, which is indexed by $i+1$, two solutions can be adopted. The first one assumes that the past prices and other marks are known until the i^{th} trade (Ghysels and Jasiak, 1998). The second solution assumes that, in addition to this information, the duration between the i^{th} and the $i^{th}+1$ trades is also known (Engle, 2000; Meddahi, Renault and Werker, 2005). This difference is important to HFR management; the market risk is generally evaluated after each transaction for the time horizon of the future transaction and not one millisecond before its realization. Besides, the liquidity risk, as measured by the TaR, does not exist if we assume that the duration before the next trade is known. Therefore, we consider only the information available at the current trade to compute our HFR forecasts, which is one of the primary differences of our approach from the IVaR of Dionne, Duchesne and Pacurar (2009).

Additionally, we propose three original concepts of VaR: the *total* VaR, the *marginal* (or *per-*

[4] In fact, the link between the returns volatility and the conditional intensity process is valid only for this type of event. As noted by Giot (2001), "*a high trading intensity does not imply a corresponding price movement. For example, if the price goes back and forth between 100 $ and 100 and 1/4 $ every second, this asset features a large trading intensity, but no price movement. It can thus be argued that [trade] durations give relatively few information about the [volatility] of this asset*".

time-unit) VaR and the *instantaneous* VaR. The total VaR corresponds to the maximum expected loss for the time horizon of the next trade or microstructure event, whereas the per-time-unit VaR refers to the time horizon of the next second. The third concept corresponds to the maximum expected loss for the time interval between t_i, and $t_i + \Delta$, for an infinitesimal time increase Δ. We demonstrate that the third VaR corresponds to the intraday VaR proposed by Giot (2005) and thus differs from the total and per-time-unit VaRs that we recommend for risk management.

We present a simple forecasting algorithm for the HFR measure based on an EACD model for the conditional duration process and a time-varying GARCH model (Ghysels and Jasiak, 1998) for the conditional volatility. The two models are estimated by QML. To compute then the intraday TaR and VaR, we propose a semi-parametric approach similar to that considered by Engle and Manganelli (2001) in the day-to-day VaR perspective. No specific assumptions (except those required by the QML estimation method) are made regarding the conditional distributions of durations and returns.

Finally, we perform a backtesting procedure specifically designed to assess the validity of the VaR and TaR forecasts for each trade. In contrast to previous studies, we do not rescale the VaR forecasts to fixed-time intervals to apply the typical backtesting procedures. The performance of the HFR measure is evaluated at each transaction up to the time horizon of the next transaction, as is generally the approach in HFR management. We use three backtests that are compatible with irregularly spaced data: the LR test of Christoffersen (1998) based on a Markov-chain model, the duration-based test of Christoffersen and Pelletier (2004) and the GMM duration-based test proposed by Candelon et al. (2011).

The performance of the HFR measure is illustrated in an empirical application to three financial assets: Bank of America and Microsoft stocks and an exchange-traded fund (ETF) that tracks the Standard and Poor's (S&P) 500 index. The use of an ETF is justified by the increasing importance of these assets in the fund management industry.[5] For each asset, we compute a sequence of one-step-ahead forecasts for 1%-VaR and 1%-TaR based on the trade or price events recorded in September 2010. In all of the cases, the VaR and TaR forecasts accurately capture the volatility and duration dynamics. The VaR and TaR violations, which are defined as circumstances in which the ex-post tick-by-tick returns (durations for TaR) are smaller (higher for TaR) than the ex-ante VaR (TaR) forecasts, satisfy the unconditional coverage and independence assumptions (Christoffersen, 1998). The frequency of violations is always statistically not different from to the level of risk, *i.e.*, 1% in our case. More importantly, these violations are not clustered, which indicates that TaR and VaR forecasts adjust rapidly to the changes observed in past durations and

[5] At the end of August 2011, 2,982 ETFs worldwide were managing USD 1,348 bn, which represents 5.6% of the assets under the management of the fund management industry. Additionally, the total ETF turnover that occurs on-exchange via the electronic order book was 8.5% of the equity turnover (BlackRock, 2011).

returns. Finally, we demonstrate that these backtesting results are valid throughout the day and the week, even if the EACD-GARCH model is not re-estimated. In addition, our findings are also robust to the choice of the split point used to separate the in-sample and out-of-sample periods (Hansen and Timmermann, 2012).

The remainder of the paper is organized as follows. In the first section, we define the HFR measure and introduce the method used to model the dynamics of returns and durations. The second section proposes the forecasting algorithm and the intraday periodicity adjustment. The third section presents the empirical application and the backtesting procedure, whereas the fourth section discusses marginal VaR and instantaneous VaR. The last section presents conclusions and suggests further research.

2 High-frequency Risk Measure

The tick-by-tick data for a given stock are described by two variables: the time of the transactions and a vector of marked point processes. The latter variable contains, for example, the volume, the bid-ask spreads, the price of the contract observed at the time of the transaction. Consider a trade that occurred at time t_i at a log-price p_i, and denote z_i the corresponding vector of marks other than the price. Thus, the duration between two consecutive trades is defined as $x_i = t_i - t_{i-1}$ and the corresponding continuously compounded return is $r_i = p_i - p_{i-1}$. The information set available at time t_{i-1} is denoted by $\mathcal{F}_{i-1}$ and includes all past durations and marked point processes: $\mathcal{F}_{i-1} = \{x_j, p_j, z_j\ , j \leq i-1\}$. In this framework, we propose an intraday risk measure defined as follows.

Definition 1 *For a shortfall probability $\alpha \in\ ]0,1[\,$, the HFR measure for the i^{th} trade is the combination of VaR and TaR, such that:*

$$\Pr\left(r_i < -VaR_i(\alpha) \mid \mathcal{F}_{i-1}\right) = \alpha, \tag{1}$$

$$\Pr\left(x_i > TaR_i(\alpha) \mid \mathcal{F}_{i-1}\right) = \alpha. \tag{2}$$

The HFR measure takes into account the irregular timing of trades and can be considered a risk measure for market risk and liquidity risk. The HFR measure's first component is an intraday irregularly spaced VaR. In this context, the VaR is defined as the maximum expected loss that will not be exceeded under normal conditions for a given confidence level $1-\alpha$ at the time horizon of the next trade. Note that VaR is defined in *event* time, not calendar time. This definition differentiates our measure from the intraday VaRs proposed by Giot (2005) and Dionne, Duchesne and Pacurar (2009), which are generally averaged for regularly time-spaced intervals for backtesting purposes. The second dimension of our risk measure is the TaR. TaR is defined as the minimum duration before the next trade may occur for a confidence level of $1-\alpha$ and can be interpreted as a liquidity risk measure. That is, the length of intra-trade durations reveals the speed of activity and is a natural indicator of market liquidity. The risk levels for TaR and VaR can differ. However, in the rest of this study, we will consider an identical level of risk α.

Another advantage of the HFR measure is that it can be extended to any type of market microstructure event. The previous definition was proposed for a trade for which the event arrival corresponds to the transaction arrival time. Alternatively, the HFR measure can be considered for a subset of trades with specific characteristics or marks.[6] For instance, we can define an HFR for price events. First introduced by Engle and Russell (1997), the price duration represents the time necessary for the price of an asset to change by a given amount C, *i.e.,* $|p_i - p_{i-1}| \geq C$, where i is the index of the most recently selected process among the initial point processes. x_i and r_i denote the price duration and the log-return defined in price-event time, respectively, and the HFR measure corresponds to the combination of *conditional* VaR and TaR (with respect to the condition $|r_i| \geq C$). This concept is similar to the CoVaR (conditional value at risk) proposed by Adrian and Brunnermeier (2011) in the context of systemic risk. The interpretation is the following: the conditional VaR is the maximum expected loss that will not be exceeded under normal conditions for a confidence level α at the time horizon of the next trade for which the log-price variation will exceed the threshold C. The conditional TaR becomes the minimum duration before the next price event for a confidence level of $1-\alpha$.

The HFR measure can be extended to any type of thinned point process and can be defined, e.g., in price events, volume events (the time necessary for trading a given number of shares). The only changes required are the interpretation of the HFR measure and the introduction of a conditioning event, denoted $\mathbb{C}_i$.

$$\Pr\left(r_i < -VaR_i(\alpha) \mid \mathcal{F}_{i-1}; \mathbb{C}_i\right) = \alpha, \tag{3}$$

$$\Pr\left(x_i > TaR_i(\alpha) \mid \mathcal{F}_{i-1}; \mathbb{C}_i\right) = \alpha. \tag{4}$$

where $\mathbb{C}_i = \varnothing$ in the case of trade events, $\mathbb{C}_i = \{|r_i| \geq C\}$ in the case of price events, etc.

Both VaR and TaR are defined as quantiles of two conditional distributions: the conditional distribution of the intraday returns and the conditional distribution of durations. These distributions are not independent. Several theoretical models and empirical studies in the microstructure literature report a link between market activity, which is measured by the time interval between two consecutive transactions, and price dynamics. For instance, Diamond and Verrecchia (1987) and Easley and O'Hara (1992) analyze the trading mechanism in an asymmetric-information model

[6] In the point process literature, retaining only the arrival times that are thought to carry special information is referred to a thinning the point process (Pacurar, 2008).

and demonstrate that the trading volume, the bid-ask spread and the trading activity are directly related to the behavior of informed traders. They conclude that market activity plays an important role in price adjustment because time conveys information. In Diamond and Verrecchia, the intensity of durations is a signal of the quality of news (“good" or “bad"), whereas in Easley and O'Hara, it indicates a news occurrence (e.g., *long durations* are likely to be associated with *no news*) because variations in trading rates are associated with changes in the number of informed traders.

3 Methodology

The aim of the HFR measure is to quantify the risk on durations and the risk on price returns, while considering the irregular timing of trades or other types of market microstructure event. These risks are likely to be related. In this section, we present the models used to describe the dynamics of durations (and TaR) and tick-by-tick returns (and VaR).

3.1 TaR and Durations

The first component of the HFR measure, *i.e.*, the TaR, is determined by the dynamics of durations. We assume that the duration dynamics are described by by an autoregressive conditional duration (ACD) model (Engle and Russell, 1998). Let $\psi_i = \mathbb{E}(x_i|\mathcal{F}_{i-1})$ denote the expectation of the i^{th} duration conditional on the information set available at time t_{i-1}. We assume the following:

$$x_i = \psi_i v_i, \tag{5}$$

where v_i is an *i.i.d.* positive-valued process with a pdf denoted by $f_v(.)$ such that $\mathbb{E}(v_i) = 1$. Various specifications (e.g., linear, log-linear, switching regime) can be considered here for the dynamics of the conditional duration ψ_i and for the density of v_i: each configuration corresponds to a specific ACD model (for a survey, see Pacurar, 2008 and Bauwens and Hautsch, 2009). In order to illustrate our HFR measure, we consider here the simplest and highly successful member of this family, i.e., the EACD(p, q) model, where E accounts for the exponential distribution of the errors:

$$\psi_i = w + \sum_{j=1}^{p} \alpha_j x_{i-j} + \sum_{j=1}^{q} \beta_j \psi_{i-j}, \tag{6}$$

$$f_v(v_i) = \exp(-v_i), \tag{7}$$

with $w > 0$, $\alpha_j \geq 0$ and $\beta_j \geq 0$ to ensure the positivity of ψ_i, as in ARCH-type models. Under these assumptions, the TaR is given by:[7]

$$TaR_i(\alpha) = -\psi_i \ln(\alpha). \tag{8}$$

[7] The quantile function of an exponential distribution with a rate parameter λ, is defined as $F_{\exp}^{-1}(p;\lambda) = -\ln(1-p)/\lambda$. In our case, $\lambda = 1$ and $p = 1-\alpha$.

The exponential distribution implies a flat conditional hazard function, which is restrictive. An alternative to specifying another parametric distribution for v_i (e.g., Weibull, generalized gamma, Burr, etc.) is to view the EACD as a QML. A significant advantage of the exponential distribution is that it provides consistent quasi-maximum likelihood (QML) estimators for the ACD parameters.[8] This property is particularly important for computing the TaR. Therefore, if the conditional duration is correctly specified, even if the conditional density (Equation 7) is misspecified, consistent QML estimates of ψ_i can be obtained when the true density (unknown) of v_i belongs to the linear exponential family (Gouriéroux, Montfort and Trognon, 1984). Then, the TaR can be alternatively defined by:

$$TaR_i(\alpha) = \psi_i \; F_v^{-1}\left(1-\alpha\right), \tag{9}$$

where $F_v^{-1}(p)$ denotes the p-quantile associated with the (unknown) distribution of v_i. This quantile can be estimated by a non-parametric method based on the standardized durations x_i/ψ_i.

The TaR can be computed for trade durations but also for any other thinned point process, such as price durations and volume durations. For thinned point processes, two solutions can be adopted. The first solution estimates the conditional intensity for the trades and then deduces the intensity of the thinned point process. The relationship between both intensity functions depends on the probability of each trade being selected (Daley and Vere-Jones, 2007). For instance, for price durations, the relationship depends on the probability of observing a trade with a price change larger than the threshold C. The second approach is more direct (Engle and Russell, 1997, 1998). It estimates the conditional mean function of price (or volume) durations directly from the price (or volume) events. In the remainder of the study, we will use the second approach to compute the conditional TaR for a particular subset of trades.

3.2 VaR and Volatility

The HFR measure's second component is the VaR. The VaR depends on the total variance of the return process. For simplicity, let us consider the demeaned return process $r_{c,i} = r_i - \mathbb{E}\left(r_i | \mathcal{F}_{i-1}\right)$ associated with the i^{th} trade and assume that:[9]

$$r_{c,i} = \sqrt{h_i}\varepsilon_i, \tag{10}$$

where $h_i = \mathbb{E}(r_{c,i}^2|\mathcal{F}_{i-1})$ denotes the *total* (conditional) variance of the tick-by-tick returns and ε_i is a strong white noise process with $\mathbb{E}\left(\varepsilon_i\right) = 0$ and $\mathbb{E}\left(\varepsilon_i^2\right) = 1$. If the distribution of ε_i belongs

[8] More generally, Drost and Werker (2004) demonstrate that consistent estimates are obtained when the QML estimation is based on the standard gamma family of distributions, to which the exponential distribution also belongs.

[9] As in most empirical studies based on tick-by-tick data, we consider continuous prices and returns. However, intraday transaction prices often take only a limited number of different values because of certain institutional features with regard to price restrictions. An alternative consists of using an autoregressive conditional multinomial (ACM) model to account for the discreteness of transaction prices (Engle and Russell, 2005).

to the location-scale family, the $\mathcal{F}_{i-1}$-conditional VaR can be defined as a linear function of the conditional volatility (Jorion, 2007):

$$VaR_{c,i}(\alpha) = -\sqrt{h_i} F_{\varepsilon}^{-1}(\alpha), \tag{11}$$

where $F_{\varepsilon}(.)$ denotes the cdf of the innovation process ε_i.

