

2000s-19

**Rolling-Sample Volatility
Estimators: Some New Theoretical,
Simulation and Empirical Results**

Elena Andreou, Eric Ghysels

Série Scientifique
Scientific Series



CIRANO
Centre interuniversitaire de recherche
en analyse des organisations

Montréal
Mai 2000

CIRANO

Le CIRANO est un organisme sans but lucratif constitué en vertu de la Loi des compagnies du Québec. Le financement de son infrastructure et de ses activités de recherche provient des cotisations de ses organisations-membres, d'une subvention d'infrastructure du ministère de la Recherche, de la Science et de la Technologie, de même que des subventions et mandats obtenus par ses équipes de recherche.

CIRANO is a private non-profit organization incorporated under the Québec Companies Act. Its infrastructure and research activities are funded through fees paid by member organizations, an infrastructure grant from the Ministère de la Recherche, de la Science et de la Technologie, and grants and research mandates obtained by its research teams.

Les organisations-partenaires / The Partner Organizations

- École des Hautes Études Commerciales
- École Polytechnique
- Université Concordia
- Université de Montréal
- Université du Québec à Montréal
- Université Laval
- Université McGill
- MEQ
- MRST
- Alcan Aluminium Ltée
- Banque Nationale du Canada
- Banque Royale du Canada
- Bell Québec
- Développement des ressources humaines Canada (DRHC)
- Fédération des caisses populaires Desjardins de Montréal et de l'Ouest-du-Québec
- Hydro-Québec
- Imasco
- Industrie Canada
- Raymond Chabot Grant Thornton
- Télélobe Canada
- Ville de Montréal

© 2000 Elena Andreou et Eric Ghysels. Tous droits réservés. All rights reserved.

Reproduction partielle permise avec citation du document source, incluant la notice ©.

Short sections may be quoted without explicit permission, provided that full credit, including © notice, is given to the source.

Ce document est publié dans l'intention de rendre accessibles les résultats préliminaires de la recherche effectuée au CIRANO, afin de susciter des échanges et des suggestions. Les idées et les opinions émises sont sous l'unique responsabilité des auteurs, et ne représentent pas nécessairement les positions du CIRANO ou de ses partenaires.

This paper presents preliminary research carried out at CIRANO and aims at encouraging discussion and comment. The observations and viewpoints expressed are the sole responsibility of the authors. They do not necessarily represent positions of CIRANO or its partners.

Rolling-Sample Volatility Estimators: Some New Theoretical, Simulation and Empirical Results*

Elena Andreou[†], Eris Ghysels[‡]

Résumé / Abstract

Nous proposons des extensions de la théorie asymptotique de Foster et Nelson pour l'estimation de variance. Nous proposons une approximation asymptotique qui permet de comparer des estimateurs obtenus à partir de données avec fréquences d'échantillonnage différentes. Une autre extension consiste à appliquer les arguments de Foster et Nelson à des processus plus généraux tels que la volatilité intégrée.

We propose different extensions of the continuous record asymptotic analysis for rolling sample variance estimators developed by Foster and Nelson (1996). First, despite the difference in information sets we are able to compare the asymptotic distribution of volatility estimators involving data sampled at different frequencies. We focus on traditional historical volatility filters involving monthly, daily and intra-daily observations. Second, we introduce a continuous record asymptotics approach for estimating the so called integrated volatility, which represents the cumulative integral of instantaneous volatility. The new approach treats integrated volatility as a stochastic process sampled at high frequencies and suggests rolling sample estimators which share many features with spot volatility estimators. We discuss optimal weighting schemes for integrated volatility estimators. Thirdly, we establish the links between various spot and integrated volatility estimators. Theoretical results are complemented with extensive Monte Carlo simulations and an empirical investigation.

Mots Clés : Données haute fréquence, volatilité, Monte Carlo

Keywords: High-frequency data, volatility, continuous record asymptotics, Monte Carlo simulations

JEL: G10, C15, C13

Corresponding Author: Eric Ghysels, Pennsylvania State University, 523 Kern Graduate Building, University Park, PA 16802-3305 Tel.: (814) 865 7924 Fax: (814) 863 4775 email: eghysels@grizzly.la.psu.edu
The authors would like to thank the hospitality and financial support of CIRANO and the University of Manchester. We would like to thank Tim Bollerslev and Denise Osborn for comments on an earlier draft of this paper. We would also like to thank Michel Dacorogna of Olsen and Associates for providing us the data.

[†] University of Manchester

[‡] Pennsylvania State University

Introduction

There are many strategies for estimating time-varying conditional variances and covariances of financial market data. Some involve parametric models, the most popular belong either to the ARCH or the SV class of models, while others are purely data-driven. We focus exclusively on the latter type of estimation strategies which typically involve a rolling or block sample approach. Data-driven volatility estimates can be distinguished by (1) the sampling frequency, (2) the data window length and (3) the weighting scheme. There are indeed various schemes, some consist of slicing the returns data into non-overlapping blocks (of equal size) while others rely on a moving window. The block-sampling approach estimates a time series process of conditional variances or covariances using exclusively data from a single block. This approach has its roots in various academic papers, such as Merton (1980), Poterba and Summers (1986), French, Schwert and Stambaugh (1987), Schwert (1989 1990a,b) and Schwert and Seguin (1990). Another commonly used approach relies on sliding spans of data generating rolling regressions. Early contributions include Fama and MacBeth (1973), Officer (1973) and Merton (1980). Asset pricing applications often involve sampling at a monthly frequency with rolling windows between 5 to 10 years. Recent examples include Campbell and Lettau (1999), Chan, Karceski and Lakonishok (1999), Fleming, Kirby and Ostdiek (2000) and Malkiel and Xu (1999). Practitioners compute daily volatilities using the same schemes applied to a daily and monthly sampling frequency, a prominent example are the *RiskMetrics* volatility measures.

There are no clear rules regarding the choice of data sampling frequency nor the number of lags to include. Some authors use monthly data and take a 60-month filter, following Fama and MacBeth (1973). Others have used daily data and take monthly sums and cross-products of (squared) returns, following French, Schwert and Stambaugh (1987). Some authors report estimators involving different data frequencies to check for robustness or emphasize different features. For instance, Schwert (1990b) considers volatility estimates based on yearly averages of monthly squared returns, daily returns squared and averaged across a month and finally 15-minute squared returns across a trading day. Likewise, Campbell *et al.* (2000) consider a 12-month span rolling sample estimator of volatility, daily averages across a month as well as a quarter. How do these various volatility estimation schemes compare and could we verify robustness via different sampling frequencies? The main purpose of this paper is to provide answers to some of these questions. The theoretical work of Foster and Nelson (1996) formally establishes that with sufficiently high frequency data one can estimate instantaneous volatility, denoted by σ_t and sometimes also called spot volatility, using rolling and block sampling filters. They provide some powerful results about the estimation of spot volatility and establish the efficiency of different weighting schemes. They also show that for a large class of continuous path stochastic volatility models the optimal weighting scheme is exponentially declining and provide formulas for the optimal lag (lead) length of various estimation procedures. The advent of intra-day high frequency financial data has made it possible to carry this approach a step further. In a recent set of papers Andersen and Bollerslev (1998), Barndorff-Nielsen and Shephard (1999) and Andersen *et al.* (1999) suggest the sum of squared high frequency (say 5 minutes) returns as an approximation to the integrated volatility, $\int_{t-H}^t \sigma_\tau^2 d\tau$, i.e. an integral of instantaneous volatilities which we will call Q_t^H since it relates to the quadratic variation

of a process over a time interval from t to $t + H$. Various related earlier measures were also proposed such as the cumulative absolute returns studied by Hsieh (1991) and the (monthly) sum of squared returns augmented by cross-products to account for correlation as suggested by Merton (1980), Poterba and Summers (1986) and French, Schwert and Stambaugh (1987).

We propose different extensions of the continuous record asymptotic analysis of Foster and Nelson for data-driven variance estimators. First, despite the difference in information sets we are able to compare the asymptotic distribution of estimators involving data sampled at different frequencies. A priori it may seem impossible to compare rolling sample estimators involving data sampled at different frequencies. However, Foster and Nelson impose certain regularity conditions, which enable us to maintain the same asymptotic distributional properties of the mean squared errors of volatility filters involving data sampled at different frequencies. This insight is the key to the first extension of the Foster and Nelson results presented in this paper. Next we transfer the Foster and Nelson continuous record asymptotic analysis to the efficient estimation of the integrated volatility process such as Q_t^H . We view Q_t^H as a continuous time process in t for fixed H which is sampled discretely at high frequencies. The advantage of this approach, also suggested by Barndorff-Nielsen and Shephard (1999), is that the optimal weight design for integrated volatility estimators can be obtained directly from Foster and Nelson's analysis with the appropriate modifications. This leads to the introduction of a large new class of integrated volatility estimators. The same arguments are also applied to cumulative absolute return measures which also yields a new class of such estimators.

It obviously takes Monte Carlo simulations to back up the accuracy of any asymptotic distribution theory in sample sizes encountered in practice. For our comparison of sampling frequencies this wisdom certainly applies even more so than is usually the case. The continuous record asymptotic arguments provide only an approximation to what we encounter in practice. Moreover, the core idea of the asymptotics is ever finer sampling while we use the distribution theory to compare asymptotic efficiency of schemes involving different sampling frequencies. To validate our asymptotic reasoning we report an extensive Monte Carlo study. The design is built on stochastic volatility models which describe FX and equity markets data. We follow Andersen and Bollerslev (1998) and take advantage of the elegant GARCH temporal aggregation results of Drost and Nijman (1993) and Drost and Werker (1996).

The paper is organized as follows: In the first section we briefly review the relevant theoretical results from Foster and Nelson which support the comparison of rolling sample estimates with various sampling frequencies. The second section introduces the continuous record asymptotics for integrated volatility estimation. The third section provides the details about the Monte Carlo design. The simulation results are described in the fourth section. The fifth section reports empirical comparisons of intra-day sampling schemes using 5-, 15- and 30-minute FX series to examine the efficiency of daily volatility filters and daily S&P 500 data to study the efficiency of alternative monthly data-driven filters. A final section concludes the paper.

1 Comparing Estimators for Spot Volatility with Different Sampling Frequencies

The focus of this first section is on data-driven estimators for spot volatility which figure prominently in many asset pricing applications. Spot volatility estimators are also a natural starting point since the work of Foster and Nelson refers to this type of estimators. For convenience we assume a continuous time process and model the instantaneous returns $r_t \equiv dp_t$, where p_t is the log price, as a stochastic volatility process.¹ In particular, let:

$$\begin{aligned} dp_t &= \mu(p_t, \sigma_t)dt + \sigma_t dW_{1t} \\ d\sigma_t^2 &= \zeta(p_t, \sigma_t)dt + \delta(p_t, \sigma_t)dW_{2t} \end{aligned} \tag{1.1}$$

where W_{1t} and W_{2t} are standard Brownian motions (possibly correlated) and the functions $\mu(\cdot, \cdot)$, $\zeta(\cdot, \cdot)$ are continuous and $\delta(\cdot, \cdot)$ strictly positive. In a first subsection we propose extensions of the continuous record asymptotic analysis for rolling sample variance estimators of Foster and Nelson (1996) which enable us to compare the asymptotic distribution of estimators for σ_t involving data sampled at different frequencies. In a second subsection we discuss some comparisons of filters which maintain asymptotic efficiency across various sampling frequencies.

1.1 Theoretical Extensions

The powerful results in Foster and Nelson are driven by a continuous record asymptotic theory which assumes that a fixed span of data is sampled at ever finer intervals. The basic intuition driving the results is that normalized returns, r_t/σ_t , over short intervals appear like approximately i.i.d. with zero conditional mean and finite conditional variance and have regular tail behavior which make the application of Central Limit Theorems possible. Foster and Nelson impose several regularity conditions for the diffusion process appearing in (1.1) which we will not review here. Instead, we will only highlight certain assumptions which are critical for our analysis. We will adopt a notation slightly different from Foster and Nelson, but similar to that used by Drost and Nijman (1993), Drost and Werker (1996) and Andersen and Bollerslev (1998). Namely, let $r_{(m),t} \equiv p_t - p_{t-1/m}$ be the discretely observed time series of continuously compounded returns with m observations per day, per month, or whichever benchmark applies. Henceforth, we will call $m = 1$ the benchmark frequency, which in the context of our paper will be either daily or monthly. Hence, the unit interval $r_{(1),t}$ is assumed to yield the daily or monthly return.² The $r_{(m),t}$ process is a discrete step function with the $1/m$ horizon returns

¹As noted by Foster and Nelson, their analysis applies not only to SV diffusions but also, with appropriate modifications, to discrete time SV, to ARCH models and to certain types of random coefficient models. While we start with a continuous time SV framework, we will focus later on a particular case which yields a GARCH(1,1) model using exact discretization methods (see also Drost and Werker (1996), Meddahi and Renault (1997) and Andersen and Bollerslev (1998)).

²The notion of a benchmark frequency will be used extensively, particularly when simulating models. We use the daily and monthly examples as they are most commonly encountered in applications. The benchmark frequency will also serve as a reference frequency to normalize simulation results and make them comparable across frequencies as will be discussed later.

determining the step size. Therefore, when calculating monthly volatility estimates, one can either rely on daily data, i.e. use $r_{(22),t}$ with approximately 22 trading days per month, such as Merton (1980), French, Schwert and Stambaugh (1987), Schwert (1990a), among others, or else use sliding spans of squared monthly returns $r_{(1),t}$ such as in Officer (1973), Merton (1980), Schwert (1989), among others. Similarly, to obtain daily volatility estimates, one can rely on high frequency financial data, e.g. five minute data (with $m = 288$ for FX data, as for instance in Andersen and Bollerslev (1998) and Andersen *et al.* (1999) or $m = 78$ for equity markets as in Chin, Chan and Karolyi (1991)), or use daily data as for instance in the *RiskMetrics* schemes. In general, with ever finer sampling intervals, i.e. $m \rightarrow \infty$, we approach the continuously compounded returns, or equivalently $r_{(\infty),t} \equiv r_t$. The process $\{r_{(m),t}\}$ is adapted to the filtration $\{\mathcal{F}_{(m),t}\}$ and conditional expectations and variances will be denoted as $E_{(m),t}(\cdot)$ and $Var_{(m),t}(\cdot)$ respectively, whereas unconditional moments follow a similar notation, $E_{(m)}(\cdot)$ and $Var_{(m)}(\cdot)$. From (1.1) we obtain the discrete time dynamics:

$$r_{(m),t} = \mu_{(m),t}m^{-1} + M_{(m),t} - M_{(m),t-1/m} \equiv \mu_{(m),t}m^{-1} + \Delta_{(m)}M_{(m),t},$$

which is the so called Doob-Meyer decomposition of the $1/m$ horizon returns into a predictable component $\mu_{(m),t}$ and a local martingale difference sequence. Consequently:

$$Var_{(m),t}(r_{(m),t}) \equiv E[(\Delta_{(m)}M_{(m),t} - \mu_{(m),t})^2 | \mathcal{F}_{(m),t}] = \sigma_{(m),t}^2 m^{-1}$$

where $\sigma_{(m),t}^2$ measures the conditional variance per unit of time. Various data-driven estimators for $\sigma_{(m),t}^2$ can generically be written as:

$$\hat{\sigma}_{(m),t}^2 = \sum_{\tau} w_{(\tau-t)} (r_{(m),t} - \hat{\mu}_{(m),t})^2 \quad (1.2)$$

where $w_{(\tau-t)}$ is a weighting scheme and $\hat{\mu}_{(m),t}$ is a (rolling sample) estimate of the drift. To facilitate the discussion, we restrict our analysis to flat weighting schemes involving $n_L m^{-1/2}$ lags and $n_R m^{-1/2}$ leads. Other weighting schemes, including optimal ones, will be covered at the end of this section. When $n_R = 0$, the filter is one-sided and backward-looking, a case of most practical interest. Please note also that $m = 1$ implies that n_L is simply the number of days, or months, of squared returns used to compute the conditional volatility (again assuming $n_R = 0$).³ These schemes include the most commonly used volatility estimators involving equally weighting observations throughout the day, across all days of a month, or a sliding span of daily or monthly returns. The asymptotic efficiency of $\hat{\sigma}_{(m),t}^2$ only depends on the process characteristics once the filter weights are fixed, in this particular case once n_L and n_R are determined.⁴ Theorem 2 of Foster and Nelson establishes that $m^{1/4}(\hat{\sigma}_{(m),t}^2 - \sigma_t^2) \rightarrow N(0, C_{(m),t})$

³This includes the zero lag, since we assume end-of-day, or end-of-month, volatilities. For the 24-hour FX market we pick a particular time, like 21:00 hours GMT every day, similar to Andersen and Bollerslev(1998).

⁴Foster and Nelson assume that the $Sup_{\tau}\{\tau : w_{\tau-t} > 0\} - Inf_{\tau}\{\tau : w_{\tau-t} > 0\} = O(m^{-1/2})$ and hence shrinks as m increases and is bounded in probability by $m^{-1/2}$ (see Foster and Nelson (1996, Assumption D)). For flat weights involving $n_L m^{-1/2}$ lags and $n_R m^{-1/2}$ leads the weighting scheme can be characterized as $w_{(m),t} = m^{-1/2}(n_L + n_R)^{-1} I\{\tau \in [-n_L m^{-1/2}, n_R m^{-1/2}]\}$.

as $m \rightarrow \infty$, where the continuous record asymptotic variance for a flat weighting scheme equals:

$$C_{(m),t}^F \equiv \frac{\theta_{(m),t}}{n_R + n_L} + \sqrt{\theta_{(m),t}\Lambda_{(m),t}} \rho_{(m),t} \frac{n_R - n_L}{n_R + n_L} + \Lambda_{(m),t} \frac{n_R^3 + n_L^3}{3(n_R + n_L)^2}. \quad (1.3)$$

The superscript F in (1.3) refers to the flat weighting scheme. Besides the window length parameters n_L and n_R , three other elements determine the asymptotic efficiency of the filter $\hat{\sigma}_{(m),t}^2$. They are $\Lambda_{(m),t}$, $\theta_{(m),t}$ and $\rho_{(m),t}$, each depending on the sampling frequency m and the characteristics of the underlying process (1.1). The process $\Lambda_{(m),t}$ represents the “variance of the variance”, and therefore any increase of its value will increase $C_{(m),t}^F$ and deteriorate the asymptotic efficiency of filtering. The process $\theta_{(m),t}$ represents the conditional fourth moment. When the data span increases, namely when $n_R + n_L$ increases, then the first term on the right hand side of equation (1.3) decreases, as the usual asymptotics would predict. Note, however, that the third term in the same equation increases with wider data spans, a result driven by the fact that only local cuts of the data exhibit a relatively stable variance. Finally, the process $\rho_{(m),t}$ measures the correlation between the empirical second moment and the conditional variance. As Foster and Nelson observe, the correlation is unity for ARCH-type processes and zero for continuous path diffusions. To streamline the discussion we do not provide explicit characterizations of the three processes since details appear in Foster and Nelson.