Numerous approaches can be used to model the dynamics of h_i. In the context of intraday VaR measures, Giot (2005) proposes to re-sample the data along a pre-specified time grid, which yields temporally equidistant observations. Then, standard time-series models (RiskMetrics- or GARCH-type models) can be used to forecast h_i and the corresponding VaR for equidistantly time-spaced returns.[10] However, this approach neglects the irregular timing of trades and can be considered to be a marginal approach because the dynamics of returns depend only on past returns and the time between market events does not enter into the analysis.

Alternatively, the use of tick-by-tick data facilitates the joint modeling of durations and volatility. Here, a key element is the conditioning information set. The conditioning information $\mathcal{F}_{i-1}$, which is used to define TaR and VaR, includes past returns, durations and other marks observed until the event $i-1$, *i.e.*, $\mathcal{F}_{i-1} = \{x_j, p_j, z_j, j \leq i-1\}$. Consequently, conditionally on $\mathcal{F}_{i-1}$, the current duration x_i and return r_i are stochastic. Generally, a particular form of exogeneity is assumed, which substantially simplifies the estimation and forecasting procedures (Meddahi, Renault and Werker, 2005). For instance, Engle (2000), Meddahi, Renault and Werker (2005) and Dionne, Duchesne and Pacurar (2009) propose to model the conditional volatility by considering an information set $\mathcal{G}_{i-1}$ that includes $\mathcal{F}_{i-1}$ and the current duration x_i such that $\mathcal{G}_{i-1} = \{x_j, p_j, z_j, j \leq i-1; x_i\}$. In the context of intraday risk measures, Dionne, Duchesne and Pacurar (2009) use a UHF-EGARCH-type model (Engle, 2000) to define their $\mathcal{G}_{i-1}$-conditional IVaR. Therefore, the duration prior to the next trade is assumed to be known when evaluating the risk associated with the future price change. In contrast, the $\mathcal{F}_{i-1}$-conditional VaR defined in Equation (11) depends only on the information available at t_{i-1}, *i.e.*, the time of the last trade. This difference is important in the context of HFR management (Figure 1). The use of $\mathcal{G}_{i-1}$ (the bottom part of Figure 1) corresponds to a particular setting. In general, high-frequency traders and risk managers are interested in conditionally forecasting the market risk until the next event based on the information available at the time of the last event (the top part of Figure 1). Finally, HFR measures are necessarily defined with respect to $\mathcal{F}_{i-1}$ because if the duration x_i is known, TaR is useless.

[10] As we will discuss in the final section, Giot (2005) also proposes to derive an irregularly time-spaced VaR from the instantaneous volatility and the price intensity.

Figure 1: Information sets $\mathcal{F}_{i-1}$ and $\mathcal{G}_{i-1}$

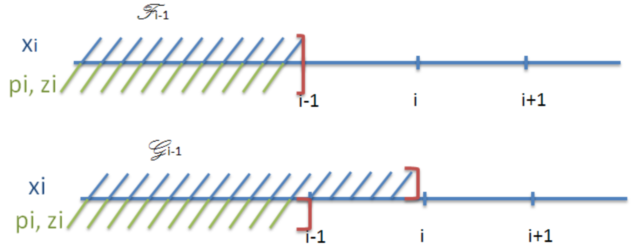

A solution here consists in considering the ACD-GARCH model proposed by Ghysels and Jasiak (1998). Based on the temporal aggregation of GARCH process results of Drost and Nijman (1993) and Drost and Werker (1996), Ghysels and Jasiak prove that the volatility model can be described by a time-varying coefficient GARCH model:

$$h_i = \omega_{i-1} + \alpha_{i-1}\varepsilon_{i-1}^2 + \beta_{i-1}h_{i-1}, \tag{12}$$

where the time-varying parameters $\omega_{i-1} = f_\omega\left(\omega, \alpha, \beta, \kappa, \psi_i\right)$, $\alpha_{i-1} = f_\alpha\left(\omega, \alpha, \beta, \kappa, \psi_i\right)$ and $\beta_{i-1} = f_\beta\left(\omega, \alpha, \beta, \kappa, \psi_i\right)$ are functions of the expected duration $\psi_i = \mathbb{E}(x_i|\mathcal{F}_{i-1})$, the kurtosis of the returns κ and three structural parameters denoted ω, α and β.[11] These parameters correspond to those of a GARCH model defined on regularly time-spaced data with a frequency equal to the time scale, *i.e.* one second in our case. The time-varying parameters ω_{i-1}, α_{i-1} and β_{i-1} correspond to the parameters of the weak GARCH process associated with the cumulated per-second returns observed over a period of ψ_i seconds (for more details, see Drost and Nijman, 1993). The formal expressions of ω_{i-1}, α_{i-1} and β_{i-1} are provided in Appendix A.

The parameters of the EACD and the GARCH models can be estimated using a GMM procedure or a two-step estimation method. The latter method first estimates the parameters of the ACD and then the parameters of the GARCH model. The ACD-GARCH is genuinely bivariate because past volatilities may determine the transaction durations and vice versa. The GMM procedure is suitable when trading and volatility are not mutually exogenous. In contrast, the use of the two-step approach is based on the assumption that the spacing between transactions is weakly exogenous in the sense of Engle, Hendry and Richard (1983) with respect to the return process (Appendix A).

Finally, the VaR defined in Equation (11) can be forecast for each trade and for any other subset of transactions. Then, the index i corresponds to the index of the most recently selected process among the initial point processes, and the information set includes both $\mathcal{F}_{i-1}$ and the appropriate conditioning event $\mathbb{C}_i$ (for instance, $|r_i| \geq C$ in the case of a price event).

[11] Given the functional forms derived by Drost and Nijman (1993) for a weak GARCH representation, Ghysels and Jasiak (1998) consider a restriction of the form $\alpha_{i-1} + \beta_{i-1} = (\alpha + \beta)^{\psi_i}$.

4 HFR Measure and Irregularly Spaced Data

The definition of an HFR measure requires a careful treatment of the intraday tick-by-tick data. In this section, we address two specific dimensions: (i) the use of the mid-quotes and/or transaction prices and (ii) the intraday periodicity adjustment. Then, we present the forecasting procedures that we propose for the HFR measures.

4.1 Data

The relevance of our HFR measures is assessed based on the tick-by-tick data for two financial stocks (Microsoft (MSFT) and Bank of America (BAC)) and one ETF (SPY), provided by SPDR ETFs, that tracks the S&P 500 index. The data are extracted from the Trade and Quote (TAQ) database and include information on each trade and quote during the period from September 1 to October 29, 2010 (42 trading days). The database consists of two parts. The trade database summarizes the trading process and contains the date and time stamp t_i for the i^{th} trade and additional marks, such as transaction log-prices (p_i) and volume. The quote database contains the quoting process and reports the date and time t_j of the j^{th} quote in addition to the bid (b_j) and ask (a_j) prices.

Because high-frequency data could contain incorrectly recorded elements, all of the series were cleaned prior to use by applying the set of baseline rules proposed by Brandorff-Nielsen et al. (2009). The two parts of the TAQ database were corrected and merged to adjust the trade prices using rules that include the bid-ask spread. Negative trade prices or volumes, zero bid and ask values were deleted. After merging the two databases, several simultaneous transactions were observed. Therefore, we considered only the first time stamp. The price and quote series are reported every trading day from 9:30 am to 4 pm Eastern Standard Time (EST), and the effect of the overnight or the opening auction is removed by deleting the opening trades. A total of 223,502 observations for the S&P 500, 78,265 for BAC and 432,626 for MSFT remained.

The HFR measures can be applied to either transaction prices or mid-quotes defined as $p_i = (b_i + a_i)/2$. In the high-frequency volatility literature, mid-quotes are commonly used instead of transaction prices to avoid the bid-ask bounce exhibited by the trade process. Consequently, Giot (2005) defines his IVaR for mid-quotes. However, Dionne, Duchesne and Pacurar (2009) argue that the use of mid-quotes may understate the true VaR. Because we are interested in estimating the VaR for real transactions, we also consider the transaction prices. Nevertheless, prices may be affected by various microstructure noises, and the corresponding returns are generally serially autocorrelated. To capture such microstructure effects, an ARMA(p, q) model is generally used for the tick-by-tick returns. Given the autocorrelation and partial autocorrelation functions of the returns, we use a simple AR(1) model for all of the assets and denote by $\mu_i = \mathbb{E}(r_i|\mathcal{F}_{i-1})$ the conditional mean of the returns with:

$$\mu_i = \theta + \rho\, r_{i-1}. \tag{13}$$

Then, the HFR measure's VaR component is expressed as a function of the VaR obtained for the demeaned process (Equation 11) as follows:

$$VaR_i(\alpha) = -\mu_i + VaR_{c,i}(\alpha). \tag{14}$$

Simple descriptive statistics on durations and returns are presented in Table 1. For each of the three series under analysis, we report statistics on the trade durations and returns as well as on the price durations and the corresponding returns. Price durations are defined for two stock-specific thresholds C. These thresholds are determined by analyzing the empirical distributions of the absolute returns (Appendix B, Figure B1) and do not necessarily correspond to the price increments. The total number of observations decreases by approximately 50% and 80% between the entire sample and the samples of price events.

Among the three stocks, Microsoft is the most traded asset; the average trade duration is 12.55 seconds for BAC, 4.39 seconds for the S&P 500 and only 2.27 seconds for MSFT. For all of the assets, the average price duration is larger than the trade duration and increases with the threshold C; the larger the price change is, the more significant is the time required for the price to change by this amount. As already observed in other studies, the trade and price durations for the three assets exhibit a positive autocorrelation, an overdispersion (*i.e.*, the standard deviation is greater than the mean) and a right-skewed shape (Pacurar, 2008). The overdispersion suggests that the exponential distribution is not appropriate for the unconditional distribution of trade or price durations but not that conditional durations cannot be exponentially distributed. The high values of the Ljung-Box Q-statistics obtained for 10 or 20 lags indicate the presence of ACD effects (duration clustering) at any reasonable level.

The mean transaction return is extremely small for all of the assets and much lower than its standard deviation. The return associated with the price event is not always larger (on average) than the trade return. However, its volatility is naturally higher. For all of the market microstructure events considered, BAC is the most volatile asset. The volatility of the BAC trade returns is nearly five times larger than that of the market index and two times that of MSFT. All of the trade return series display a kurtosis higher than that of a normal distribution. However, the leptokurticity decreases with the threshold C. The skewness is always positive for the S&P 500, negative for MSFT, and dependent on the threshold C for BAC. Finally, as previously mentioned, the returns are autocorrelated as the Q-statistics obtained for 10 or 20 lags result in the rejection of the null hypothesis of no autocorrelation. However, the values of the Q-statistics tend to decrease with the threshold C (except for MSFT).

4.2 Intraday Periodicity Adjustment

In addition, it is important to consider the intraday seasonal patterns of the data. It is widely documented that high-frequency data exhibit a strong seasonality because of the progression of the market activity during the same day (intraday periodicity) and during each day of the week (interday periodicity). Generally, this pattern must be removed from the data prior to fitting the model to avoid distortions in the results.

The price duration exhibits significant diurnal patterns that could affect the quality of our results. Consequently, we define the price duration variable as follows:

$$x_i = \bar{\omega}_d\,(t_i)\,\widetilde{x}_i, \tag{15}$$

with $\bar{\omega}_d\,(t_i)$ is the seasonal component and $\widetilde{x}_i$ is the seasonally adjusted duration. Many procedures have been proposed to estimate and adjust for the seasonal component (see Wood et al., 1985; Haris, 1986; Andersen and Bollerslev, 1997, 1998; Bauwens and Giot, 2000; Dufour and Engle, 2000; Drost and Werker, 2004; Hecq, Laurent and Palm, 2012; and others). We follow the same approach as Dionne, Duchesne and Pacurar (2009) and Anatolyev and Shakin (2007). The deterministic component is identified as the expected trade duration conditional on the time-of-day and the day-of-week effects. For each day of the week, the seasonal factor is computed by averaging the durations over 30-minute intervals. Then, cubic splines are used to smooth the time-of-day effect and to extrapolate the latter effect for any time during the day.

Figure 2 displays the estimated intraday seasonal components for the duration series of the S&P 500, BAC and MSFT. The overall conclusion is that a day-of-week effect exists because the time-of-day component for Monday differs from the time-of-day component for Tuesday and so on throughout the week. The well-documented inverted-U shape (see, for example, Engle and Russell, 1998; Bauwens and Giot 2000; etc.) indicates that the market activity (resp. trade duration) is higher (resp. shorter) at the opening and closing of the trading day than around midday. To remove this seasonal component, we divide the raw data of observed price durations by the time-of-day effect and run an EACD model on the stochastic deseasonalized component $\widetilde{x}_i$.

In addition, returns must be diurnally adjusted because volatility present also a daily-specific shape. Engle (2000) proposes to adjust the volatility *per-time-unit* by considering $r_i/\sqrt{x_i}$ and a spline procedure identical to that used for durations. However, this approach could result in a residual intraday seasonality in the *total* volatility. Therefore, we prefer to directly adjust the squared returns (instead of returns per-time-unit) for intraday periodicity. We assume that:

$$r_{c,i} = \sqrt{\bar{\omega}_r\left(t_i\right)}\, \widetilde{r}_{c,i}, \tag{16}$$

where $\widetilde{r}_{c,i}$ denotes the seasonally adjusted demeaned return and $\sqrt{\bar{\omega}_r\left(t_i\right)}$ is the corresponding seasonal component. As for the durations, for each day of the week, the seasonal factor is computed by averaging the corresponding values over 30-minute intervals and then using a cubic spline. The right panel of Figure 2 displays the estimated intraday seasonal components for volatility. The patterns are analogous to those found in previous studies. The seasonal component exhibits a U-shape, which indicates that the return volatility is generally lower in the middle of the day than at the beginning and end of the day. However, even if the effect of the overnight or the opening auction is removed, the volatility is generally higher at the opening than at the close of the trading day.