We are interested in comparing $C_{(1),t}^F$, which is based on the benchmark sampling frequency, with $C_{(m),t}^F$ for $m > 1$, or equivalently compare the asymptotic efficiency of volatility estimators involving data sampled at different frequencies. The comparison can be demonstrated by the following illustrative example: Suppose one starts with a 30-day historical volatility estimate and instead of sampling only once a day one slices the trading period in half and therefore has twice as many returns. Now we ask the following hypothetical question: How many lags of half-day returns does it take to attain the *same* efficiency as the historical 30-day filter? Please note that we do not change the weighting scheme. We only sample twice more often and try to find out how many lags of half-daily returns attain the same efficiency as a 30-day filter using daily returns. The answer is not 15 days worth of lagged returns sampled twice daily, i.e. the same number of observations which was thirty. Indeed, to maintain the same (asymptotic) efficiency we obviously do not have a simple linear trade-off between sampling frequency and number of observations. It takes in fact more than 15 days (or 30 half-day observations) of data to maintain the same efficiency. Hence, we need more and more lags as the sampling becomes finer to maintain a particular level of efficiency. It is important to note here that we only try to maintain a certain level of efficiency, and therefore sidestep the issue whether the efficiency one attains is adequate. Bai, Russell and Tiao (1999) have recently argued that Merton’s (1980) result may be an utopia since the precision of volatility estimates may vary substantially depending on the features of the return process. As the result in (1.3) clearly shows, one may have a very large or small asymptotic MSE depending on the magnitude of the conditional higher moment terms. Bai *et al.* (1999) examine the accuracy of several data-driven and GARCH filters as one changes the sampling frequency and the process characteristics. Instead, we investigate how one can maintain a certain level of efficiency as one changes the sampling frequency. While the theoretical analysis in this section sidesteps the

issue of efficiency levels, it should be noted that the Monte Carlo simulations reported later will examine the actual efficiency levels attained for data generating processes which resemble actual observed series. We will also join the discussion of Bai *et al.* (1999) in examining the effects of excess kurtosis in our simulation and empirical analysis.

The asymptotic efficiency of $\hat{\sigma}_{(m),t}^2$ only depends on the process characteristics once the filter weights are fixed, in this particular case once n_L and n_R are determined. Therefore, to compare $C_{(1),t}^F$ with $C_{(m),t}^F$ we must be able to appraise the difference between $\theta_{(m),t}$ and $\theta_{(1),t}$, and also compare $\Lambda_{(1),t}$, $\rho_{(1),t}$ with $\Lambda_{(m),t}$ and $\rho_{(m),t}$. Fortunately, this comparison is possible as a result of a critical assumption in Foster and Nelson, namely their Assumption B, which states that:

$$\begin{aligned} \text{Sup}_{t \leq \tau \leq t+1/\sqrt{m}} \left| \theta_{(m),\tau} - \theta_{(m),t} \right| &= o_p(1) \\ \text{Sup}_{t \leq \tau \leq t+1/\sqrt{m}} \left| \Lambda_{(m),\tau} - \Lambda_{(m),t} \right| &= o_p(1) \\ \text{Sup}_{t \leq \tau \leq t+1/\sqrt{m}} \left| \rho_{(m),\tau} - \rho_{(m),t} \right| &= o_p(1). \end{aligned}$$

This assumption implies that the conditional fourth moments, variance of variance and correlation roughly stay constant over small intervals, where small is interpreted as an interval of size $1/\sqrt{m}$. This assumption guarantees that the process (1.1) is regular enough with higher moments changing slowly over time. We will replace Assumption B of Foster and Nelson with a slightly different condition, namely:

$$\begin{aligned} \text{Sup}_{t \leq \tau \leq t+1/\sqrt{m}} \left| \theta_{(m),\tau} - \theta_{(1),t} \right| &= o_p(1) \\ \text{Sup}_{t \leq \tau \leq t+1/\sqrt{m}} \left| \Lambda_{(m),\tau} - \Lambda_{(1),t} \right| &= o_p(1) \\ \text{Sup}_{t \leq \tau \leq t+1/\sqrt{m}} \left| \rho_{(m),\tau} - \rho_{(1),t} \right| &= o_p(1) \end{aligned}$$

This assumption is in the same vein as that used by Foster and Nelson, the main difference is that in Assumption B of Foster and Nelson the filtration is kept constant at $\{\mathcal{F}_{(m),t}\}$. We require the stronger condition that relative to the benchmark filtration $\{\mathcal{F}_{(1),t}\}$ we have local stability of the conditional higher moments at all finer sampling frequencies. In particular, we can write $\text{Sup}_{t \leq \tau \leq t+1/\sqrt{m}} \left| \theta_{(m),\tau} - \theta_{(1),t} \right| \leq \left| \theta_{(m),t} - \theta_{(1),t} \right| + \text{Sup}_{t \leq \tau \leq t+1/\sqrt{m}} \left| \theta_{(m),\tau} - \theta_{(m),t} \right|$. Hence, Assumption B of Foster and Nelson implies our condition provided $\left| \theta_{(m),t} - \theta_{(1),t} \right|$ is $o_p(1)$. One can interpret this condition as saying that at the daily (monthly) level (and beyond) we have reached stability of all relevant conditional higher moments. To a certain degree this is what underlies the empirical application in Foster and Nelson who consider optimal filtering of daily volatility for the S&P 500 index.

Though the regularity condition is fairly mild it is not innocuous of course, and warrants some further discussion. First and foremost, it should be noted that the assumption pertains to *conditional* higher moments after the volatility dynamics are taken into account. Unconditional kurtosis can vary dramatically as one changes the sampling frequency (see Bai *et al.* (1999) for recent empirical evidence). Several observations can be made about the assumption. First, for $\rho_{(m),t}$ this condition is often trivially satisfied when ARCH or SV type processes are considered, since $\rho_{(m),t}$ is constant across m and t . Second, Foster and Nelson propose

estimators for $\theta_{(m),t}$ and $\Lambda_{(m),t}$ (see formula (10) and (11) in their paper). Hence, one can inspect the time series of point estimates of conditional higher moments, say $\theta_{(1),t}$ and $\theta_{(m),t}$ for $m \geq 1$ where t is for instance at the daily frequency. Third, there are several papers which document the behavior of conditional higher moments, particularly the kurtosis, as sampling frequencies increase. For instance, Baillie and Bollerslev (1989) find that GARCH parameter estimates and tail characteristics for FX daily data carry over to weekly, fortnightly, whereas the degree of leptokurtosis and time-dependent heteroscedasticity is reduced as the length of the sampling interval increases to monthly data. Engle, Ito and Lin (1990) also examine four hourly FX data series and find that although all have excess kurtosis, these do not, however, deviate dramatically from the daily kurtosis levels encountered in the literature.⁵ Hsieh (1991) presents results pertaining to equity markets, namely he examines 15-minute and daily S&P 500 returns and finds comparable results. Chin, Chan and Karolyi (1991) also examine S&P 500 returns at 5-minute sampling frequency and report sample kurtosis which do not substantially deviate from the daily sample kurtosis levels. This evidence suggests that at least for daily benchmark frequencies our assumption appears to be satisfied. For data sampled at the monthly frequency it is well known that the conditional kurtosis increases when one moves from monthly to weekly and daily sampling frequencies. Hence, one may be more comfortable with this assumption when the benchmark frequency is daily and one samples at finer intra-daily frequencies. Therefore, when the monthly frequency is the benchmark frequency and we compare filters involving, say daily data, with monthly data filters our comparison may not be very accurate, even on asymptotic grounds, since we ignore the variation of higher moments. It should also be noted that for some models, such as the GARCH diffusion which we will consider in the Monte Carlo section, we can compute for each m the entries to (1.3).

The great advantage of these results is that we can simply drop all the subscripts in (1.3) and write the expression as:

$$C_{(m),t}^F \approx \frac{\theta}{n_R + n_L} + \sqrt{\theta\Lambda\rho} \frac{n_R - n_L}{n_R + n_L} + \Lambda \frac{n_R^3 + n_L^3}{3(n_R + n_L)^2} \quad (1.4)$$

for $m \geq 1$. It is worth noting that to facilitate their discussion Foster and Nelson also simply drop all the subscripts to (1.3), see in particular their equation (9). This representation allows us to make relatively simple comparisons of asymptotically equivalent sampling schemes involving sampling at different frequencies m . For instance, when n_R , n_L , θ , Λ and ρ are fixed, then a one-sided window of length n_L with daily data yields the same asymptotic efficiency as a one-sided window of length $n_L m^{-1/2}$ with intra-daily sampled data at frequency $1/m$.

1.2 Some Examples of Equivalent Filters

We present asymptotically MSE-equivalent one-sided volatility filters for different sampling frequencies. First we consider the case where the benchmark frequency is daily data. Panels A through C of Table I report numerical comparisons for both the 24-hour foreign exchange and

⁵More recently, Goodhart and Figliuoli (1991), Bollerslev and Domowitz (1993) and Andersen and Bollerslev (1998) also report FX intra-daily results which confirm the stability of the conditional kurtosis at high sampling frequencies.

6.5 hour trading equity markets. The top panel pertains to 22-day historical volatility filters, Panel B covers the 26-day filter and the third panel pertains to 30-day filters. These three cases of n_L correspond roughly to the range of lags often encountered in practice. We examine filters which yield the same asymptotic $C_{(m),t}^F$ for intra-day sample frequencies such as $m = 2, 24, 13, 288, 78, 1440$ and 390 . These sampling frequencies correspond respectively to half-daily sampling, hourly FX, half-hourly equity, five-minute FX, five-minute equity and finally, one-minute FX and equity markets. We report the number of lags $n_L m^{-1/2}$ for $n_L = 22, 26$ and 30 days and all values of m and translate these lags in number of days of FX and equity market high frequency observations. We note from Table I that for a 22-day filter we need 374 lags of five-minute observations in FX markets and 195 lags in equity markets to maintain the same asymptotic MSE. This amounts to 1.3 days FX data and 2.5 days stock data. When four daily observations are added, i.e. n_L equals 26, the trade-off gets obviously worse, with respectively 442 (FX) and 230 (equity) 5-minute lags (or 1.5 and 2.9 days, respectively). Adding four more daily returns yields lags equal to 510 and 265, i.e. roughly 2 days of five-minute FX data or 3.5 days of equity returns sampled at the same high intra-day frequency. Sampling at the less realistic one-minute level gets the lag length data for FX under one day for all values of n_L but fails to reach below one day of equity markets (the shortest historical volatility filter $n_L = 22$ yields 1.1 days of one-minute data).

The remaining two panels in Table I cover two monthly benchmark frequency cases. These are the 60-month and 12-month filters, again two cases commonly found in the literature. From Panel D we note that a 60-month filter and a 282-lag filter of daily returns are asymptotically equivalent, i.e. it takes 13 months of daily data to attain the same efficiency as a filter with 60 monthly observations. Half-daily sampling yields a four month gain, namely 398 lags or 9 months of data are necessary to maintain the same efficiency. Taking hourly data reduces this further to slightly less than a quarter of observations (i.e. 2.6 months). The final panel shows a similar trade-off. The commonly used annual lag length of monthly returns is equivalent to about the same quarter length of daily data. This is in fact an interesting comparison. For instance, Campbell *et al.* (2000) use the annual filter with monthly data to extract volatility and decompose it into a market, industry- and firm-specific component and also use a quarterly block sampling scheme to compare the three volatility components with GNP growth rates. The results in Panel E of Table I show that these two schemes are roughly asymptotically equivalent (at the monthly and hence also quarterly frequency).

The arguments presented so far can be reversed as well. In Table II we report asymptotically equivalent historical window one-day and one-month lengths for spot volatility filters. The entries to Table II report numerical calculations based on equation (1.4) using $n_L = \sqrt{m}$ as the number of daily or monthly observations in a one-sided historical volatility filter which is asymptotically equivalent to a one-day or one-month filter with sampling frequency $1/m$. The top panel covers the daily benchmark frequency whereas the second panel covers the monthly benchmark frequency. From Table II we learn that an hourly FX filter (i.e. $m = 24$) is equivalent to a 5-day historical volatility filter, while half-hourly equity market data filter with $m = 13$ is worth one day less. The more interesting case of a five minute FX market filter (i.e. $m = 288$) is asymptotically equivalent to a 17-day historical volatility filter. The equity market filter with the same frequency, is as efficient as a 9-day filter. We know from the figures

reported in Table I that we only gain the efficiency of the usual historical volatility once we sample at one minute intervals. Indeed, as the results in Table II indicate a one minute FX filter with $m = 1440$ is as efficient as a 38-day historical volatility estimate and a one-minute stock market filter is equal to a 20-day one.

The lower panel of Table II deals with one month worth of data and reports comparable monthly frequency historical filters. For example, the first entry in the lower panel shows that a 22-day historical volatility filter is equivalent to a 5-month filter. One month of hourly FX data corresponds to a filter of monthly data almost two years long. The most extreme case reported is that of five-minute FX data when sampled for an entire month correspond to an 80-month historical volatility filter.

Obviously, at this point we do not know whether the theoretical asymptotic trade-offs described in Tables I and II are a good approximation of what we encounter in practice, hence the need to conduct a thorough Monte Carlo investigation to which we return in Section 3. Before turning our attention to simulations it is worth noting that the arguments presented in this section easily extend to weighting schemes other than the flat scheme discussed so far. For instance, Theorem 6 of Foster and Nelson covers the so called dominating flat weights, which have the same sliding span of data as the flat scheme but where the actual weights are reshaped (see formula (17) in Foster and Nelson). The resulting $C_{(m),t}^D$, where D stands for the dominating flat scheme, is again a function of $\rho_{(m),t}$, $\theta_{(m),t}$ and $\Lambda_{(m),t}$. Hence, under the same regularity conditions we can compare dominating flat weighting schemes for different m on the basis of n_L and n_R . The optimal exponentially declining weighting scheme considered in Theorem 5 of Foster and Nelson is slightly more complicated as it involves, at least in principle infinite weighting schemes. It is noted though that in practice such weights need to be truncated (otherwise they would also violate Assumption D of Foster and Nelson). Our analysis would also apply to truncated optimal exponentially declining weights.

2 Continuous record asymptotic analysis of integrated volatilities

The analysis of Foster and Nelson applies not only to SV diffusions but also to a variety of other processes, including the integrated volatility processes over some horizon H . For instance, Taylor (1986), Hsieh (1991) and Schwert (1989) estimate conditional volatility as the distributed lag of the absolute value squared) residuals from an autoregressive model for returns.⁶ Hence, one examines $\sum_{t-H}^t |r_t|$ (ignoring the conditional mean correction) as a measure of integrated volatility, which sometimes also called the cumulative absolute returns. Andersen and Bollerslev (1998) considered $\int_{t-H}^t \sigma_\tau^2 d\tau$, i.e. an integral of instantaneous volatilities which we will call Q_t^H since it relates to the quadratic variation of a process. They suggest the use of 5-minute squared returns FX data to approximate the integral. French, Schwert and Stambaugh (1987), Schwert (1989, 1990a) used a similar measure involving the (monthly) sum of squared daily

⁶Absolute returns are taken instead of squared returns since it is argued that in the presence of deviations from normality absolute values could be more robust than squared values for conditional variance estimation (see e.g. Davidian and Carroll (1987)). This will be discussed later in greater detail.

returns augmented by cross-products to account for correlation. We will first focus on the case of Q_t^H , and subsequently broaden our analysis to alternative measures like the cumulative absolute returns.

The analysis in this section considers Q_t^H , or related processes, as a continuous time process in t for fixed H which is sampled discretely at high frequencies. Hence, instead of block-sampling the quadratic variation once a day, as in Andersen and Bollerslev (1998), we construct rolling samples, say every 5 minutes, of integrated volatilities and rely on the Foster and Nelson results to establish optimal estimators for integrated volatilities. The daily (block sampling) scheme is a particular case of the generic context considered in this section. The advantage of the approach taken here is that the optimal weight design for integrated volatility estimators can be directly obtained with appropriate modifications from Foster and Nelson's analysis. In a first subsection we discuss the extension of the Foster and Nelson analysis to the quadratic variation process. Then, in a second subsection we discuss the link between rolling sample estimators for conditional spot volatility and its quadratic variation. In a final third subsection we cover the cumulative absolute return estimator.

2.1 Asymptotic distribution of Quadratic Variation Filters

We consider the integrated volatility process:

$$Q_t^H = \int_{t-H}^t \sigma_\tau^2 d\tau \quad (2.1)$$

which relates to the quadratic variation of a process. The latter is a semi-martingale (see for instance Shiryaev (1999, p. 303-304) for details) whenever the original returns process is a semi-martingale, an assumption made by Foster and Nelson. The process has its own natural (continuous time) filtration $\{\mathcal{F}_t^Q\}$. One can also consider the approach of Barndorff-Nielsen and Shephard (1999) who introduce a class of non-Gaussian Ornstein-Uhlenbeck processes for which one can easily characterize analytically the quadratic variation.

Throughout this section we will assume that the benchmark frequency is daily data. We will be interested in discretizations $Q_{(m_1, m_2), t}^H$ where the double index refers to the fact that we sample m_1 times a day integrated volatilities using sums of squared returns sampled m_2 times a day. We will at first look at the $Q_{(m, m), t}^H$ discretization and then examine the Andersen and Bollerslev scheme $Q_{(1, m), t}^H$. The former has a Doob-Meyer decomposition (see again Shiryaev (1999) for details):

$$Q_{(m, m), t}^H = \mu_{(m), t}^Q m^{-1} + M_{(m), t}^Q - M_{(m), t-1/m}^Q \equiv \mu_{(m), t}^Q m^{-1} + \Delta_{(m)} M_{(m), t}^Q$$

where the process $Q_{(m, m), t}^H$ is adapted to the filtration $\{\mathcal{F}_{(m), t}^Q\}$ which contains the filtration of discretely sampled spot volatilities. Extending the Foster and Nelson analysis to integrated volatilities suggests the following generic data-driven estimators:

$$\hat{Q}_{(m_1, m_2), t}^H = \sum_i w_{(m_1), i} \sum_j r_{(m_2), (t-i/m_1)-j/m_2}^2 \equiv \sum_i w_{(m_1), i} IV_{(m_2), t-i/m_1}^H \quad (2.2)$$

where $w_{(m_1),i}$ is a weighting scheme involving discrete (rolling) sampling at frequency $1/m_1$, of integrated volatilities $IV_{(m_2),t}^H$ defined as:

$$IV_{(m_2),t}^H \equiv \sum_{j=0}^{Hm_2} r_{(m_2),t-j}^2$$

which are based on sampling squared returns at a frequency $1/m_2$. It is important to note that the class of estimators defined in (2.2) is very large and contains many new estimators which were not considered in the literature. Moreover, the direct link with the design of spot volatility estimators makes their structure very transparent and attractive. Our main focus will start again with equal weighting schemes applied to a finite data window. Analogous to the analysis in the previous section, we have for flat weights the scheme characterized as $w_{(m_1),i} = m^{-1/2}(n_L + n_R)^{-1}I\{\tau \in [-n_L m_1^{-1/2}, n_R m_1^{-1/2}]\}$. An extension of Theorem 2 in Foster and Nelson yields that $m^{1/4}(\hat{Q}_{(m,m),t}^H - Q_t^H) \rightarrow N(0, C_{(m),t}^Q)$ as $m \rightarrow \infty$ where the continuous record asymptotic variance for a flat weighting scheme equals:

$$C_{(m),t}^{QF} \equiv \frac{\theta_{(m),t}^Q}{n_R + n_L} + \sqrt{\theta_{(m),t}^Q \Lambda_{(m),t}^Q} \rho_{(m),t}^Q \frac{n_R - n_L}{n_R + n_L} + \Lambda_{(m),t}^Q \frac{n_R^3 + n_L^3}{3(n_R + n_L)^2} \quad (2.3)$$

and the superscript QF refers to the equal weighted flat weighting scheme (applicable to the quadratic variation). The formal proof of the result in (2.3) is omitted as it is a straightforward extension of Theorem 2 in Foster and Nelson applied to the quadratic variations. The asymptotic variance of the normalized extraction error $m^{1/4}(\hat{Q}_{(m,m),t}^H - Q_t^H)$ has again three components, which have the same interpretation as those in the previous section, though applied to the quadratic variation. For instance, the process $\Lambda_{(m),t}^Q$ represents the “variance of the quadratic variation”, whereas $\theta_{(m),t}^Q$ represents the conditional second moment of the quadratic variation, which is implicitly related to the conditional fourth moment of the process. We have the same effects of the window parameters on the asymptotic efficiency as in the rolling sample estimation of conditional variances.

We noted in Section 1 that the regularity condition allowing us to treat higher moments as being roughly constant is not innocuous. We noted that for FX market volatility we have some evidence regarding the behavior of, say the conditional kurtosis, as one changes the sampling frequency. Most of these results apply to spot as well as integrated volatility. Therefore, we should be equally at ease (or not at ease) with the maintained regularity condition to compare integrated volatility estimators across different sampling frequencies.