4.3 Forecasting Algorithm

This section describes the algorithm required to forecast our HFR measures. The last event of the in-sample period is n, and $n+1$ is the first event of the out-of-sample period. The procedure to forecast $TaR_{n+1}(\alpha)$ and $VaR_{n+1}(\alpha)$ is as follows:

i) The parameters of the EACD-GARCH model are estimated (Appendix A) on deseasonalized duration and tick-by-tick demeaned return series, $\{\widetilde{x}_i\}_{i=1}^n$ and $\{\widetilde{r}_{ic}\}_{i=1}^n$.

ii) We compute the one-step-ahead out-of-sample expected duration $\widehat{\psi}_{n+1}$ and volatility $\widehat{h}_{n+1}$ for the seasonally adjusted series.

iii) The quantiles of the innovation processes, ε_i and v_i, are estimated by the corresponding empirical quantiles of the in-sample series of standardized and deseasonalized returns $\{\widehat{\varepsilon}_i\}_{i=1}^n$ and durations $\{\widehat{v}_i\}_{i=1}^n$.

$$\widehat{F}_{\varepsilon}^{-1}\left(\alpha\right) = percentile\left(\{\widehat{\varepsilon}_{t_i}\}_{i=1}^n, 100\alpha\right), \tag{17}$$

$$\widehat{F}_{v}^{-1}\left(1-\alpha\right) = percentile\left(\{\widehat{v}_i\}_{i=1}^n, 100\left(1-\alpha\right)\right), \tag{18}$$

with $\widehat{\varepsilon}_i = \widetilde{r}_{c,i}/\widehat{h}_i^{1/2}$ and $\widehat{v}_i = \widetilde{x}_i/\widehat{\psi}_i$ for $i = 1, ..., n$.

iv) Finally, the VaR and TaR out-of-sample forecasts are given by:

$$TaR_{n+1}(\alpha) = \widehat{F}_{v}^{-1}\left(1-\alpha\right)\ \widehat{\psi}_{n+1}\, \bar{\omega}_d\left(t_{n+1}\right), \tag{19}$$

$$VaR_{n+1}(\alpha) = -\widehat{\mu}_{n+1} - \widehat{F}_{\varepsilon}^{-1}\left(\alpha\right)\ \widehat{h}_{n+1}^{1/2}\, \bar{\omega}_r^{1/2}\left(t_{n+1}\right). \tag{20}$$

where $\bar{\omega}_z\left(t_{n+1}\right)$ denotes the deterministic seasonal component (estimated in-sample) associated with time t_{n+1}, for $z = \{d, r\}$. Because the time of the next event t_{n+1} is unknown, a feasible estimator of VaR and TaR can be obtained by replacing $\bar{\omega}_z\left(t_{n+1}\right)$ with $\bar{\omega}_z\left(t_n\right)$.

By replicating this procedure M times, we obtain a sequence of out-of-sample, one-step-ahead forecasts $\{VaR_{n+j}(\alpha)\}_{j=1}^{M}$ and $\{TaR_{n+j}(\alpha)\}_{j=1}^{M}$. The model is not re-estimated at each step, as in a standard rolling windows scheme. The parameters are estimated on the first n observations and then considered as fixed for $n+1, ., n+M$. For S&P 500 and BAC, the estimation period is September 1 to September 22, 2010 (15 trading days). For MSFT, because of the large number of observations, the estimation period ends on September 15. The estimation samples include 80,845 observations for S&P 500, 27,972 for BAC and 99,155 for MSFT.

5 Empirical Analysis

In this section, we apply the forecasting procedure described in the previous section to the three assets by considering the trade and the price events successively.

5.1 Estimation Results and HFR forecasts

Estimation Results: The estimation results for the various EACD-GARCH models are reported in Tables 2 and 3. For each asset, the first column corresponds to the estimates for the trades and the second column for the price events. We use a two-stage estimation method, which relies on a duration exogeneity assumption. In the first step, we estimate the EACD model by QML with an exponential density function. The results are reported in Table 2. In all of the cases, the specification tests on durations (not reported) lead us to consider an EACD(2,2) model. All of the parameters are significant, and the sum of autoregressive parameters is generally close to one, which confirms the high persistence in durations. The model's performance in capturing the latent structure of durations can be assessed by comparing the autocorrelations of the seasonally adjusted durations (cf. Table 1) with those of the standardized residuals $\widehat{v}_i = \widetilde{x}_i / \widehat{\psi}_i$. In Table 2, we report the p-values of the Ljung-Box test Q-statistics based on 10 and 20 lags, respectively, for the series of standardized residuals. For all of the assets, the null hypothesis of no autocorrelation cannot be rejected at the 95% confidence level, which indicates that the models successfully remove the autocorrelation observed in the original durations.

In the second step, we use the sequence of predicted durations to evaluate the time-varying parameters ω_i, α_i and β_i of the GARCH(1,1) process. More precisely, we use the conditional raw durations, which are defined as $\widehat{\psi}_{n+1}\ \bar{\omega}_d\left(t_n\right)$, because the GARCH parameters depend on the expected (and uncorrected) duration before the next trade (Drost and Nijman, 1993).[12] The structural parameters ω, α and β are estimated by QML (Appendix A). The estimation results are reported in Table 3. These parameters can be interpreted as the parameters of a GARCH model based on regular high-frequency (one-second) data. Note that their estimated values are similar to those obtained for daily data; the ARCH parameter is much smaller than the GARCH parameter and their sum is close to unity. In contrast, for IBM, Ghysels and Jasiak (1998) obtain a relatively small value for the estimated GARCH parameter β, which is close to α. They explain this outcome by the fact that the drift ω_i depends directly on $\widehat{\psi}_i$ and the expected duration absorbs all of the persistence. In our case, the persistence captured by ω_i seems to be insufficient to reproduce the persistence of volatility.[13] From these estimates, we compute the duration-dependent parameters ω_i, α_i and β_i of the time-varying GARCH model (Equation 12) by applying the formulas of temporal aggregation of Drost and Nijman (1993). For example, for BAC the parameters α_i and β_i have means equal to 0.075 and 0.753, respectively, while their values are within the intervals [0.046; 0.076] and [0.343; 0.927]. Finally, this model does not permit the removal of the serial dependence from the tick-by-tick standardized returns $\widehat{\varepsilon}_i$, as shown by the Ljung-Box test Q-statistics. However, the remaining dependencies are much less important than those observed for the raw returns (Table 1).

Out-of-sample VaR and TaR Forecasts: Figure 3 displays the first 3,000 out-of-sample forecasts of the 1%-VaR and 1%-TaR for the S&P 500 and BAC trades on September 23, 2010 and the MSFT trades on September 16.[14] These observations correspond to the time period between 9:30 am and 12:57 am for SPY, until 12:09 am for BAC and until 11:17 for MSFT. First, we observe that the TaR and VaR forecasts accurately capture the volatility and duration dynamics. For BAC, the minimum duration (over this period) before the next trade may occur (with a confidence level of 99%) ranges between 21.45 seconds and 4 minutes and 25 seconds (or 265 seconds). For the same period, the maximum expected loss that will not be exceeded (for a given confidence level of 99%) at the time horizon of the next trade ranges from 0.072% to 0.256%. Second, the frequencies of VaR violations (1% for S&P 500, 0.8% for BAC and 0.7% for MSFT) and TaR violations (1% for S&P 500, 0.9% for BAC and 0.6% for MSFT) are close to the nominal level of 1%, and, more importantly, they are not clustered. The same results are observed for the VaR and TaR forecasts on the price events (Figure C1, Appendix C).

We observe that the TaR and VaR are negatively correlated. When the market risk is low, the

[12] Similar results (available upon request) are obtained when we use the conditional seasonally adjusted duration $\widehat{\psi}_{n+1}$.

[13] Another difference from Ghysels and Jasiak (1998) is that we use seasonally adjusted returns and conditional raw durations in the GARCH model, whereas they consider raw returns and seasonally adjusted durations.

[14] To compare the VaR forecasts with the returns, we display the opposite of VaR, *i.e.*, $-VaR_i(\alpha)$.

duration between two transactions is important. It implies a lower intensity in financial trading activity and, consequently, an important liquidity risk as measured by TaR. However, the correlation between the two measures for BAC (-0.151) and MSFT (-0.111) is not high, as confirmed by the scatter plot in Figure 4. In contrast, the ETF correlation reaches -0.407. This negative correlation is not necessarily caused by the model's structure. Given the values of the GARCH parameters, it is possible to observe an increase in the conditional duration (and thus of TaR) simultaneously with a decrease in volatility (and thus of VaR).[15] This negative relationship revealed by the data confirms the previously mentioned theoretical findings.

So far the analysis compares the VaR with the returns but does not accurately reveal the ultra-high-frequency market risk. A VaR of 0.2% associated with a transaction that will occur in ten seconds does not present the same risk for the investor as the same VaR associated with a transaction that will occur in ten minutes. A solution consists of comparing the ex-ante per-time-unit VaR, which is defined as $VaR_i\left(\alpha\right)/\psi_i$, with the ex-post returns standardized by the conditional duration (Figure 5). The per-time-unit VaR can be interpreted as the maximum expected loss for the next second. Consequently, both measures (total VaR and per-second VaR) provide further insight into high-frequency market risk and may be useful for risk management.

5.2 Backtesting

Traditionally, the quality of an economic variable forecast is assessed by comparing its *ex-post* realization with the *ex-ante* forecasted value. However, this approach is not suitable for VaR and TaR forecasts because the true quantiles of the corresponding distributions are not observable. Therefore, VaR assessment is generally based on the concept of violation (also known as hit or exception). A violation is said to occur if the *ex-post* realization of the return (resp. duration) is more negative (resp. larger) than the *ex-ante* VaR (resp. TaR) forecast. Let us define two binary variables associated with the *ex-post* observation of an α%-VaR violation and a α%-TaR violation:

$$H_i\left(r\right)=\begin{cases}1 & \text{if } r_i<-VaR_i(\alpha)\\ 0 & \text{otherwise}\end{cases}, \tag{21}$$

$$H_i\left(d\right)=\begin{cases}1 & \text{if } x_i>TaR_i(\alpha)\\ 0 & \text{otherwise}\end{cases}. \tag{22}$$

As emphasized by Christoffersen (1998) and Berkowitz, Christoffersen and Pelletier (2011), VaR and TaR forecasts are valid if the violations satisfy the following two hypotheses:[16]

[15] On the contrary, the use of a volatility measure based on the price intensity (as in Engle and Russell, 1998 or Giot, 2005) implies a negative relationship between the volatility and the conditional duration.

[16] The unconditional coverage hypothesis is intuitive. In fact, if the empirical frequency of violations is significantly lower (respectively higher) than the coverage rate α, the risk is overestimated (respectively underestimated). However, the unconditional coverage hypothesis sheds no light on the possible dependence of violations. Therefore, the independence property of violations is an essential one because it is related to the ability of the VaR and TaR models to accurately model the dynamics of the returns and durations. A model that does not satisfy the independence property can result in clusters of violations even if the model has the correct average number of violations.

(i) The unconditional coverage (UC) hypothesis: the probability of an *ex-post* violation must be equal to the α coverage rate:

$$\Pr\left(H_i(z)=1\right)=\mathbb{E}\left(H_i(z)\right)=\alpha, \quad \forall z \in \{r,d\}. \tag{23}$$

(ii) The independence (IND) hypothesis: VaR and TaR violations observed at two different dates for the same coverage rate must be distributed independently. The variable $H_i(z)$ associated with a VaR or TaR violation for the i^{th} trade should be independent of the variable $H_{i-k}(z)$, $\forall k \neq 0$. Thus, past violations should not influence current and future violations.

When the UC and IND hypotheses are simultaneously valid, VaR and TaR forecasts are said to have a correct conditional coverage (CC), and the corresponding violation process is a martingale difference, with:

$$\mathbb{E}\left(\left. H_i(z)-\alpha \right| \mathcal{F}_{t-1}\right)=0, \quad \forall z \in \{r,d\}. \tag{24}$$

Several remarks are in order. First, one could also define a joint violation, *i.e.*, circumstances in which the return and duration exceed the VaR and TaR. However, because VaR violations and TaR violations might not be independent, the probability of observing a joint violation is not necessarily equal to α^2. Therefore, we test the TaR and VaR forecasts independently. Second, the VaR forecasts are not rescaled to fixed-time intervals to apply usual backtesting procedures as in Giot (2005) or Dionne, Duchesne and Pacurar (2009). The summation of the tick-by-tick returns and the VaR forecasts over a fixed-time interval (for instance, five minutes) may obscure the potential misspecification of the VaR model. More fundamentally, the HFR measure is designed to forecast the duration and market risks at the time horizon of the next trade. Consequently, the measure's performance must be assessed for the same time horizon.

However, the use of irregularly spaced returns and VaR forecasts implies that some of the usual backtesting procedures (for a survey, see Campbell, 2007; Hurlin and Pérignon, 2012) cannot be used in this context.[17] Here, we consider three backtests that are compatible with irregularly spaced data. The first is the LR test proposed by Christoffersen (1998). This test is based on a Markov-chain model with two states (violation or no violation). The UC, IND and CC assumptions are tested through parameter restrictions on the transition probability matrix. The UC test corresponds to Kupiec's test (1995) for the percentage of failures, which is also embedded in the regulatory requirements based on the backtesting of daily VaR models.

[17] For instance, the dynamic quantile test proposed by Engle and Manganelli (2004) is no longer adapted. This test consists of testing certain linear restrictions in a linear regression model of the violation variable on a set of lagged explanatory variables (e.g., past violations, returns, VaRs). In our context, the autoregressive structure of the model neglects the irregular dimensions of the data.

The second test is the duration-based LR test proposed by Christoffersen and Pelletier (2004). This test exploits the duration between two consecutive violations.[18] Under the CC hypothesis, this duration follows a geometric distribution with parameter α. Exploiting this property, the authors consider a continuous exponential distribution and propose a LR test for the null hypothesis of CC. A lifetime distribution that nests the exponential distribution (for instance, a Weibull distribution) is specified under the alternative, and the memoryless property of the violations can be tested by means of parameter restrictions.

The third test, proposed by Candelon et al. (2011), also exploits the duration between two consecutive violations. However, this test is based directly on the geometric distribution. The GMM framework, which was proposed by Bontemps and Meddahi (2012) to test for distributional assumptions, is applied to the VaR and TaR forecasts. The test statistic is a J-statistic based on the moments defined by the orthonormal polynomials associated with the geometric distribution.

Backtesting Results: Table 4 reports the results of the backtests for a 1%-VaR (Panel A) and a 1%-TaR (Panel B). These tests are based on a sequence of 1,000, 2,000 or 3,000 out-of-sample one-step-ahead forecasts obtained according to the forecasting procedure described in Section 4. For each test and asset (S&P 500, BAC or MSTF), we distinguish the UC, IND and CC null hypotheses and report the corresponding test p-values. Entries in italics denote a failure of the model at the 95% confidence level. Additionally, we report the empirical frequency of violations for each case.