2.2 The Relationship between Rolling Quadratic Variation and Spot Volatility Filters

Andersen and Bollerslev (1998) propose an estimator for integrated volatility which is based on a very different argument, namely an estimator of Q_t^H using the following approximation of the quadratic variation by a finite sum:

$$\text{plim}_{m \rightarrow \infty} \left(\int_{t-H}^t \sigma_\tau^2 d\tau - \sum_{j=1, \dots, m} r_{(m),t-j/m}^2 \right) = 0$$

which suggests (using our notation) the following estimator for the m sampling frequency estimated IV:

$$Q_{(1,m),t}^H \equiv IV_{(m),t}^H = \sum_{j=0}^{Hm} r_{(m),t-j}^2$$

In this subsection we would like to compare the $Q_{(1,m),t}^H$ estimator with $Q_{(m,m),t}^H$ proposed in the previous section. We are particularly interested in estimators which involve equal weighting schemes, i.e. $Q_{(m,m),t}^H$ involving weights $w_{(m),j} = 1/n_L$, which we will call for convenience historical integrated volatility denoted $HIV_{(m),t}^H$, namely:

$$HIV_{(m),t}^H = \frac{1}{n_L} \sum_{j=0}^{n_L-1} IV_{(m),t-j/m}^H.$$

There are two immediate relationships that can be established between HIV -type estimators and estimators previously used in the literature. First, it is straightforward to note that the case when n_L equals one corresponds to the IV estimator. More interestingly, we can also express the HIV estimator as an IV -type estimator replacing $r_{(m),t-j}^2$ by $\tilde{\sigma}_{(m),t-j}^2$, i.e. replacing squared returns by estimated *instantaneous* volatilities:

$$HIV_{(m),t}^H \equiv \sum_{j=1}^{Hm} \tilde{\sigma}_{(m),t-j}^2$$

where $\tilde{\sigma}_{(m),t-j}^2 = \sum_{l=1}^{l_s} \tilde{w}_{(m),l} r_{(m),t-j-l}^2$. In general it is not straightforward to find a direct relationship between the weights $\tilde{w}_{(m),j}$ involved in the estimation of the intra-daily spot volatility components of the daily integrated volatility and the weights of the HIV estimator, namely $w_{(m),i}$ in (2.2). It also follows that the weights for the spot volatility estimators are not particularly optimal weights. Nevertheless, one expects that the historical integrated volatility estimator will be more efficient since one replaces a noisy $r_{(m),t-j}^2$ by (albeit sub-optimal) $\tilde{\sigma}_{(m),t-j}^2$. The analysis in Andersen and Bollerslev (1998) indeed suggests that replacing $r_{(m),t}^2$ by volatility estimates yields a more efficient scheme.

There is one particular case where the relation between the historical integrated volatility weights and the spot volatility weights is rather straightforward. This is the case when n_L equals m , i.e. one takes a day's length of integrated volatilities. In this particular case the weighting scheme for the $\tilde{\sigma}_{(m),t-j}^2$ is also an equal weighting scheme of length m . Hence, one takes a daily sum of spot volatility estimates each involving equal weights covering one day of data points.⁷ This particular case is worth pursuing further since it allows us to say more about the comparison of the estimator $IV_{(m),t}^H$ of Andersen and Bollerslev and our historical integrated volatility estimator. In some sense, one can view the former as being a one-sided filter with one lag, i.e. $n_L m^{-1/2} = 1$, whereas the $HIV_{(m),t}^H$ has $n_L m^{-1/2} = m$ lags. Using (2.3) one can calculate the relative efficiency as the ratio of the asymptotic mean squared error of the HIV estimator relative to that of the IV estimator. This ratio of MSEs is:

⁷Note that the HIV estimator involves two days worth of data because the last component $\tilde{\sigma}_{(m),t-m}^2$ involves an entire extra day of data. This will be taken into account when we compare simulated data-driven filters.

$[3\theta^Q m^3 - 3\sqrt{\theta^Q \Lambda^Q} \rho^Q m^{7/2} + \Lambda^Q m^4] / [3\theta^Q - 3\sqrt{\theta^Q \Lambda^Q} \rho^Q m^{3/2} + \Lambda^Q m^3]$. Unfortunately, we cannot make any efficiency comparisons unless we specify θ^Q , Λ^Q and ρ^Q . This precludes us from calculating generic efficiency gains similar to the asymptotic equivalence results reported in Tables I and II.

It should also be noted that the discussion of optimal weights for extracting integrated volatility is easier to address when Q_t^H is viewed as a continuous time process since it leads us to the results of Foster and Nelson regarding optimal window and weighting schemes. Before we discuss these extensions let us note that it would not be easy to proceed from $Q_{(m,m),t}^H = \sum_{j=1}^{Hm} \tilde{\sigma}_{(m),t-j}^2$ and formulate optimal weighting schemes for $\tilde{\sigma}_{(m),t-j}^2$ for all $j = 1, \dots, Hm$. Indeed, it would be difficult to handle the intra-daily cross-correlation between spot volatilities. Approaching the design of optimal filter weights directly via (2.2) is relatively straightforward as it applies directly to the mean squared error of the quadratic variation instead of the mean squared error of the individual spot volatilities. Hence, we can apply Theorem 6 of Foster and Nelson which covers dominating flat weights (as noted before, using the design of formula (17) from Foster and Nelson). The resulting $C_{(m),t}^{QD}$ obviously depends again on $\rho_{(m),t}^Q$, $\theta_{(m),t}^Q$ and $\Lambda_{(m),t}^Q$. Last, but certainly not least, we can consider the optimal exponentially declining weighting scheme considered in Theorem 5 of Foster and Nelson.

2.3 Cumulative Absolute Returns

To conclude this section we should note that we can also investigate other processes in continuous time, such as:

$$CAR_t^H = \int_{t-H}^t |r_\tau| d\tau \quad (2.4)$$

which is the cumulative absolute return (CAR) process and relates to estimators of conditional volatility used by Taylor (1986), Hsieh (1991) and Schwert (1989). In order to avoid repetition we will only briefly mention that the analysis of the previous two subsections can be modified to handle CAR processes. Moreover, one can also derive by analogy the following estimator for the m sampling frequency estimated CAR:

$$CAR_{(m_1, m_2), t}^H = \sum_i w_{(m_1), i} \sum_j |r_{(m_2), (t-i/m_1)-j/m_2}| \equiv \sum_i w_{(m_1), i} CAR_{(m_2), t-i/m_1}^H \quad (2.5)$$

where $w_{(m_1), i}$ is again a weighting scheme involving discrete (rolling) sampling at frequency $1/m_1$, of Cumulative Absolute Returns $CAR_{(m_2), t}^H$ defined as:

$$CAR_{(m_2), t}^H = \sum_{j=0}^{Hm_2} |r_{(m_2), t-j}|$$

which are based on sampling at frequency $1/m_2$. Similar to $Q_{(m,m),t}^H$ we are particularly interested in $CAR_{(m,m),t}^H$ involving weights $w_{(m),j} = 1/n_L$, which we will call historical cumulative

absolute return denoted $HCAR_{(m),t}^H$, namely:

$$HCAR_{(m,m),t}^H \equiv HCAR_{(m),t}^H = \frac{1}{n_L} \sum_{j=0}^{n_L-1} CAR_{(m),t-j/m}^H.$$

As noted before, there is interest for this type of estimators when the process features conditional heavy tail behavior. In light of these arguments, and the arguments of Bai *et al.* (1999) noted earlier, we will also study these filters in the context of the Monte Carlo simulations and the empirical analysis. The novelty of the estimators in (2.5) is that we use sliding spans of cumulative absolute returns over fine sampling intervals, similar to the historical integrated volatility estimators. The filters of Taylor (1986), Hsieh (1991) and Schwert (1989) can again be viewed as special cases, similar to IV being a special case of HIV -type filters.

3 Monte Carlo Design

The objective of the Monte Carlo study is twofold. First we want to examine whether the predictions of the continuous record asymptotic theory describe adequately the sampling behavior of filters when applied to actual data. Therefore, we aim for a design tailored to (1) applications routinely found, and (2) predictions derived from continuous record asymptotics. The second scope of the Monte Carlo study is to examine the optimal filter design of integrated volatility estimators. We organize the section in subsections. The first subsection describes the processes we simulate. The second subsection covers issues pertaining to the implementation of the design. A third subsection introduces the various filters which will be considered. A fourth subsection presents the sample sizes and diagnostics.

3.1 Simulated Models

The models used for the simulation study are representative of the FX and equity financial markets, popular candidates of which are taken to be returns on DM/US\$, YN/US\$ exchange rates and S&P 500 stock index. We consider the following continuous time stochastic volatility model which is based on the results of Drost and Nijman (1993) and Drost and Werker (1996):

$$\begin{aligned} d \ln Y_t &= \sigma_t dW_{pt} \\ d\sigma_t^2 &= \theta(\omega - \sigma_t)dt + (2\lambda\theta)^{1/2}dW_{\sigma t}. \end{aligned} \tag{3.1}$$

The so-called GARCH diffusion yields exact GARCH(1,1) discretizations which are represented by the following equations:

$$\begin{aligned} \ln Y_t - \ln Y_{t-1/m} &= r_{(m),t} = \sigma_{(m)} z_{(m),t} \\ \sigma_{(m),t}^2 &= \phi_{(m)} + \alpha_{(m)} r_{(m),t-1/m}^2 + \beta_{(m)} \sigma_{(m),t-1/m}^2 \end{aligned} \tag{3.2}$$

where $z_{(m),t}$ is $NIID(0, 1)$ and $r_{(m),t}$ is the returns process sampled at frequency $1/m$. The diffusion parameters of (3.1) and the GARCH process parameters of (3.2) are related via formulas

appearing in Drost and Werker (1996, Corollary 3.2). Likewise, Drost and Nijman (1993) derive the mappings between GARCH parameters corresponding to processes with $r_{(m),t}$ sampled with different values of m . This allows us to estimate a GARCH process using, say daily data with $m = 1$, and computing the GARCH parameters $\alpha_{(m)}$, $\beta_{(m)}$, $\phi_{(m)}$, for any other frequency m as well as the diffusion parameters θ , ω and λ .

For the two FX series we take the results of Andersen and Bollerslev (1998), mainly for comparison purposes. The DM/US\$ parameter estimates are $\theta = 0.035$, $\omega = 0.636$ and $\lambda = 0.296$ and the YN/US\$ parameters are $\theta = 0.054$, $\omega = 0.476$ and $\lambda = 0.480$. Based on these the implied GARCH(1,1) parameters $\alpha_{(m)}$, $\beta_{(m)}$ and $\phi_{(m)}$ are computed for alternative intra-day frequencies. For instance, in the FX markets the daily, 30-minute, 5-minute and 1-minute frequencies are based on $m = 1, 48, 288$ and 1440 , respectively. These high intra-day sampling frequencies are intended to mimic the continuous record asymptotic analysis and to gauge its accuracy given certain frequencies encountered in practice. The computations of the GARCH parameters for alternative m are reported in Table III and were obtained using the software available from Drost and Nijman (1993).⁸

Following the same paradigm we consider an analogous example for the equity market with only a 6.5 hours of trading as opposed to the 24 hour trading in FX markets. We, therefore, estimate a GARCH(1,1) model for daily S&P 500 returns and using the above disaggregation results we consider the equivalent intra-day frequencies for $m = 1, 13, 78$ and 390 . The results in the top panel of Table III refer to the daily frequency Normal-GARCH(1,1) estimated parameters for the S&P 500 and the disaggregated models are also reported in the panels that follow. Note that the following sample sizes are considered: 02/04/86-29/08/97 ($T = 2884$) and the post stock market crash sample 04/01/88-29/08/97 ($T = 2443$).

In the light of the most widely early as well as recent empirical applications of data-driven volatility filters (outlined in the Introduction), we carry this analysis to the monthly frequency. We now aggregate the daily GARCH parameters for the monthly frequency using the approximation of 22 trading days per month (see for instance, French et al., 1987 and Schwert, 1989) to obtain the monthly GARCH parameters in the last panel of Table III.

3.2 Asymptotics and Reality

Certain adjustments are required in order to translate some of the continuous record asymptotic results into a meaningful simulation design. There are two issues we need to highlight.

First, it should be noted that when Foster and Nelson discuss spot volatility, they consider a measure of volatility normalized by the sampling interval, i.e. a measure of volatility per unit of time. We need to incorporate this in our simulation design in order to make for instance comparisons across sampling frequencies. We have tailored our discussion around a benchmark, or reference, frequency, i.e. $m = 1$, which refers to the daily and monthly frequencies. We will use the reference frequency as a benchmark to measure volatility. This implies that if

⁸Note that the daily frequency GARCH parameters, obtained from Andersen and Bollerslev (1998, Table 1, p.889), refer to the QMLE estimates under the assumption of conditional Normality for the daily FX returns from 01/10/87-30/09/92. The estimated daily GARCH(1,1) parameters for the extended sample until 30/11/96 (in the empirical section), does not produce significantly different estimation results.

we sample at, say 5-minute intervals and the benchmark frequency is daily we will actually rescale the 5-minute volatility estimates by the sampling frequency m so that they have a daily volatility interpretation. Schwert (1990b) is a practical example of a comparison of daily volatility computed from rescaled intra-daily returns (see Schwert (1990b, Figure D)). The rescaling will make some of the spot volatility estimators introduced in the next subsection look like integrated volatilities. Despite the rescaling we will still call them spot volatility estimators, as they are at the intersection between the classes of spot and integrated volatility estimators.

Second, the distinction between block and rolling sampling schemes may get blurred once we vary the sampling frequency. The continuous record asymptotic analysis involves ever shrinking sampling intervals. In practice, we deal with a discrete grid of points. Recall that for flat weights involving $n_L m^{-1/2}$ lags and $n_R m^{-1/2}$ leads the weighting scheme can be characterized as $w_{(m),t} = m^{-1/2}(n_L + n_R)^{-1}I\{\tau\epsilon[-n_L m^{-1/2}, n_R m^{-1/2}]\}$. Take for instance the case of a daily benchmark sampling frequency. In practice, a filter with $m = 1$ and $n_L = 26$ appears like a 26-day rolling sample estimator. However, when we shrink m such that $n_L m^{-1/2}$ is below one, i.e. the lag length is shorter than a day, then the filter starts to look like a block-sampling scheme since we measure spot volatility at the end of the day with only intra-day non-overlapping blocks of data. This will make the Monte Carlo simulations deviate somewhat from the ideal setup of the theoretical analysis. From the results in Foster and Nelson we know that block-sampling schemes tend to be less efficient, so we would expect that the calculations in Tables I and II will be less favorable for the block-sample intra-daily filters in our simulation experiment.

3.3 Data-Driven Volatility Filters

We define both spot and integrated volatility estimators. We begin by the former which are most widely used in practice and are covered in a first subsection, whereas integrated volatility estimators appear in a second subsection.⁹

3.3.1 Spot Volatility Filters

In light of the asymptotic properties of various types of data-driven spot volatility measures discussed in Section 1.1 we will consider the following selection of filters, starting with the daily benchmark frequency:

(1) *Exponentially Weighted Moving Average Volatility* $\hat{\sigma}_t^{RM}$ which is defined following the industry standard introduced by J.P. Morgan (see Riskmetrics Manual, 1995) as:

$$\hat{\sigma}_t^{RM} = \lambda \hat{\sigma}_{t-1}^{RM} + (1 - \lambda) r_{(1),t}^2 \quad t = 1, \dots, n_{days}. \quad (3.3)$$

where $\lambda = 0.94$ for daily data, $r_{(1),t}$ is the daily return and n_{days} is the number of trading days. Henceforth we will denote for convenience this filter by RM .

⁹The main simulation analysis does not assume conditional mean effects for the return process. Hence, we do not allow for them in the data-driven volatility definitions. Nevertheless, the study addresses the MA effects in intra-daily returns for the simulated filters.

(2) *One-day Spot Volatility* $\hat{\sigma}_t^{SV1}$ as the intra-daily (rescaled) mean of the log of squared returns $r_{(m),t}$ for different values of m , to produce the end-day spot volatility measure:

$$\hat{\sigma}_t^{SV1} = \sum_{j=1}^m r_{(m),t+1-j/m}^2 \quad t = 1, \dots, n_{days}. \quad (3.4)$$

where for the 5-minute sampling frequency the lag length take values, $m = 288$ for financial markets open 24 hours per day (e.g. FX markets) and $m = 78$ for a stock market open 6.5 hours per day. The filter will be denoted *SV1* although, as discussed in Section 3.2, this filter is rescaled to yield daily volatility estimates and therefore resembles an integrated volatility measure.

(3) *k-day Spot Volatility*, $\hat{\sigma}_t^{SVk}$, is a one-sided moving k -day average of $\hat{\sigma}_t^{SV1}$. The purpose of selecting k is to make efficiency comparisons with the daily return filters. In our analysis we will set $k = 2$ and 3 days. Here we will for convenience refer to *SVk*, where k takes the values 2 or 3.

(4) *One-sided Rolling daily window Volatility*, $\hat{\sigma}_t^{RV}$, defined as:

$$\hat{\sigma}_t^{RV} = \sum_{j=1}^{n_L} w_j \left(r_{(1),t+1-j} \right)^2 \quad t = 1, \dots, n_{days}. \quad (3.5)$$

where n_L is the lag length of the rolling window in days. When the weights w_j are equal to n_L^{-1} then one considers flat weights. Geometrically declining weights are $w_j = \exp(-\alpha j)$ with $\alpha = 0.0665$. In our simulations we will consider n_L equal to 26 and 52 days to conform with the optimality in Foster and Nelson and the common practice of taking (roughly) one month worth of data (see e.g. Schwert (1989) among others). We also consider the 52-day window to make an indirect comparison with the 26-day two-sided filter discussed in Foster and Nelson. Note that we mainly consider one-sided windows since we also wish to examine their long-run forecasting performance. This filter will be denoted *RV* followed by the number of lags involved in the filter. We will consider flat weights in the simulation and empirical design.

The above filters are also defined for the monthly benchmark frequency: In the one-month spot volatility, we have 22 trading days, which we extend to 2 and 3 months (or approximately 44 and 66 days). In the rolling monthly volatilities we define windows, n_L , of 12, 24 and 60 months (e.g. Officer (1973), Merton (1980), Campbell *et al.* (2000), Chan *et al.* (1999)). In addition we also consider in the daily case a two-sided filter, namely:

(5) *Two-sided rolling window*, $\hat{\sigma}_t^{RV-2}$, defined by:

$$\hat{\sigma}_t^{RV-2} = \sum_{j=1}^{n_L} w_j \left(r_{(1),t-j} \right)^2 + \sum_{j=1}^{n_R} w_j \left(r_{(1),t+j} \right)^2, \quad t = 1, \dots, n_{days} \quad (3.6)$$

where n_L lags and n_R leads are involved and alternative weighting schemes can be considered again. In accordance with the one-sided filters we take $n_L = n_R = 26$. This filter will be denoted *RV26 - 2*.

3.3.2 Integrated Volatility Estimators

Following the analysis of various types of data-driven integrated volatility measures discussed in Section 2 we will consider the following selection of filters:

(1) *One-day Integrated Volatility* $\hat{\sigma}_t^{IV1}$ (defined by Andersen et al., 1999) as the sum of the log of squared returns $r_{(m),t}$ for different values of m , to produce the daily volatility measure:

$$Q_{(1,m),t}^1 = \sum_{j=1}^m r_{(m),t+1-j/m}^2 \quad t = 1, \dots, n_{days}. \quad (3.7)$$

where for the 5-minute sampling frequency the lag length take values, $m = 288$, for financial markets open 24 hours per day (e.g. FX markets) and $m = 78$ for a stock market open 6.5 hours per day. We will refer to this filter as *IVk* with k equal to one.¹⁰ The *IV2* and *IV3* filters will also be considered and are defined analogous to *SVk*.