The results indicate that the intraday VaR and TaR forecasts are satisfactory. First, the empirical frequencies of VaR and TaR violations are statistically not different from the 1% risk level for all of the assets. In all of the cases considered, the p-values associated with the UC tests exceed 5%. These p-values tend to decrease with sample size. For a frequency of violations close to, but different from the 1% level (for instance, 1.01%), the null hypothesis tends to be rejected more often when the sample size increases. Second, the VaR and TaR violations are not clustered, which indicates that the VaR and TaR forecasts accurately capture the volatility and duration dynamics. The IND tests indicate that the null hypothesis of independence between the current and past violations cannot be rejected. The only exception is observed in the VaR for MSFT when we consider a sample size of 3,000 observations. This independence property in the time dimension does not indicate that both violation processes are independent. We observe that the VaR and TaR violations occur for the *same trades* in fewer than 0.05% of the total number of violations for all of the assets. Extreme returns and extreme durations do not generally occur at the same time, which confirms the negative relationship between market risk and trading activity. Third, the CC

[18] Note that we must distinguish the duration between two consecutive violations from the trade duration used to compute the TaR. As a consequence, the first one can be used to backtest both the VaR and the TaR.

tests do not allow us to reject the null hypothesis of a martingale difference both for the VaR and TaR violations, which supports the conclusion that the VaR and TaR forecasts are valid.

Because we consider irregularly spaced risk measures, backtesting a sample of 1,000 observations does not necessarily span the same period for BAC or MSFT (the least and the most traded stocks in our analysis). Additionally, one could be interested in testing the validity of the HFR measures over a fixed period. Table 5 reports the results of the backtests over a period of 30 minutes, one hour and four hours after the opening time on September 23, 2010 (for BAC and S&P 500) and on September 16, 2010 (for MSFT), respectively. For a period of 30 minutes, the backtesting sample includes 609 observations for S&P 500, 225 for BAC and 907 for MSFT. For both the VaR and the TaR, the null hypothesis of UC cannot be rejected at the 95% confidence level, which indicates that the frequency of violations is statistically not different from the level of 1%. Besides, these violations are not clustered. These results confirm that the HFR measure accurately assesses the intraday market and liquidity risks.

Periodicity and Split Point: The previous backtesting procedure was conducted only for the first trading hours of the first forecasting day (September 23, 2010). To determine if the VaR and TaR forecasts remain valid for the next trading days, we propose a backtesting analysis based on a fixed rolling window of one hour over a period of seven days (September 23 to October 1) without re-estimation of the model. Figure 6 displays the p-values of Christoffersen's LR test (1998) for the frequency of violations (UC) of both 1%-TaR and 1%-VaR forecasts for the S&P 500 ETF. The results confirm the validity of the HFR forecasts. Regardless of the day or hour considered over this period, the p-values are larger than 5%, which indicates that the null hypothesis of UC cannot be rejected for the VaR (upper panel of Figure 6) and TaR (middle panel of Figure 6) forecasts. For the first two trading days, the p-values exhibit an inverted-U shape; the probability is relatively small at the beginning and end of the day and larger at midday. This observation highlights the fact that the reliability of the VaR and TaR forecasts is less certain when the trading activity is particularly important and the returns are more volatile. It can also be explained by the number of observations used for each backtest (bottom panel of Figure 6). Because we consider a fixed period of one hour, the number of observations is more important at the beginning and end of the day. Consequently, even if the frequency of violations is constant over the day and assumed to be close to (but different from) the level of 1%, the inverted-U shape of the p-values reflects the increase in the power of the test with the number of observations.

Finally, our backtesting results may depend on how our sample is split into estimation and evaluation periods (Hansen and Timmermann, 2012). To assess the robustness of the results, we propose an analysis based on a rolling window in which the split point between the in-sample and out-of-sample periods moves by 1,000 observations at each iteration between September 22 and October 10. As shown in Figure 7, the rolling window scheme is chosen so that there is no overlapping of the backtesting sample. At each step, the EACD-GARCH parameters are estimated using the past 25,000 observations, and we compute 1,000 out-of-sample VaR and TaR forecasts for BAC. Figure 8 displays the corresponding p-values from Christoffersen's LR test for UC. The figure confirms the robustness of our results because whatever the split point used, the UC hypothesis cannot be rejected for a 95% confidence level.

Figure 7: Rolling Window Scheme

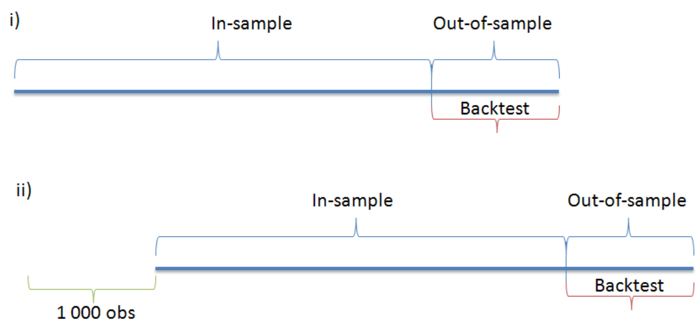

6 Marginal VaR and Instantaneous VaR

In the previous sections, we considered the *total* VaR, which is defined as the maximum expected loss that will not be exceeded under normal conditions for a $1-\alpha$ confidence level at the time horizon of the next trade or microstructure event i. Alternatively, we can define two additional conditional VaR concepts for HFR management: (i) the *per-time-unit* VaR or *marginal* VaR and (ii) the *instantaneous* VaR.

Definition 2 *For a shortfall probability $\alpha \in \,]0,1[$, the per-time-unit α-VaR (or marginal VaR) is defined as follows:*

$$\Pr\left(\frac{r_i}{x_i} < -VaR_i(\alpha; x_i) \mid \mathcal{F}_{i-1}; \mathbb{C}_i\right) = \alpha. \tag{25}$$

The instantaneous α-VaR is defined for a constant time variation Δ:

$$\lim_{\Delta \to 0} \frac{1}{\Delta} \Pr\left(r_{i-1}(\Delta) < -VaR_i(\alpha; \Delta) \mid \mathcal{F}_{i-1}; \mathbb{C}_i\right) = \alpha, \tag{26}$$

where $r_{i-1}(\Delta) = p(t_{i-1} + \Delta) - p(t_{i-1})$ denotes the instantaneous return measured between t_{i-1} and $t_{i-1} + \Delta$.

Let us consider the case in which i denotes a trade and the time is measured in seconds. The per-time-unit VaR corresponds to the maximum expected loss for the next second. The instantaneous VaR corresponds to the maximum expected loss for the time interval between t_{i-1} and $t_{i-1}+\Delta$, for an infinitesimal time increase Δ.

Marginal VaR The marginal VaR depends on the conditional marginal volatility or the conditional volatility per-time-unit. In the case of a trade event ($\mathbb{C}_i = \varnothing$), the marginal volatility (Gerhard and Hautsch, 2002) is defined as follows:

$$\sigma_i^2 = \underset{r,x}{\mathbb{E}}\left(\left. \frac{r_{c,i}^2}{x_i} \right| \mathcal{F}_{i-1} \right). \tag{27}$$

where $r_{c,i} = r_i - \mathbb{E}\left(r_i | \, \mathcal{F}_{i-1} \right)$ denotes the demeaned return process. To our knowledge, no dynamic model has been proposed for σ_i^2. Such a model would require the consideration of a bivariate process that accounts for the stochastic nature of the current returns and durations (conditionally on $\mathcal{F}_{i-1}$) to derive the conditional distribution of their ratio. In contrast, most of the volatility per-time-unit models consider an information set $\mathcal{G}_{i-1} = \{x_j, p_j, z_j \,, j \leq i-1; x_i\}$ that includes the current duration x_i (Engle, 2000; Meddahi, Renault and Werker, 2005). Subsequently, the conditional volatility per-time-unit becomes:

$$\sigma_i^2 = \frac{1}{x_i} \underset{r}{\mathbb{E}}\left(r_{c,i}^2 \right| \mathcal{G}_{i-1} \left. \right) = \frac{1}{x_i} Var(r_i | \mathcal{G}_{i-1}). \tag{28}$$

When the durations are seasonally adjusted, the corresponding $\mathcal{G}_{i-1}$-conditional marginal VaR can be defined as follows:[19]

$$VaR_i(\alpha; x_i; \mathcal{G}_{i-1}) = -\frac{\mu_i}{x_i} - \sigma_i \; F_{\varepsilon}^{-1}(\alpha) \sqrt{\frac{\bar{\omega}_r(t_i)}{x_i}}, \tag{29}$$

where $F_{\varepsilon}^{-1}(\alpha)$ denotes the quantile of the standardized and adjusted returns $\varepsilon_i = \bar{\omega}_r(t_i) \; r_{c,i}/\sigma_i$. Modeling the $\mathcal{G}_{i-1}$-conditional marginal volatility enables the identification of the $\mathcal{G}_{i-1}$-conditional total volatility as $x_i \sigma_i^2$. This quantity differs from the $\mathcal{F}_{i-1}$-conditional total volatility h_i defined in the equation (10). Consequently, we can also derive a ($\mathcal{G}_{i-1}$-conditional) total VaR, denoted $VaR_i(\alpha; \mathcal{G}_{i-1})$, as:

$$VaR_i(\alpha; \mathcal{G}_{i-1}) = -\mu_i - \sigma_i \; \sqrt{x_i} \; F_{\varepsilon}^{-1}(\alpha) \sqrt{\bar{\omega}_r(t_i)}. \tag{30}$$

This VaR corresponds to the IVaR proposed by Dionne, Duchesne and Pacurar (2009). However, Dionne, Duchesne and Pacurar use a simulation-based method to infer the VaR at any fixed-time horizon for backtesting purposes. Thus, two differences arise between the IVaR and the HFR-VaRs defined in Equations (11) and (30). First, the IVaR is not forecast for each trade i (or any other

[19] As for the total VaR, the $\mathcal{G}_{i-1}$-conditional marginal VaR and total VaR can be forecast for each trade, price event, volume event by adapting the information set $\mathbb{C}_i$ and the interpretation of the index i.

microstructure event) but for a regular time interval even if it is computed using tick-by-tick data and adapted to irregular time intervals. Second, the IVaR is conditional on the current trade duration x_i, which is not the case for the $\mathcal{F}_{i-1}$-conditional VaR defined in the equation (11).

When given $\mathcal{G}_{i-1}$, various models for σ_i^2 can be considered. We propose to consider the simplest model, *i.e.,* the Ultra High-Frequency (UHF) GARCH model proposed by Engle (2000). The conditional volatility per-time-unit is then assumed to follow a simple GARCH$(1,1)$ equation:[20]

$$\sigma_i^2 = \omega + \alpha\left(\widetilde{r}_{c,i-1}/\sqrt{\widetilde{x}_{i-1}}\right)^2 + \beta\sigma_{i-1}^2 + \gamma\widetilde{x}_i, \tag{31}$$

where $\widetilde{r}_{c,i}$ and $\widetilde{x}_i$ represent the seasonally adjusted demeaned return and duration, respectively. Other models could be used. Dionne, Duchesne and Pacurar (2009) consider an EGARCH-type of model in which the log-volatility depends on the current duration, past returns and past durations.[21] However, Meddahi, Renault and Werker (2005) argue that it is better to model the variance per-time-unit instead of the total variance over the next event because total variances are primarily influenced by the associated duration. They propose an AR(1) model for σ_i^2 with time-varying parameters that depend on x_i. This specification is based on an exact discretization of a continuous-time stochastic volatility process observed at irregularly spaced times. The time-varying feature of the parameters contrasts with the specification (Equation 31) used by Engle (2000) who does not consider the effects of temporal aggregation on the model's parameters.

Instantaneous VaR Finally, an instantaneous VaR can be forecast for each event. Following the notations of Gerhard and Hautsch (2002), the instantaneous volatility of the price change is defined for a constant and infinitesimal time increase:[22]

$$\widetilde{\sigma}_{i+\Delta}^2 = \lim_{\Delta \to 0} \frac{1}{\Delta} \underset{r}{\mathbb{E}}\left(r_{c,i-1}^2(\Delta) \middle| \mathcal{F}_{i-1}\right), \tag{32}$$

where $r_{i-1}(\Delta) = p(t_{i-1}+\Delta) - p(t_{i-1})$ and $r_{c,i-1}(\Delta) = r_i(\Delta) - \mathbb{E}\left(r_i(\Delta)\middle|\mathcal{F}_{i-1}\right).$ The corresponding VaR for the instantaneous demeaned return process can be expressed as follows:

$$VaR_{c,i}(\alpha;\Delta) = -\,\widetilde{\sigma}_{i+\Delta}^2\; F_{\varepsilon(\Delta)}^{-1}(\alpha), \tag{33}$$

where $F_{\varepsilon(\Delta)}(.)$ denotes the cdf of the innovations process associated with $r_i(\Delta)$.

[20] Engle (2000) proposes various other specifications that depend on the current duration, for instance $\sigma_i^2 = \omega + \alpha\left(r_{c,i-1}/\sqrt{x_{i-1}}\right)^2 + \beta\sigma_{i-1}^2 + \gamma x_i$.

[21] For $\gamma = 1$ (see Dionne, Duchesne and Pacurar, 2009, page 780) their total volatility model can be rewritten as a model for the logarithm of the volatility per time-unit that encompasses the UHF GARCH.

[22] By considering arithmetic returns, Engle and Russell (1998) define the instantaneous volatility as follows:

$$\widetilde{\sigma}_i^2 = \lim_{\Delta\to 0}\frac{1}{\Delta}\mathbb{E}_i\left(\left(p_{i+\Delta}-p_i\right)/p_i\right)^2 = \frac{1}{p_i^2}\lim_{\Delta\to 0}\frac{1}{\Delta}\mathbb{E}_i\left(p_{i+\Delta}-p_i\right)^2$$

This expression resembles the expression obtained for logarithmic returns and log prices (except for the term p_i^2).

Using tick-by-tick data, the instantaneous volatility is derived directly from the expected duration or the intensity. However, trade intensity may provide relatively little information on the price volatility. For instance, if the price fluctuates between two values every second, the asset exhibits a large trading intensity but no price change (Giot, 2001). Consequently, the instantaneous volatility is generally derived from the *price* durations, and the corresponding VaR is only reported for price events.

Thinning the market point process for trades with respect to a minimum change in price enables us to derive the close relationship between price intensity and instantaneous intraday volatility (Engle and Russell, 1998). Numerous trades over a short duration result in a high trading intensity, which intuitively could generate more frequent price revisions. Let $N(t)$ be a counting variable equal to the total number of price events that have occurred by time t. The conditional price intensity function is defined as follows:

$$\lambda(t_i | \ \mathcal{F}_{i-1}; \mathbb{C}_i) = \lim_{\Delta \to 0} \frac{1}{\Delta} \Pr\left(N(t_i + \Delta) > N(t_i) \mid \mathcal{F}_{i-1}; \mathbb{C}_i\right), \tag{34}$$

where $\mathbb{C}_i = \{|r_i| \geq C\}$. As shown by Engle and Russell, the instantaneous volatility can be expressed as a linear function of the conditional intensity:[23]

$$\widetilde{\sigma}^2_{i+\Delta} = \lambda(t_i | \ \mathcal{F}_{i-1}; \mathbb{C}_i) \ C^2. \tag{35}$$

If an EACD model is applied to price durations, this expression becomes:

$$\widetilde{\sigma}^2_{i+\Delta} = \frac{C^2}{\psi_i}. \tag{36}$$

The intuition is the following. If the price change (observed over a time interval Δ) is smaller than the threshold C, we assume that there is no price change. In fact, the change is unobservable because the thinned point process considers only the trades for which $|r_i| \geq C$. In other cases, Engle and Russell assume that the instantaneous price change is precisely equal to C. The price change can then be viewed as a binomial process; the probability of observing a return equal to C is equal to $\lambda(t_i | \ \mathcal{F}_{i-1}; \mathbb{C}_i)$, otherwise there is no change. Expression (35) corresponds to the variance of this binomial process.

Giot (2005) uses a similar approach to compute an intraday VaR for price events. Because the innovations of the instantaneous returns cannot be identified, Giot substitutes $F^{-1}_{\varepsilon(\Delta)}(\alpha)$ with the empirical α-quantile of the standardized total returns $r_{c,i}/\widetilde{\sigma}_{i+\Delta}$. Thus, the corresponding VaR is not an instantaneous VaR as defined in the equation (33), because it combines the instantaneous volatility and the quantile of the innovations of the total returns. Giot compares these estimates

[23] For arithmetic returns, Engle and Russell (1998) obtain a similar expression with $\widetilde{\sigma}^2_{i+\Delta} = \lambda(t_i | \ \mathcal{F}_{i-1}; \mathbb{C}_i) \ C^2 / p^2_{i-1}$.

with the usual GARCH based estimates obtained by interpolating the prices on a grid of regularly spaced time points. The backtests demonstrate that GARCH-based predictions are better measures of risk than the ACD-based predictions. However, the two VaR approaches are not directly comparable. First, the GARCH approach does not consider the irregular timing of price events. Second, the time horizon of the GARCH-based VaR is one time period, whereas the time horizon of the instantaneous VaR is only the infinitesimal time interval Δ. A solution that facilitates the comparison of these measures involves deriving the total volatility (and VaR) from the instantaneous volatility. The total volatility of the price returns derived from an EACD model can be expressed as follows (Tse and Yang, 2012):

$$h_i = \int_{t_{i-1}}^{t_i} \widetilde{\sigma}^2_{i+\Delta}\ d\Delta = C^2 \left(\frac{t_i - t_{i-1}}{\psi_i} \right), \tag{37}$$

Giot (2005) considers a simpler approach and computes the average of the instantaneous VaR measured in a regular sampling between two dates.

Figure 9 shows the three types of 1%-VaR (*i.e.*, the total, the instantaneous and the time per-unit) for the S&P 500 for the period between September 23 at 9:30 am and September 24 at 10:00 am. For comparison, all of the VaRs are computed for price events with a threshold C equal to 0.0075%. The total VaR is logically larger than the per-time-unit VaR, which is generally larger than the instantaneous VaR. The total VaR measures the maximum expected loss that will be exceeded for a level of risk of 1% at the time horizon of the next price change, whereas the marginal VaR is defined for the time horizon of one second. Because the average price duration is equal to 8.8 seconds (Table 1), the total VaR is approximately eight times larger than the marginal VaR. The instantaneous VaR, defined for an infinitesimal time increase, should always be smaller than the marginal VaR. Here, the instantaneous VaR is based on the empirical quantile of the standardized per-time-unit returns $r_{c,i}/\left(x_i \widetilde{\sigma}_{i+\Delta}\right)$, and not on the standardized returns $r_{c,i}/\widetilde{\sigma}_{i+\Delta}$, as in Giot (2005). However, the difference between the instantaneous and marginal VaRs is smaller than it would be if we could identify the innovation of the instantaneous return process. Beyond their values, which are not directly comparable, the three types of VaR exhibit strong correlations, even using different computation methods (UHF-GARCH versus EACD). The correlation between the marginal and the total VaR is equal to 0.866. The correlation coefficients between the two other combinations of VaR measures are slightly smaller but still high (0.714 for the correlation between the instantaneous VaR and the total VaR and 0.711 between the instantaneous VaR and time per-unit VaR). This result confirms the robustness of our risk measures.

7 Conclusion

In this paper, we propose an intraday High Frequency Risk (HFR) measure specifically designed for HFR management. The HFR measure is a joint conditional measure of market risk and liquidity risk for irregularly spaced high-frequency data based on two well-known measures: VaR and TaR. These risk measures are computed and evaluated at each transaction for the time horizon of the next trade. The VaR is hence not rescaled at fixed-time intervals for backtesting purposes. One advantage of the HFR measure is that it can be extended to any type of market microstructure event (e.g., price events or volume events) in contrast to the IVaR proposed by Giot (2005), which is relevant only for price events. Additionally, our forecasting procedure complies with HFR management practices. When forecasting the VaR for the next trade, we do not assume that the duration prior to this trade is known, in contrast to Dionne, Duchesne and Pacurar (2009). Thus, TaR and VaR are estimated with an EACD-GARCH model (Ghysels and Jasiak, 1998), in which the parameters of the volatility model depend on the expected durations.

In addition, we also distinguish three concepts of intraday VaR: the total, marginal (or per-time-unit) and instantaneous VaRs. These different VaRs provide a complementary view of the market risk and may be used in HFR management. Finally, we propose a backtesting procedure specifically designed to assess the validity of the VaR and TaR forecasts for each microstructure event. We test for the frequency of VaR and TaR violations and for the independence of these violations.

The HFR measure is empirically applied to two stocks (BAC and MSFT) and an ETF based on the S&P 500 index. We prove that the VaR and TaR accurately capture the volatility and duration dynamics of these three assets. The backtests show that the VaR and TaR forecasts are valid because their violations occur with a frequency close to the level of risk and they are not clustered. The VaR and TaR out-of-sample forecasts remain valid throughout the day and the week, although the model is not re-estimated. In addition, these results are robust with respect to the choice of the split point between in-sample and out-of-sample periods.

A Appendix A: EACD-GARCH model

The ACD-GARCH, proposed by Ghysels and Jasiak (1998), is a time varying coefficient GARCH model, where the durations between transactions determine the dynamic of the parameters. The parameter behavior is described by the temporal aggregation formulas of a weak GARCH proposed by Drost and Nijman (1993) and Drost and Werker (1996).

Let us denote $x_i = t_i - t_{i-1}$ the duration between two consecutive trade times and $\psi_i = \mathbb{E}(x_i|\mathcal{F}_{i-1})$ the conditional expected duration given the information set $\mathcal{F}_{i-1} = \{x_j, p_j, z_j\}_{j=1}^{i-1}$ (see section 3.1). Using the temporal aggregation formula of Drost and Nijman (1993) for the case of a weak GARCH process and a flow variable, Ghysels and Jasiak propose an ACD-GARCH model for the conditional volatility h_i of returns defined as follows:

$$h_i = \omega_{i-1} + \alpha_{i-1} r_{i-1}^2 + \beta_{i-1} h_{i-1}, \tag{38}$$

where the parameters w_{i-1}, α_{i-1}, β_{i-1} are functions of the expected duration ψ_i:

$$\omega_{i-1} = \psi_i \ \omega \frac{1 + (\alpha + \beta)^{\psi_i}}{1 - (\alpha + \beta)}, \tag{39}$$

$$\alpha_{i-1} = (\alpha + \beta)^{\psi_i} - \beta_{i-1}, \tag{40}$$

and $|\beta_{i-1}| < 1$ is the solution of the quadratic equation:

$$\frac{\beta_{i-1}}{(1 + \beta_{i-1}^2)} = \frac{a(\alpha, \beta, \kappa, \psi_i)(\alpha + \beta)^{\psi_i} - b(\alpha, \beta, \psi_i)}{a(\alpha, \beta, \kappa, \psi_i)\left((\alpha + \beta)^{2\psi_i}\right) - 2b(\alpha, \beta, \psi_i)}, \tag{41}$$

$$\begin{aligned} a(\alpha, \beta, \kappa, \psi_i) \quad = \quad & \frac{\psi_i (1 - \beta)^2 + 2\psi_i (\psi_i - 1)(1 - \alpha - \beta)^2 \times \left(1 - (1 - \alpha - \beta)^2 + \alpha^2\right)}{(\kappa - 1)(1 - \alpha - \beta)^2} \\ & + \frac{4c(\alpha, \beta, \psi_i)}{(1 - \alpha - \beta)^2}, \end{aligned} \tag{42}$$

$$b(\alpha, \beta, \psi_i) = \frac{\left(\alpha\left(1 - (\alpha + \beta)^2\right) + \alpha^2(\alpha + \beta)\right)\left(1 - (\alpha + \beta)^{2\psi_i}\right)}{1 - (\alpha + \beta)^2}, \tag{43}$$

$$c(\alpha, \beta, \psi_i) = \left(\psi_i (1 - \alpha - \beta) - 1 + (\alpha + \beta)^{\psi_i}\right) \times \left(\alpha\left(1 - (\alpha + \beta)^2\right) + \alpha^2(\alpha + \beta)\right). \tag{44}$$

where κ denotes the kurtosis of the returns. The parameters ω, α, β correspond to the GARCH model defined for the returns sampled at a regular frequency of one second. The time varying parameters ω_{i-1}, α_{i-1}, β_{i-1} correspond to the weak GARCH process obtained for the aggregated returns over a period of ψ_i seconds.

Let $\theta = (\theta^d, \theta^r)$ be the complete parameter vector, with θ^d the subvector of parameters of the ACD model and θ^r a subvector containing the parameters of the conditional mean and variance for the unequally time-spaced returns. As shown by Ghysels and Jasiak (1998), there are many ways to estimate the ACD-GARCH model. In this article we adopt a sequential procedure by estimating in a first step an ACD(p,q) model as in Engle and Russell (1996) by using in a second step the sequence of expected duration estimates $\psi_i(\hat{\theta^d})$ to estimate the weak GARCH model along the lines of Ghysels and Jasiak (1998).

B Appendix B: Empirical distribution of absolute returns

Figure B1: Empirical distribution of absolute returns

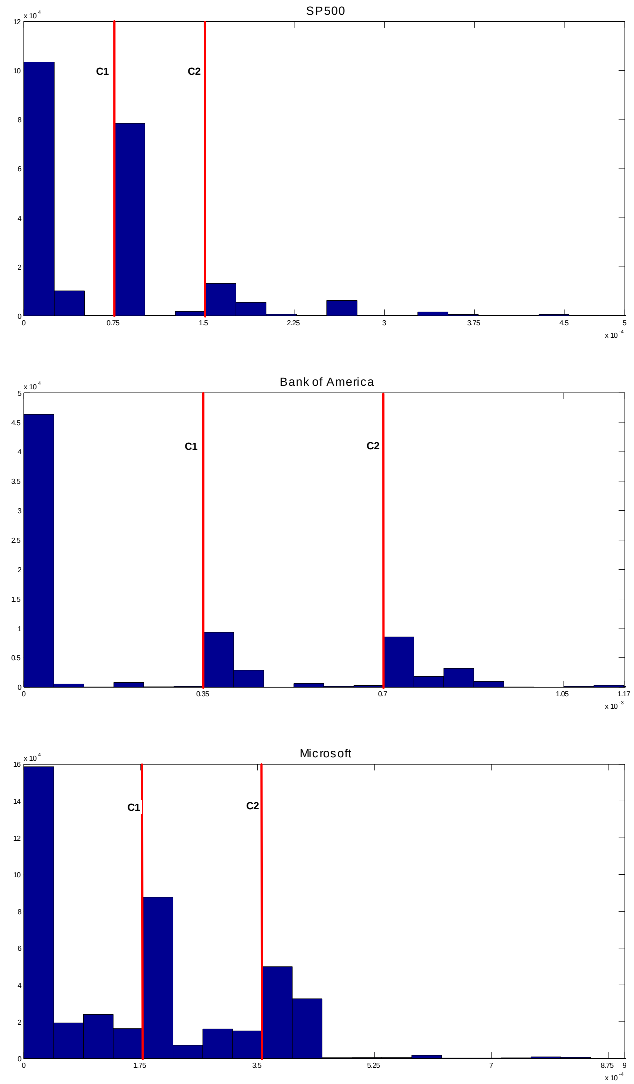

Note: This figure reports the empirical distribution of absolute returns for each of the three stocks under analysis. The histogram is computed with 100 classes, but to better observe the price changes we report here a zoom on the first 20 classes. The red vertical lines mark the thresholds C_1 and C_2 used to define price events.