(2) *One-day Historical Integrated Volatility*, which was discussed at length in Section 2, namely (for $H = 1$ and $n_L = m$):

$$HIV_{(m),t}^1 = \frac{1}{m} \sum_{j=1}^m IV1_{(m),t+1-j/m}. \quad t = 1, \dots, n_{days}. \quad (3.8)$$

In the remainder of the paper we will refer to this filter as *HIVk* with k equal to one. The *HIV2* and *HIV3* filters will also be considered.

(3) *Exponentially weighted Historical Integrated Volatility*, in order to appraise the efficiency gains of using optimal weighting schemes to extract Q_t^1 we consider exponential declining weights, namely:

$$EHIV_{(m),t}^1 = A^{-1} \sum_{j=1}^{km} a^{-j} IV_{(m),t-1-j/m}^1. \quad t = 1, \dots, n_{days}. \quad (3.9)$$

where we select the decay rate a , equal to a range of values (.94, .96, .99) for daily filters. Finally, A is a scaling constant to guarantee that the filter weights sum to one. These filters will be denoted *EHIVk* with k equal to 1, 2 and 3 days.

By analogy the monthly integrated volatilities and historically integrated filters of $k = 1, 2$ and 3 months are based on an approximation of 22, 44 and 66 trading days, respectively.

3.3.3 Cumulative Absolute Return Estimators

Following the analysis of various types of data-driven spot and integrated volatility measures we will also consider the following selection of CAR filters:

¹⁰Andersen et al. (1999) present an additional transformation to this by looking at the logarithm of the standard deviation of the integrated volatility. Similarly, French *et al.* (1987) show that the logarithm of monthly IV estimates the skewness effects. We do not consider this transformation as we are primarily concerned with the MSE of extraction filters and not the distribution of normalized returns.

(1) *One-day Cumulative Absolute Return* used by Taylor (1986), Hsieh (1991) and Schwert (1989) as the sum of absolute returns $|r_{(m),t}|$ for different values of m , to produce the daily volatility measure:

$$CAR_{(m),t}^1 = \sum_{j=1}^m |r_{(m),t+1-j/m}| \quad t = 1, \dots, n_{days}. \quad (3.10)$$

This filter will be called *CAR1*. The *CAR2* and *CAR3* filters will also be considered and are defined analogous to *SVk* and *IVk*.

(2) *One-day Historical Cumulative Absolute Return*, discussed in Section 2, namely (for $H = 1$ and $n_L = m$):

$$HCAR_{(m),t} = \frac{1}{m} \sum_{j=1}^m CAR_{(m),t+1-j/m}^1. \quad t = 1, \dots, n_{days}. \quad (3.11)$$

and also extended to 2 and 3 trading days, similar to *HIVk*. By analogy the monthly CAR and historically CAR filters of $k = 1, 2$ and 3 months are based on 22, 44 and 66 days, respectively.

3.4 Sample Sizes and Diagnostics

In the Monte Carlo design we consider the following sample sizes, n :

n_{years}	n_{months}	n_{days}	24hrs FX Market			6.5hrs Equity Market		
			$n_{30 \text{ min.}}$	$n_{5 \text{ min.}}$	$n_{1 \text{ min.}}$	$n_{30 \text{ min.}}$	$n_{5 \text{ min.}}$	$n_{1 \text{ min.}}$
5	60	1,250	60,000	360,000	1,800,000	16,250	97,000	487,500
10	120	2,500	120,000	720,000	3,600,000	32,500	195,000	975,000

We assume that 1 year has 250 trading days. The monthly simulation analysis is based on a sample size of 30 years, often encountered in practice.¹¹ Each experiment is performed with 1000 replications. Note that for the one-sided rolling estimates we create sufficient data before the effective sample (equivalent to one year).

3.4.1 Daily and monthly benchmark frequency simulations

We begin with the daily benchmark frequency case. First we simulate $r_{(m),t}$ for the 1-minute frequency, $m = 1440$ and $m = 390$, based on the FX and equity GARCH(1,1) models in Table III, respectively. Hence, we obtain the highest intra-day sampling frequency which we consider to be the true generating process. Next, we apply the GARCH dynamics to obtain the spot volatility, $\sigma_{(m),t}$, which for $m = 1$ refers to the daily spot volatility. The extraction error is the difference between the simulated volatility, which is model-based, and the model-free data-driven spot volatility filters:

$$\varepsilon_t^i = \sigma_{(m),t} - \hat{\sigma}_t^i \quad (3.12)$$

¹¹With $n_{years} = 30$, we have $n_{months} = 360$ and $n_{days} = 4,320$.

for $i = RM, SV1, SV2, SV3, RV26$ and $RV52$. Obviously, for the rolling window schemes we have different lag lengths. Moreover, for the SV schemes we will consider the 24-hour market cases as well as the shorter equity trading opening hours. In each case we consider five and thirty minute intra-day sampling schemes with the samples given above. A similar analysis applies to the integrated volatility filters. Since we take $H = 1$ we will denote the quadratic variation by Q_t^1 . Then we can write the extraction error as:

$$\varepsilon_t^i = Q_t^1 - Q_{(m,m),t}^i \quad (3.13)$$

for $i = IVk, HIVk$ and $EHIVk, CARk, HCARk$ $k = 1, 2, 3$.¹² For the true daily integrated volatility to approximate the integral Q_t^1 we take the sum of one-minute instantaneous GARCH(1,1) estimates over a day. Hence, we take advantage of the exact weak GARCH aggregation properties of the data generating process to compute the true quadratic variation. We will consider the different cases arising from the 24-hour FX market and the shorter-trading equity market, as well as the 5-minute and 30-minute sampling frequencies.

The same analysis is followed in the monthly Monte Carlo design where we simulate the daily GARCH models for the FX and equity markets and apply the GARCH dynamics to obtain the monthly spot volatility whereas we aggregate the daily GARCH variance to obtain the true monthly integrated volatility. To avoid further complicating the notation we will denote the monthly spot volatility filters by $SV1, SV2, SV3, RM, RV12, RV24$, and $RV60$. It will always be clear from the context that we refer to the monthly filters when SVk and RM will be discussed. The monthly integrated volatility filters will be denoted by $IV22, IV44, IV66, HIV22, HIV44$, and $HIV66$. The simulated sample is 30 years.

3.4.2 Measures of appraisal

The behavior of the extraction error, ε_t^i , is examined according to the following four dimensions:

(a) We examine the efficiency of filters using the Mean Square Error (MSE) which we compare across different filters and sample sizes. Note that we also obtain the Mean Absolute Error (MAE) since Andersen, Bollerslev and Lange (1999) argue that this criterion is more robust than the RMSE which is susceptible to outliers.¹³ The relative efficiency of one filter vis-à-vis another is studied by computing ratios of MSE's. To facilitate comparison, the MSE of all spot volatility filters will be benchmarked against the MSE of the one-day spot filter. Likewise, the MSE ratios for the family of integrated volatility filters will be computed relative to the MSE of the one-day integrated volatility. We therefore obtain the following ratios:

$$MSE^i / MSE^{SV1} \quad or \quad MAE^i / MAE^{SV1} \quad (3.14)$$

$$MSE^i / MSE^{IV1} \quad or \quad MAE^i / MAE^{IV1} \quad (3.15)$$

where i refers to the MSE's obtained from $\hat{\sigma}_t^i$ for different windows and weights, i.e. for the daily case $i = RM, SVk, RV26$ and $RV52$ for spot volatility and the MSE's obtained from

¹²We take some liberty in (3.13) with regard to notation by using $Q_{(m,m),t}^i$ to facilitate the definition of extraction process.

¹³Andersen, Bollerslev and Lange also consider analogous statistics which are adjusted for heteroskedasticity and which are found to have significant improvements in the volatility forecasting analysis.

$Q_{(m_1, m_2), t}^i$ for $i = IVk, HIVk$ and $EHIVk$ with integrated volatility (where $k = 1, 2, 3$). This analysis extends to the monthly frequency where the benchmarks are the 12-month rolling volatility and the one-month integrated volatility.

(b) We study the out-of-sample forecast performance analogous to Andersen, Bollerslev and Lange (1998). Following Baillie and Bollerslev (1992) the h -period linear projection from the weak GARCH(1,1) model with returns that span $1/m$ day(s) is expressed as:

$$P_{(m), t}(r_{(1/h), t+h}^2) = m \cdot h \cdot \sigma_{(m)}^2 + (\alpha_{(m)} + \beta_{(m)}) \cdot [1 - (\alpha_{(m)} + \beta_{(m)})^{m \cdot h}] \cdot [1 - \alpha_{(m)} - \beta_{(m)}]^{-1} \cdot (\sigma_{(m), t}^2 - \sigma_{(m)}^2) \quad (3.16)$$

where $\sigma_{(m)}^2 \equiv \phi_{(m)} \cdot (1 - \alpha_{(m)} - \beta_{(m)})^{-1}$ and $\sigma_{(m), t}^2$ would be the alternative spot and integrated volatility filters analyzed above. We shall consider $h = 20$ days as in Andersen, Bollerslev and Lange (1998) and obtain the MSE and MAE for each 20 day (or long-run) out-of-sample volatility filter forecast. For the monthly experiment we consider $h = 12$ months. It is interesting to note that if we ignore parameter uncertainty, which is the case for our Monte Carlo simulations, we can view (3.16) as a functional transformation of $\sigma_{(m), t}^2$, and therefore the asymptotic distribution of the forecast MSE is easily obtained from the asymptotic distribution of the volatility estimator using the usual delta method. This is quite useful as we can easily compare MSE's of forecasts in empirical applications, whereas MSE's of filters can only be computed in a simulation context where the true data generating process is observed. We therefore consider the forecast MSEs as a bridge between the simulation-based results and the empirical results.

(c) The coefficient of multiple determination is obtained from the regression equation of each daily data-driven volatility against the daily GARCH(1,1) volatility as suggested by Andersen and Bollerslev (1998) in the spirit of the Mincer and Zarnowitz (1969) regression:

$$\hat{\sigma}_{(m), t}^i = a + b \cdot \sigma_{(m), t} + u_t, \quad (3.17)$$

where $i = RM, SV1, SVk, RV26$ and $RV52$. A similar regression for integrated volatilities is considered involving $Q_{(m_1, m_2), t}^i$ for $i = IVk, HIVk$ and $EHIVk$ on the left hand side and Q_t^1 as regressor.