C Appendix C: HFR forecasts for price events

Figure C1: VaR and TaR out-of-sample forecasts (price events)

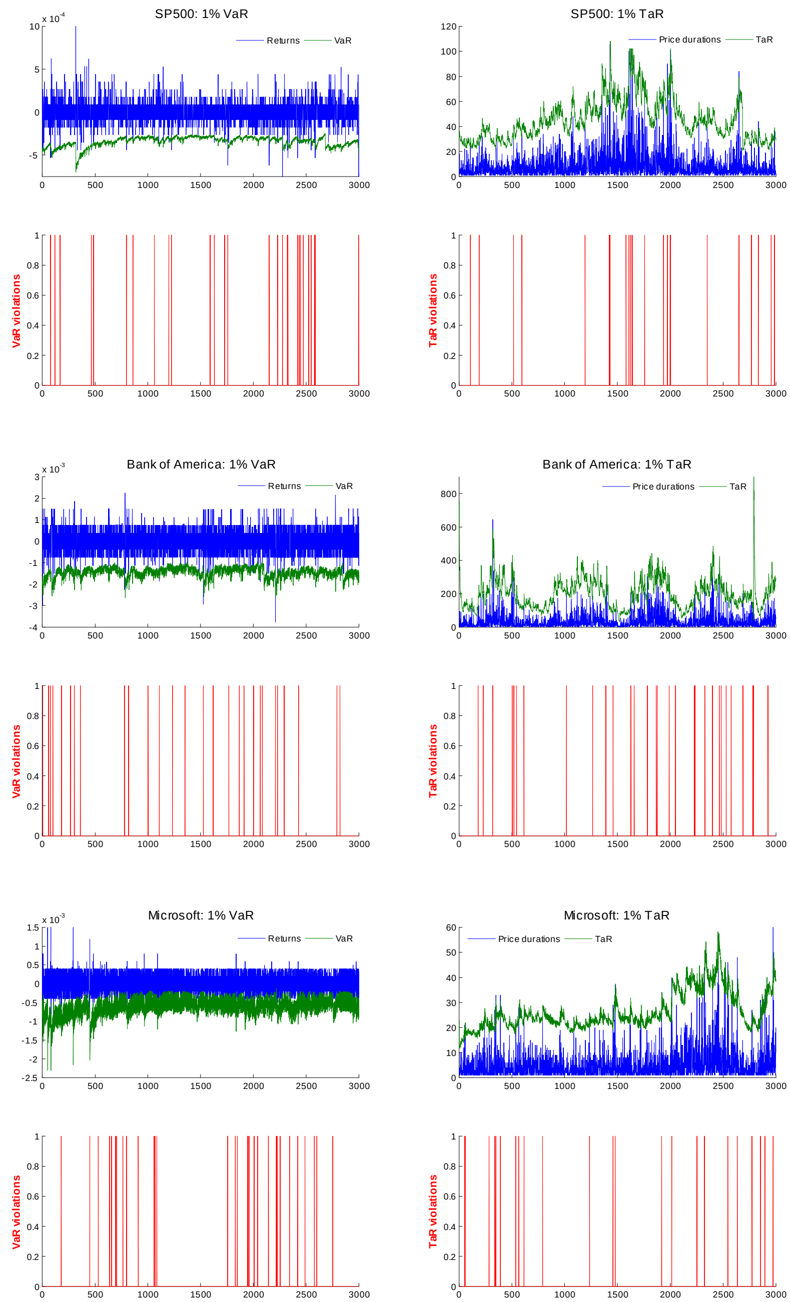

Note: The left panel of each figure displays the forecasts of the intraday 1%-VaR, the ex-post returns, as well as the corresponding violations (in red). The right panel displays the intraday 1%-TaR, the durations and the corresponding violations. The analysis is performed for price events. The out-of-sample period starts on September 16, 2010 at 9:30 am for Microsoft and on September 23, 2010 at 9:30 am for S&P 500 and Bank of America.

References

[1] Adrian, T., Brunnermeier, M., 2011. CoVaR. Working Paper, Princeton University and Federal Reserve Bank of New York.

[2] Ahlstedt, J., Villysson, J. 2012. High-frequency trading, Working Paper.

[3] Anatolyev, S., Shakin, D., 2007. Trade Intensity in the Russian Stock Market: Dynamics, Distribution and Determinants. Applied Financial Economics 17, 87-104.

[4] Andersen, T., Bollerslev, T., 1997. Intraday Periodicity and Volatility Persistence in Financial Markets, Journal of Empirical Finance, 4, 115-158.

[5] Andersen, T., Bollerslev, T., 1998. Deutsche Mark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies, Journal of Finance, 53, 219-265.

[6] Barndorf-Nielsen, O., Hansen, P.R., Lunde, A., Shephard, N., 2009. Realized Kernels in Practice: Trades and Quotes, Econometrics Journal, 12, 1-32.

[7] Bauwens, L., Giot, P., 2000. The Logarithmic ACD Model: an Application to the Bid-Ask Quote Process of Three NYSE Stocks, Annales d'Economie et de Statistique, 60, 11-149.

[8] Bauwens, L., Hautsch,N., 2009. Modelling Financial High Frequency Data Using Point Processes, in Handbook of Financial Time Series, Springer, 953-979.

[9] Berkowitz, J., Christoffersen, P., Pelletier, D., 2011. Evaluating Value-at-Risk Models with Desk-Level Data, Management Science, 57, 2213-2227.

[10] Biais. B., 2011. High-frequency trading, European Institute of Financial Regulation.

[11] BlackRock. 2011. ETF Landscape – Industry Highlights – End Q3 2011.

[12] Bontemps, C., Meddahi, N., 2012. Testing Distributional Assumptions: A GMM Approach, Journal of Applied Econometrics, 27 (6), 978-1012.

[13] Campbell, S., 2007. A Review of Backtesting and Backtesting Procedures, Journal of Risk, 9, 1-18.

[14] Candelon, B., Colletaz, G., Hurlin, C., Tokpavi, S., 2011. Backtesting Value-at-Risk: A GMM Duration-based Test, Journal of Financial Econometrics, 9 (2), 314-343.

[15] Christoffersen, P. F., 1998. Evaluating Interval Forecasts, International Economic Review, 39, 841-862.

[16] Christoffersen, P. F., Pelletier, D., 2004. Backtesting Value-at-Risk: A Duration-Based Approach, Journal of Financial Econometrics, 2 (1), 84-108.

[17] Coroneo L., Veredas, D., 2012. A Simple Two-Component Model for the Distribution of Intraday Returns, the European Journal of Finance, 18 (9), 775-797.

[18] Daley D.J., Vere-Jones D., 2007. An Introduction to the Theory of Point Processes: General Theory and Structure, Springer.

[19] Diamond, D. W., Verrecchia, R. E., 1987. Constraints on Short-selling and Asset Price Adjustment to Private Information, Journal of Financial Economics 18, 277-311.

[20] Dionne, G., Duchesne P., Pacurar, M., 2009. Intraday Value at Risk (IVaR) Using Tick-by-Tick Data with Application to the Toronto Stock Exchange, Journal of Empirical Finance, 16, 777-792.

[21] Drost, F.C., Nijman, T.E., 1993. Temporal Aggregation of GARCH Processes, Econometrica, 61, 909-927.

[22] Drost, F.C., Werker, B.J.M., 1996. Closing the GARCH gap: Continuous Time GARCH Modelling, Journal of Econometrics, 74, 31-57.

[23] Drost, F.C., Werker, B.J.M., 2004. Semiparametric Duration Models, Journal of Business and Economic Statistics, 22, 40-50.

[24] Dufour A, Engle R., 2000. Time and the Price Impact of a Trade, Journal of Finance, 6, 2467-2498.

[25] Easley, D., O'Hara, M., 1992. Time and the Process of Security Price Adjustment, Journal of Finance 47, 577-605.

[26] Engle, R.F., 2000. The Econometrics of Ultra-high-frequency Data, Econometrica, 1, 1-22.

[27] Engle R.F., Hendry D., Richard J.F., 1983. Exogeneity, Econometrica, 51 (2), 277-304.

[28] Engle, R. F., Manganelli, S., 2004. CAViaR: Conditional Autoregressive Value-at-Risk by Regression Quantiles, Journal of Business and Economic Statistics, 22, 367-381.

[29] Engle, R. F., Russell, J. R., 1997. Forecasting the Frequency of Changes in Quoted Foreign Exchange Prices with Autoregressive Conditional Duration Models, Journal of Empirical Finance, 4, 187-212.

[30] Engle, R. F., Russell, J. R., 1998. Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data, Econometrica, 66, 1127-1162.

[31] Engle, R. F., Russell, J. R., 2005. A Discrete-State Continuous-Time Model of Financial Transactions Prices and Times: The Autoregressive Conditional Multinomial - Autoregressive Conditional Duration Model, Journal of Business and Economic Statistics, 23, 166-180.

[32] Gerhard F., Hautsch N., 2002. Volatility Estimation on the Basis of Price Intensities, Journal of Empirical Finance, 9, 57-89.

[33] Ghysels, E., Gouriéroux C. and J. Jasiak, 2004. Stochastic Volatility Duration Models, Journal of Econometrics, 119, 413-433.

[34] Ghysels, E., Jasiak, 1998. GARCH for Irregularly Spaced Financial Data: the ACD-GARCH model, Studies in Nonlinear Dynamics and Econometrics, 2, 133-149.

[35] Giot, P., 2001. Time Transformations, Intraday Data, and Volatility Models, Journal of Computational Finance, 4, 31-62.

[36] Giot, P., 2005. Market Risk Models for Intraday Data, European Journal of Finance, 309-324.

[37] Gomber, P., Arndt, B., Lutat, M., Uhle, T., 2013. High-Frequency Trading, Working Paper.

[38] Gouriéroux, C., Montfort, A., Trognon, A., 1984. Pseudo-Maximum Likelihood Methods: Theory", Econometrica, 52, 681-700.

[39] Grant, J., 2010. High-frequency Trading: Up against a Bandsaw. Financial Times. Sept. 10, 2010.

[40] Hansen, P.R., Timmermann, A., 2012. Choice of Sample Split in Out-of-Sample Forecast Evaluation, Working Paper.

[41] Harris, L., 1986. A Transaction Data Study of Weekly and Intradaily Patterns in Stock Returns, Journal of Financial Economics, 16, 99-117.

[42] Hecq, A., Laurent, S., Palm, F. 2012. Common Intraday Periodicity, Journal of Financial Econometrics, 10, 325-353.

[43] Hurlin, C., Pérignon, C., 2012. Margin Backtesting, Review of Futures Market, 20, 179-194.

[44] Jorion, P., 2007. Value at Risk: The New Benchmark for Managing Financial Risk, McGraw-Hill, 3rd Edition.

[45] Khozam R., Tham W.W., 2013. Execution Risk in High-Frequency Arbitrage. Management Science, 58, 2131-2149.

[46] Kupiec, P., 1995. Techniques for Verifying the Accuracy of Risk Measurement Models, Journal of Derivatives, 3, 73-84.

[47] 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.

[48] Meddahi, N., Renault, E., Werker, B., 2005. GARCH and Irregularly Spaced Data, Economics Letters, 90, 200-204.

[49] Tse, Y., Yang, T.T., 2012. Estimation of High-Frequency Volatility: An Autoregressive Conditional Duration Approach, Journal of Business & Economic Statistics, 30(4), 2012.

[50] Wood, R.A., Mclnish, T.H., Ord, J.K., 1985. An Investigation of Transaction Data for NYSE Stocks, Journal of Finance, 25, 723-739.

Figure 2: Intraday periodicity

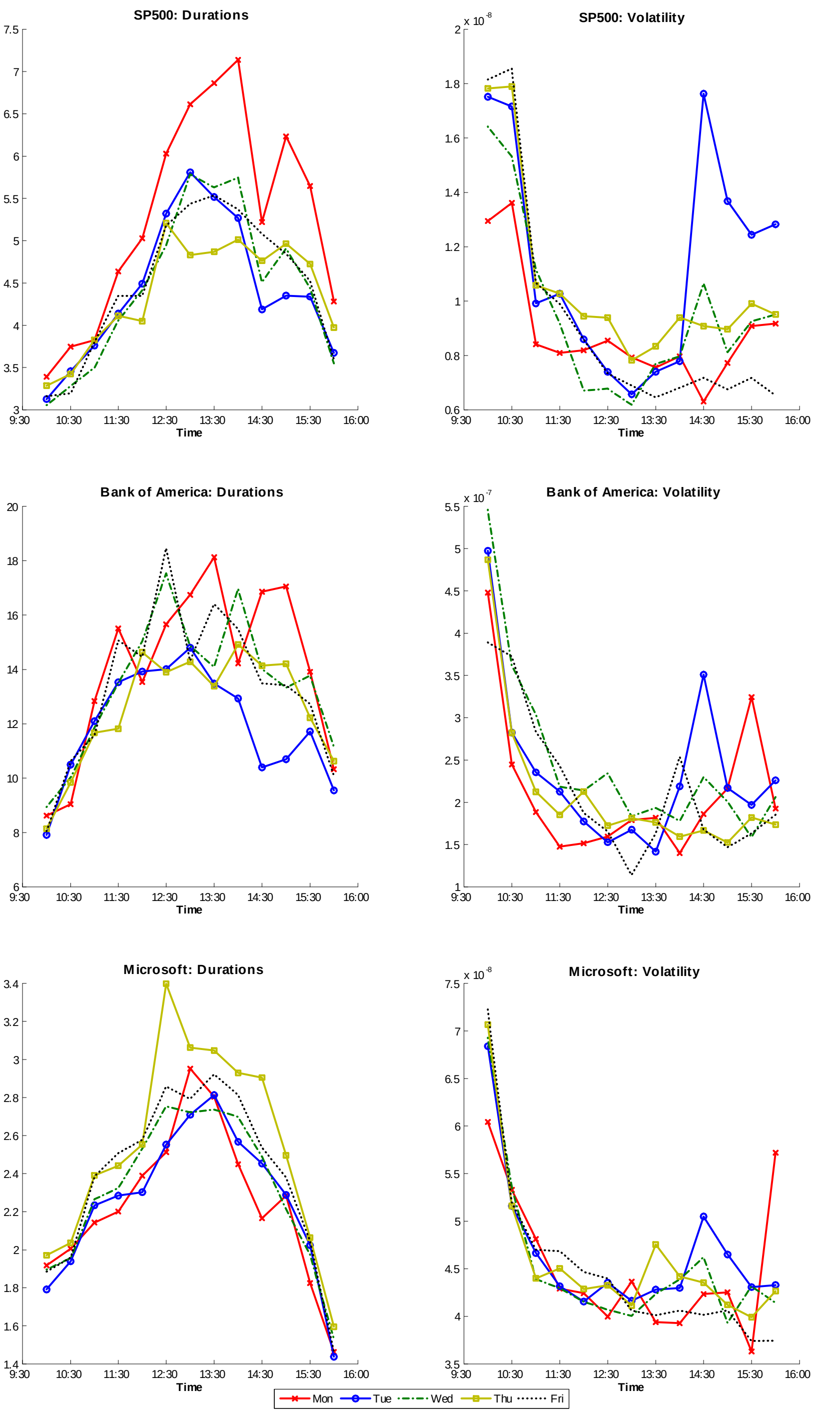

Note: This figure displays the deterministic seasonal component for S&P500, Bank of America and Microsoft. It emphasizes the presence of intraday periodicity in average durations and returns across each trading day of the week.

Figure 3: VaR and TaR out-of-sample forecasts

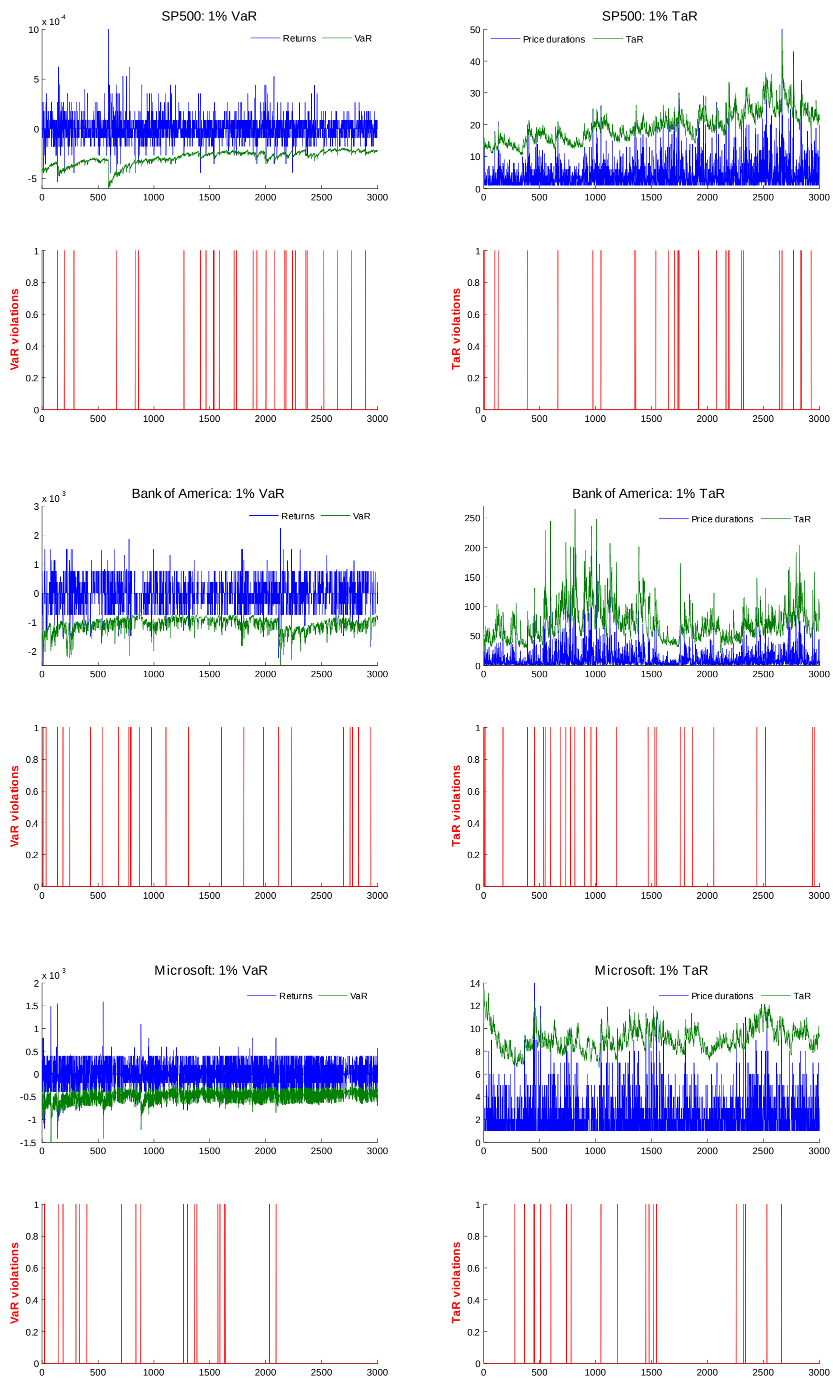

Note: The left panel of each figure displays the forecasts of the intraday 1%-VaR, the ex-post returns, as well as the corresponding violations (in red). The right panel displays the intraday 1%-TaR, the durations and the corresponding violations. The analysis is performed for trade events. The out-of-sample period starts on September 16, 2010 at 9:30 am for Microsoft and on September 23, 2010 at 9:30 am for S&P 500 and Bank of America.

Figure 4: Relationship between intraday VaR and TaR

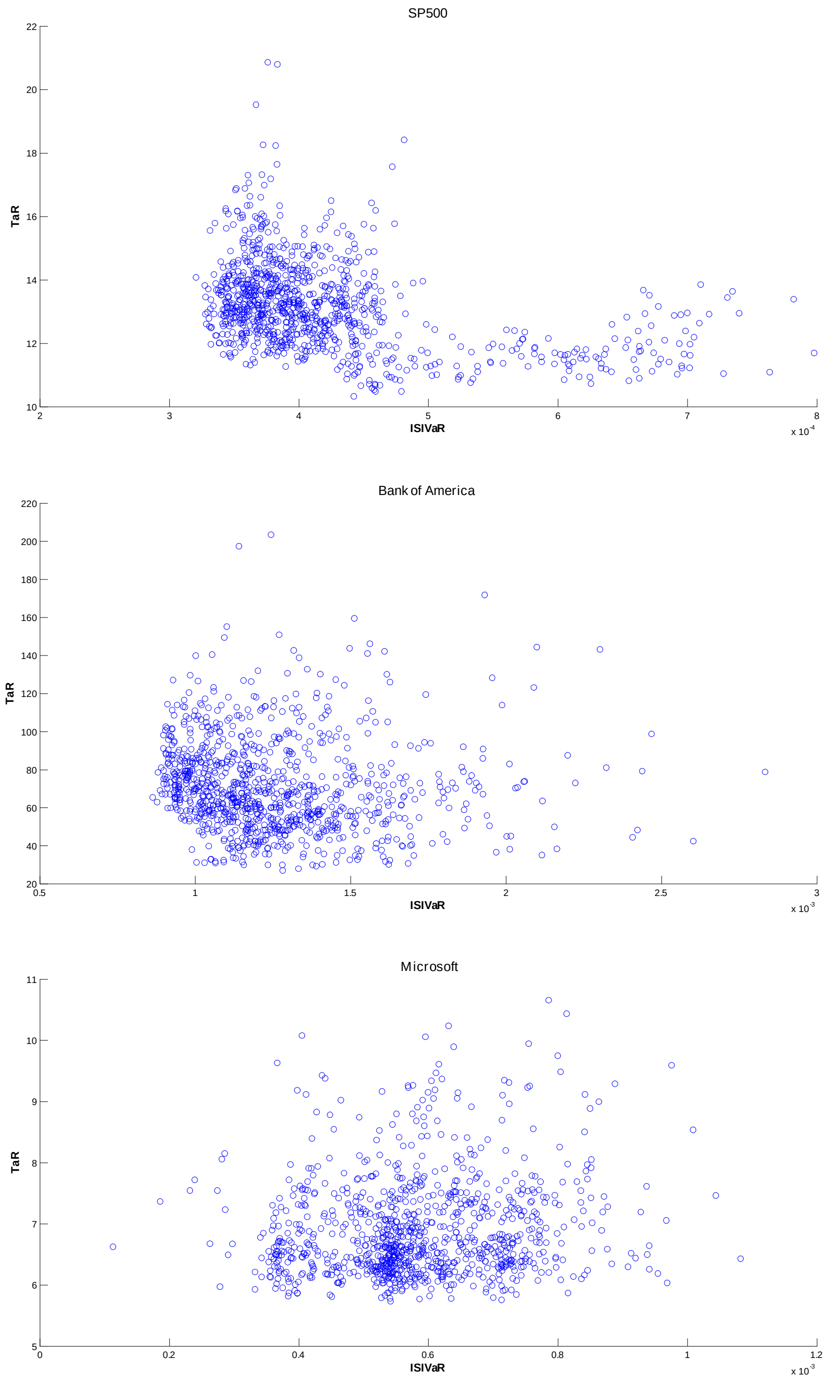

Note: This scatter plot illustrates the relation between the intraday 1%-VaR and 1%-TaR forecasts for each of the three assets under analysis. The sample period ranges from September, 23, 2010 to October 29, 2010 for S&P 500 and Bank of America and from September 16 to October 29, 2010 for Microsoft.

Figure 5: Per-time unit VaR forecasts

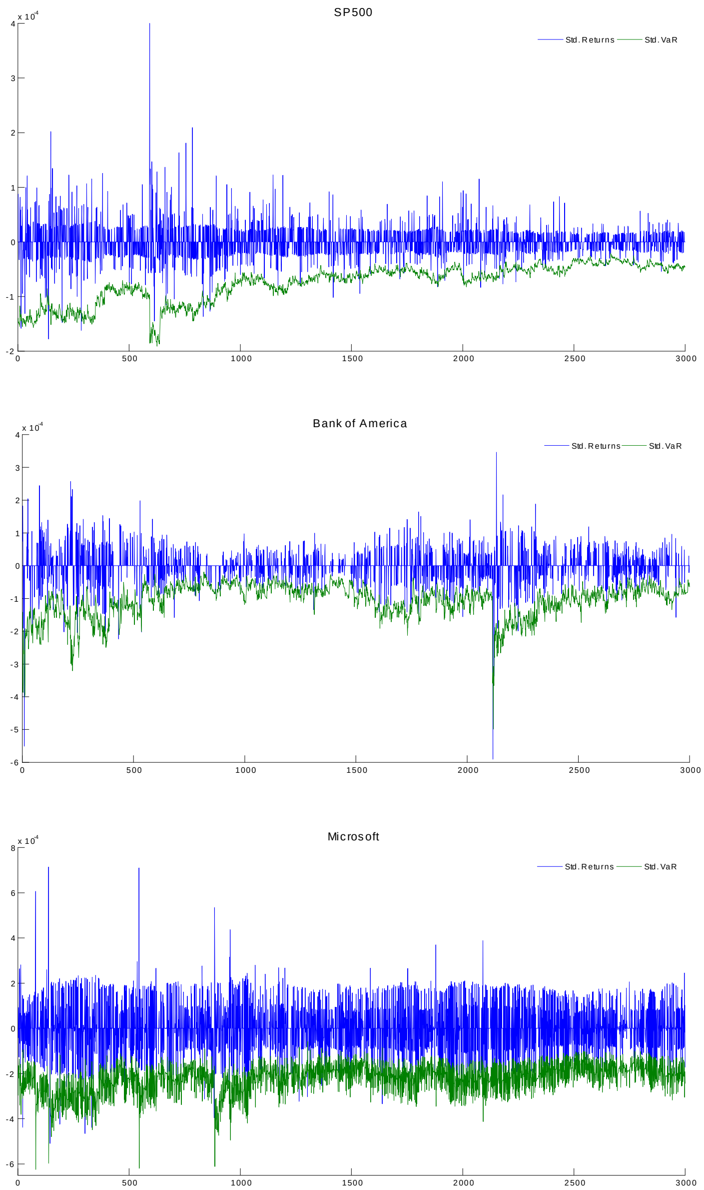

Note: This figure compares the ex-ante per-time-unit VaR to the ex-post returns standardized by the conditional duration. The out-of-sample period starts on September 16, 2010 at 9:30 am for Microsoft and on September 23, 2010 at 9:30 am for S&P 500 and Bank of America.

Figure 6: Fixed period backtesting

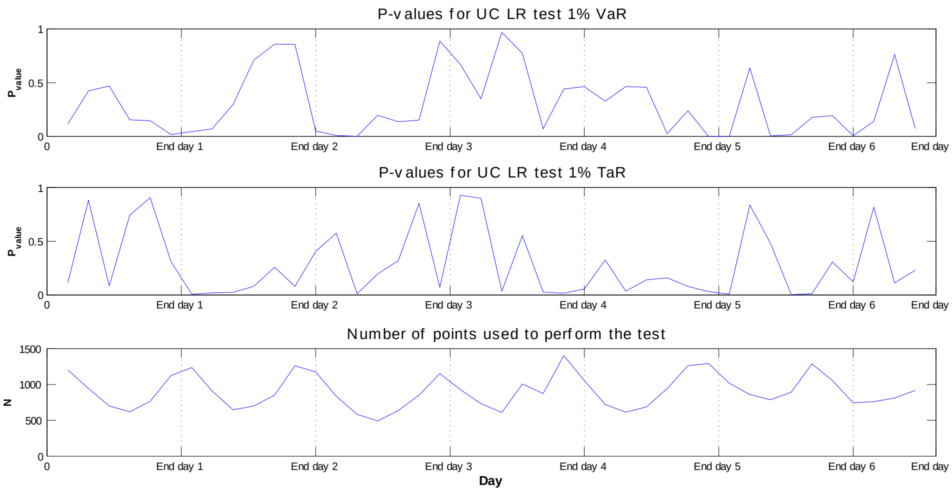

Note: The first two figures report the p-values of the Christoffersen's LR test (1998) for unconditional coverage (UC) for the S&P 500 1%-VaR and 1%-TaR forecasts.. The backtesting procedure is performed with a rolling window of one hour during seven days (September, 23 to October, 1st). The last figure displays the corresponding number of observations used to perform the aforementioned tests.

Figure 8: Rolling window backtesting

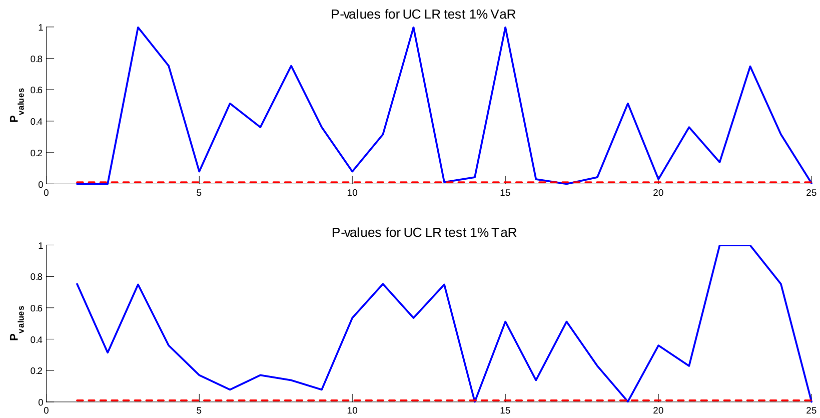

Note: This figure presents the p-values values of the Christoffersen's LR test (1998) of unconditional coverage (UC) for the BAC 1%-VaR and 1%-TaR forecasts. The analysis is performed by using at each iteration a rolling window of 1,000 observations for both the estimation and backtesting samples (with no overlapping for the latter).

Figure 9: Total, Instantaneous and Marginal 1%-VaR

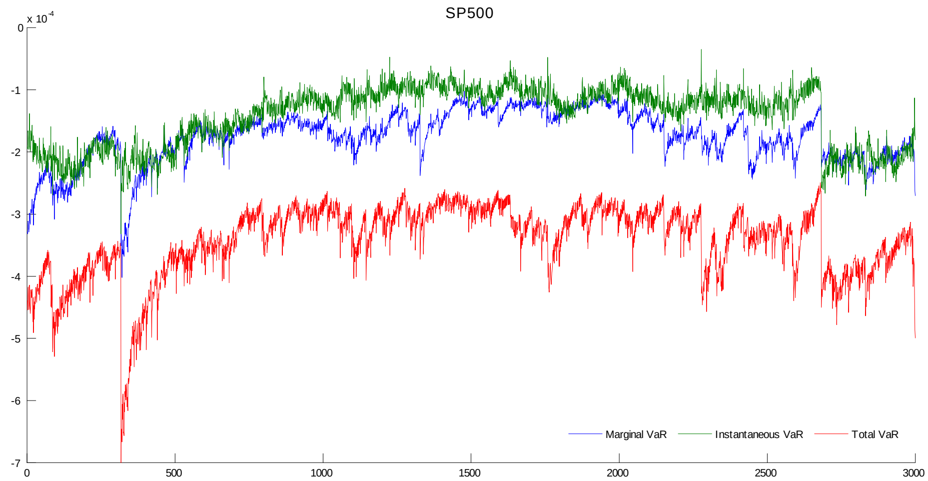

Note: This figure illustrates the total, the marginal and the instantaneous 1%-VaR for SP500. The analysis is performed for price events and a threshold C=0.0075%. The out-of-sample period starts on September 16, 2010 at 9:30 am.

Table 1: Descriptive statistics

Panel A: Durations

	S&P500			BAC			MSFT		
C	**Trade**	**0.0075%**	**0.015%**	**Trade**	**0.035%**	**0.07%**	**Trade**	**0.0175%**	**0.035%**
Mean	4.396	8.870	24.707	12.552	32.283	50.378	2.271	4.535	8.828
Std.Dev	4.706	10.226	30.112	17.638	42.377	67.220	2.067	5.039	15.601
Overdisp	1.070	1.153	1.219	1.405	1.313	1.334	0.910	1.111	1.767
Skewness	3.125	3.488	3.798	3.859	3.645	3.807	3.116	3.607	16.134
Kurtosis	19.504	26.477	30.911	27.324	27.948	28.802	19.642	32.978	592.208
Q(10)	19,038	14,633	13,939	8,232	5,015	4,920	55,613	12,651	3,770
Q(20)	31,748	26,083	24,525	11,375	8,441	8,131	98,958	21,951	4,603
N	223,501	110,774	39,756	78,265	30,423	19,484	432,626	216,641	111,272

Panel B: Returns

Mean	2.2E-07	4.3E-07	1.2E-06	-2.8E-06	-7.2E-06	-1.1E-05	6.3E-08	1.2E-07	2.3E-07
Std.Dev	1.03E-04	1.46E-04	2.36E-04	4.94E-04	7.92E-04	9.75E-04	2.35E-04	3.28E-04	4.22E-04
Skewness	0.031	0.017	0.006	-0.042	-0.010	0.004	-0.004	-0.005	-0.004
Kurtosis	14.944	7.448	3.308	9.815	3.826	2.717	4.627	2.468	1.824
Q(10)	845	696	39	1,461	964	582	72,621	172,101	143,114
Q(20)	850	702	45	1,466	977	593	72,630	172,661	155,022

Note: This table reports some descriptive statistics for the three assets under analysis, namely S&P500, Bank of America (BAC) and Microsoft (MSFT), both for durations (Panel A) and returns (Panel B) series. $Q(.)$ denotes the Ljung-Box Q-statistics computed with 10 or 20 lags. For each asset we consider the duration and return series for trade events (first column) and price events (second and third columns) defined by the threshold C. N denotes the number of observations. The sample period ranges from September 1st to October 29, 2010.

Table 2: Estimates of EACD models

	S&P500		BAC		MSFT	
C	**Trade**	**0.0075%**	**Trade**	**0.035%**	**Trade**	**0.0175%**
ω	0.002	0.008	0.008	0.088	0.000	0.014
	(9.944)	(6.453)	(10.042)	(39.331)	(8.831)	(691.270)
α_1	0.062	0.069	0.154	0.076	0.051	-0.003
	(437.221)	(120.233)	(140.059)	(33.239)	(2,917.927)	(-8.110)
α_2	-0.056	-0.050	-0.131	0.095	-0.049	0.050
	(-1,388.710)	(-158.719)	(-435.223)	(145.334)	(-654.938)	(272.490)
β_1	1.663	1.518	1.484	0.039	1.830	0.132
	(8,640.545)	(1,083.607)	(584.150)	(2.996)	(93,255.556)	(23.277)
β_2	-0.671	-0.544	-0.515	0.711	-0.832	0.808
	(-4,329.126)	(-1,521.715)	(-598.255)	(77.015)	(-8,772.861)	(145.954)
Q(10)	0.819	0.925	0.063	0.994	0.737	0.105
Q(20)	0.983	0.950	0.214	0.584	0.917	0.167

Note: This table contains the parameter estimates of the EACD models for SP&500, Bank of America (BAC) and Microsoft (MSFT). Corresponding t-statistics are in parentheses. The estimation sample covers the first three weeks of our sample (September 1st to September 22, 2010), including 80,845 trade durations for SP500, 27,972 for BAC and 99,155 for MSFT. $Q(.)$ denotes the p-values of the Ljung-Box Q-statistics for serial correlation computed with 10 and 20 lags for the standardized durations. The model is estimated both for trade durations (first column) and price durations (second column) defined for a threshold C.

Table 3: Estimates of GARCH models

	S&P500		BAC		MSFT	
C	**Trade**	**0.0075%**	**Trade**	**0.035%**	**Trade**	**0.0175%**
ω	2.90E-06	2.27E-06	1.17E-04	6.29E-06	2.81E-05	1.02E-05
	(0.001)	(0.000)	(0.020)	(0.001)	(0.009)	(0.002)
α	0.005	0.009	0.034	0.034	0.019	0.022
	(1.484)	(1.451)	(5.636)	(2.752)	(5.833)	(3.358)
β	0.992	0.989	0.951	0.963	0.974	0.974
	(282.171)	(156.343)	(159.130)	(78.980)	(306.724)	(151.694)
Mean ω_i	0.011	0.020	0.182	0.121	0.018	0.023
Range ω_i	[0.005; 0.032]	[0.007; 0.077]	[0.028; 0.686]	[0.020; 0.424]	[0.008; 0.041]	[0.007; 0.077]
Mean α_i	0.014	0.023	0.075	0.061	0.026	0.046
Range α_i	[0.009; 0.022]	[0.016; 0.031]	[0.046; 0.077]	[0.045; 0.063]	[0.019; 0.033]	[0.028; 0.068]
Mean β_i	0.976	0.957	0.753	0.822	0.957	0.933
Range β_i	[0.947; 0.987]	[0.894; 0.977]	[0.344; 0.927]	[0.547; 0.935]	[0.927; 0.973]	[0.860; 0.965]
Q(10)	45	47	65	24	5,836	468
Q(20)	49	54	70	32	5,842	486

Note: This table contains the parameter estimates of the ACD-GARCH models for S&P500, Bank of America (BAC) and Microsoft (MSFT). Corresponding t-statistics are in parentheses. The estimation sample covers the first three weeks of our sample (September 1st to September 22, 2010), including 80,845 trade durations for SP500, 27,972 for BAC and 99,155 for MSFT. We report the mean and the average of the time varying parameters ω_i, α_i and β_i. $Q(.)$ denotes the Ljung-Box Q-statistics for serial correlation computed with 10 and 20 lags for the standardized returns. The model is estimated both for trade durations (first column) and price durations (second column) defined for a threshold C.

Table 4: Backtesting of HFR out-of-sample forecasts

Panel A: VaR backtesting

S&P500

Sample size		1,000			2,000			3,000	
%Hits		0.007			0.009			0.010	
Test	UC	CC	IND	UC	CC	IND	UC	CC	IND
LR test	0.315	0.575	0.753	0.490	0.681	0.589	0.855	0.741	0.452
LR duration test		0.231	0.302		0.560	0.581		0.432	0.214
GMM duration test	0.476	0.823	0.881	0.648	0.937	0.972	0.945	0.754	0.516

BAC

Sample size		1,000			2,000			3,000	
%Hits		0.013			0.009			0.008	
Test	UC	CC	IND	UC	CC	IND	UC	CC	IND
LR test	0.360	0.554	0.558	0.649	0.766	0.567	0.346	0.520	0.517
LR duration test		0.259	0.127		0.178	0.084		0.162	0.125
GMM duration test	0.276	0.522	0.440	0.810	0.561	0.346	0.451	0.485	0.406

MSFT

Sample size		1,000			2,000			3,000	
%Hits		0.011			0.010			0.007	
Test	UC	CC	IND	UC	CC	IND	UC	CC	IND
LR test	0.752	0.016	0.004	0.822	*0.041*	*0.012*	0.081	*0.006*	*0.008*
LR duration test		0.706	0.404		0.242	0.105		*0.006*	*0.012*
GMM duration test	0.574	0.905	0.954	0.963	0.193	0.079	0.094	*0.002*	*0.015*

Panel B: TaR backtesting

S&P500

Sample size		1,000			2,000			3,000	
%Hits		0.006			0.008			0.010	
Test	UC	CC	IND	UC	CC	IND	UC	CC	IND
LR test	0.171	0.377	0.788	0.241	0.448	0.634	0.855	0.556	0.286
LR duration test		0.097	0.205		0.296	0.514		0.637	0.381
GMM duration test	0.207	0.103	0.737	0.331	0.288	0.696	0.944	0.351	0.190

BAC

Sample size		1,000			2,000			3,000	
%Hits		0.014			0.011			0.009	
Test	UC	CC	IND	UC	CC	IND	UC	CC	IND
LR test	0.229	0.398	0.528	0.657	0.709	0.484	0.577	0.670	0.484
LR duration test		0.224	0.141		0.445	0.210		0.634	0.551
GMM duration test	0.206	0.440	0.581	0.545	0.820	0.696	0.663	0.870	0.780

MSFT

Sample size		1,000			2,000			3,000	
%Hits		0.008			0.007			0.006	
Test	UC	CC	IND	UC	CC	IND	UC	CC	IND
LR test	0.512	0.756	0.719	0.155	0.329	0.657	*0.031*	0.085	0.623
LR duration test		0.606	0.970		0.240	0.808		0.057	0.735
GMM duration test	0.673	0.703	0.495	0.180	0.431	0.929	*0.027*	*0.027*	0.885

Note: This table reports the p-values of the backtesting tests of intraday 1%-VaR (Panel A) and 1%-TaR (Panel B) for the three assets under analysis. Italic entries indicate a failure of the model at the 95% confidence level. We test for the unconditional coverage (UC) hypothesis, the independence of violations (IND) and for the conditional coverage (CC). Three tests are considered: Christoffersen's LR test (1998) based on a Markov chain model, Christoffersen and Pelletier's duration-based test (2004) and the GMM-based test proposed by Candelon et al. (2011). %Hits denotes the empirical frequency of violations obtained for a sample size of 1,000, 2,000 and 3,000 observations.

Table 5: Backtesting of HFR out-of-sample forecasts (fixed time periods)

Panel A: VaR backtesting

S&P500

Sample period	30 min (608 obs)			1h (1,122 obs)			4h (3,294 obs)		
%Hits		0.007			0.006			0.010	
Test	UC	CC	IND	UC	CC	IND	UC	CC	IND
LR test	0.3679	0.6493	0.8178	0.1746	0.3807	0.7668	0.9902	0.716	0.4137
LR duration test		0.348	0.654		0.172	0.436		0.220	0.083
GMM duration test	0.591	0.911	0.931	0.264	0.474	0.871	0.808	0.549	0.286

BAC

Sample period	30 min (224 obs)			1h (470 obs)			4h (1,162 obs)		
%Hits		0.018			0.013			0.012	
Test	UC	CC	IND	UC	CC	IND	UC	CC	IND
LR test	0.284	0.523	0.702	0.560	0.781	0.693	0.495	0.668	0.559
LR duration test		0.240	0.106		0.336	0.142		0.193	0.077
GMM duration test	0.109	0.257	0.680	0.387	0.717	0.853	0.315	0.410	0.204

MSFT

Sample period	30 min (906 obs)			1h (1,744 obs)			4h (6,101 obs)		
%Hits		0.012			0.011			0.009	
Test	UC	CC	IND	UC	CC	IND	UC	CC	IND
LR test	0.529	0.017	0.005	0.709	0.050	0.016	0.292	0.000	0.000
LR duration test		0.725	0.459		0.526	0.260		*0.003*	*0.001*
GMM duration test	0.370	0.576	0.551	0.584	0.478	0.268	0.343	*0.004*	*0.008*

Panel B: TaR backtesting

S&P500

Sample period	30 min (608 obs)			1h (1,122 obs)			4h (3,294 obs)		
%Hits		0.007			0.006			0.010	
Test	UC	CC	IND	UC	CC	IND	UC	CC	IND
LR test	0.3679	0.6493	0.8178	0.1746	0.3807	0.7668	0.9902	0.6369	0.3422
LR duration test		0.344	0.635		0.099	0.192		0.744	0.453
GMM duration test	0.548	0.603	0.883	0.250	0.281	0.826	0.815	0.300	0.131

BAC

Sample period	30 min (224 obs)			1h (470 obs)			4h (1,162 obs)		
%Hits		0.013			0.011			0.013	
Test	UC	CC	IND	UC	CC	IND	UC	CC	IND
LR test	0.623	0.850	0.775	0.887	0.938	0.743	0.339	0.520	0.531
LR duration test		0.977	0.881		0.900	0.748		0.350	0.200
GMM duration test	0.425	0.368	0.190	0.510	0.571	0.244	0.271	0.518	0.473

MSFT

Sample period	30 min (906 obs)			1h (1,744 obs)			4h (6,101 obs)		
%Hits		0.009			0.008			0.008	
Test	UC	CC	IND	UC	CC	IND	UC	CC	IND
LR test	0.720	0.873	0.706	0.393	0.619	0.634	0.110	0.187	0.373
LR duration test		0.740	0.754		0.496	0.684		0.220	0.841
GMM duration test	0.881	0.646	0.381	0.542	0.590	0.544	0.119	0.165	0.779

Note: This table reports the p-values for the backtesting tests of intraday 1%-VaR (Panel A) and 1%TaR (Panel B) for the three assets under analysis. Italic entries indicate a failure of the model at the 95% confidence level. The backtesting procedure is applied over a fixed period of time (30 minutes, 1 hour, and 4 hours, respectively). The corresponding number of observations is reported for each asset. We test for the unconditional coverage (UC) hypothesis, the independence of violations (IND) and for the conditional coverage (CC) hypothesis. Three tests are considered: Christoffersen's LR test (1998) based on a Markov chain model, Christoffersen and Pelletier's duration-based test (2004) and the GMM-based test proposed by Candelon et al. (2011). %Hits denotes the empirical frequency of violations.