(d) We study the inefficiency of filters by examining the cross-covariances between extraction errors and filtered volatilities. Non-zero cross-covariance between the extraction error and the corresponding filter imply that the filter does not fully exploit all the information in the sample. We will consider cross-covariances up to 5 days, i.e. $j = 1, \dots, 5$. Finally, we examine the MSE of the autocorrelation functions between the theoretical GARCH(1,1) and the autocorrelation coefficients of the alternative data-driven volatilities.

4 Monte Carlo Results

The Monte Carlo simulation results are analyzed in view of the theoretical extentions in sections 1 and 2 for the two categories of data-driven volatilities, spot and integrated.¹⁴ We

¹⁴Some simulation results also refer to cumulative absolute returns which we compare to integrated volatilities.

investigate whether the Monte Carlo design presents supportive evidence for the theoretical predictions that relate to the following four broad results. We examine whether there is simulation support of the data window length and sampling frequency equivalence discussed in Section 1.2 and particularly in Tables I and II. Moreover, we assess whether for a given family of data-driven volatilities, there exists a relatively most efficient data window length and sampling frequency and compare block-sample and rolling-sample volatility filters. Finally, we evaluate the optimality of alternative weighting schemes for data-driven volatilities.

There are many results and to try to keep the discussion clear we will start with spot volatility estimators in a first subsection, a second one deals exclusively with integrated volatility, a third presents some additional criteria of comparison and a final subsection covers the effects of MA and excess kurtosis and presents results on cumulative absolute return filters.

4.1 Spot Volatility Estimators

Table IV reports the Monte Carlo simulation results of the contemporaneous MSE (and MAE) ratios defined in (3.14) for spot volatility. For the daily instantaneous volatility MSE (and MAE) ratios, presented in the top panel, the benchmark is the 1-day Spot Volatility ($SV1$), whereas for the monthly spot volatilities, appearing in the lower panel, it is the 12-month Rolling Volatility ($RV12$ which is often used in practice). The daily spot volatilities are obtained from inter-day and intra-day (5-minute and 30-minute) sampling frequencies. Similarly, the monthly spot volatilities are defined in terms of monthly and daily frequencies. The theoretical results are based on MSE efficiency and hence we focus our discussion on this criterion, though it should be noted that the MAE results appear to be similar to the MSE findings. Similarly, the two intraday sampling frequencies of 5- and 30-minutes also sketch the same picture. Hence, in order to provide a concise discussion we focus on the 5-minute intraday results.

4.1.1 Data window length

We focus first on the lag length selection for spot volatility estimators. This covers all estimators except for the *RiskMetrics* filter RM since it does not involve any lag length selection. The contemporaneous MSE ratios for the daily spot volatilities (in Table IV) show that the MSE efficiency for the 2-day spot volatility $SV2$ based on the 5-minute frequency relative to that of $SV1$ varies between 0.84 (for S&P 500) and 0.92 (for YN/US\$). This means that adding one extra day, i.e. $SV2$ filtering versus $SV1$, may result in efficiency gains of up to 25 % (in the case of S&P 500). Adding a second day, i.e. comparing $SV1$ with $SV3$, we note that the ratios range from 0.78 (for S&P 500) to 0.91 (for DM/US\$). Hence adding two days can result in even higher efficiency gains, namely up to 28 %. We can also examine the ratio of $SV3/SV2$ and compute the efficiency gains of adding a day to $SV2$. In the case of equities measured by the S&P 500 index this is 0.93, which implies that adding the third day yields only a 7 % efficiency gain relative to the two-day spot volatility filter. We can do the same exercise with the rolling sample daily filters, i.e. compare $RV26$ and $RV52$. Again for the S&P 500 index this yields an efficiency gain of 85 %, i.e. doubling the length of the daily filter from

$RV26$ to $RV52$ has a more pervasive impact on efficiency. Foster and Nelson argue that for the S&P 500 index, a 26-day window is optimal.¹⁵ The daily spot volatility simulation results in Table IV show that $RV26$ is not optimal. The contemporaneous MSE (and MAE) of the S&P 500 daily $RV52$ is considerably more efficient compared to the one- and two-sided 26-day filter, a relative efficiency of 0.29 versus 0.80. Finally, the contemporaneous MSE ratios for the daily spot volatilities (in Table IV) show that the 2- and 3-day spot volatilities ($SV2$ and $SV3$) based on the 5-minute frequency are not as efficient as the 26-day or the 52-day rolling volatilities ($RV26$ and $RV52$), but only approximately so for $RV26 - 2$ in the S&P 500 case. On the other hand, it is worth noting that the MAE supports the asymptotic equivalence of efficiency.

We now turn to the comparison of monthly spot volatility filters. Unlike the contemporaneous MSE results in Table IV, it appears from the MAE ratios that all filters are asymptotically equivalent since their ratios are approximately equal to unity. This result suggests that for the monthly spot volatilities we can not identify a significantly optimal window length for rolling volatilities, as opposed to the daily spot volatilities. The MSE results are a bit more heterogeneous though the variation of the ratios is considerably smaller compared to the daily filters. For equity markets, if we exclude $SV1$ (and RM not discussed here), we note efficiency gains or losses of at most 15 %. Hence the Monte Carlo results suggest that the various filters encountered in the empirical literature which employ both daily and monthly frequencies to obtain data-driven monthly volatility estimates, (such as Officer (1974), Merton (1980), French *et al.* (1987), Schwert (1989, 1990a,b) and Chan *et al.* (1999)), are approximately equally efficient.

The results so far are based on the contemporaneous MSE criterion and the h -horizon MSEs may provide some additional useful insights. Recall that if we ignore parameter uncertainty, we can view (3.16) as a functional transformation of $\sigma_{(m),t}^2$, and therefore the asymptotic distribution of the forecast MSE is easily obtained from the asymptotic distribution of the volatility estimator. We also noted that we can consider the forecast MSEs as a bridge between the simulation-based results and the empirical results. We therefore examine whether high-frequency intradaily returns improve long-run volatility forecasts. Table V reports the MSE (and MAE) ratios for the h -day ahead forecast using (3.16) for spot volatilities. We also examine whether daily data can improve the long-run forecasting performance of monthly spot volatilities.

The optimality of $RV52$ for the S&P 500 is also supported by the MSEs of the long-run forecasts. In addition, Table V shows that the daily frequency is only marginally more optimal than the monthly one for forecasting monthly spot volatilities. Furthermore the long-run forecasts show that the two frequencies are asymptotically equivalent in terms of MSE (and MAE) efficiency given that all the ratios are approximately equal to one. Hence, in

¹⁵To be more precise, the window length is calculated from $\sqrt{3\theta}/\Lambda$ (see their Theorem 3). They find an estimate for θ equal to 2.72 and for Λ equal to 0.012, yielding n_L equal to 26. Because the generated process is GARCH, the optimal window (again according to Theorem 3) is a one-sided backward looking filter, i.e. $RV26$. Note that the sample covered by Foster and Nelson is 01/1928 - 12/1990, the same as French *et al.* (1987). We use a much smaller sample to fix the parameters 04/01/86 - 29/08/97 (T=2884 observations) and the post-1987 crash period 04/01/88 - 29/08/97 (T=2443). In the empirical section we will use a large sample, namely 03/01/1928 - 29/08/1997, and find that the parameter estimates don't change very much, indicating that our simulation experiment is reasonable to compare with the Foster and Nelson setup.

the monthly spot volatilities all window lag-lengths perform equivalently and for long-run forecasting the MSE and MAE criteria shows that both daily and monthly sampling frequencies are asymptotically equivalent. This result is consolidated or verified by both the h -period ahead forecasts and contemporaneous MSEs.

4.1.2 Comparing sampling frequencies

In Section 1 we presented theoretical results pertaining to the MSE asymptotic equivalence of the data window length for different sample frequencies of spot volatility estimators. For daily spot volatilities we found, from the summary results in Table I, that a 26-day window of daily data is asymptotically equivalent to a 1.5-day and 3-day window of 5-minute data, for the FX and equity markets, respectively. Similarly, a 52-day window of daily frequency is equivalent to a 3.1-day and 6-day window for the 5-minute spot volatilities in FX and equity markets, respectively. We examine whether the simulation results of the MSE ratios defined in (3.14) provide support for these theoretical predictions.

The analysis in Section 1 pertains to the comparison of, for instance, $SV3$ and $RV26$ which are asymptotically equivalent for equity markets. While the simulations are not fully supportive of this theoretical prediction, they tell us that the relative efficiency of $RV26/SV3$ for the S&P 500 index is 0.68. For FX markets $RV26$ and $SV1$ should be roughly equivalent, while the ratio in Table IV is about 0.57. It is interesting to note here that the results obtained with the MAE are closer to the asymptotic predictions. Namely, the $RV26/SV1$ ratio for FX markets is 0.81 whereas for equity markets $RV26/SV3$ is a comparable 0.79. We consider next the monthly spot volatilities. In light of the same theoretical predictions we expect that the 12-month spot volatility based on a monthly benchmark sampling frequency is approximately equivalent to a 3-month window of daily data. The second panel in Table IV refers to monthly simulation results of the contemporaneous MSE and MAE ratios. Note that these ratios use the 12-month rolling volatility ($RV12$) benchmark. The MSE ratios for both the FX and S&P 500 markets show that the 3-month spot volatility ($SV3$) is approximately MSE-equivalent (and MAE) with $RV12$, given that the ratios of $SV3/RV12$ are close to unity.

So far, the simulation results are mixed. We find that the inter-day 26-day rolling volatility filter is a more efficient filter, contrary to the predictions of the continuous record asymptotic analysis which predicted its equivalence to the 2- and 3-day intra-day spot volatilities for the 5- or 30-minute frequency in either the FX or equity markets. The MAE ratio provides slightly better results in terms of being closer to the asymptotic predictions. The results so far are based on the contemporaneous MSE criterion. The MSEs for long-run volatility forecasting of daily spot volatilities for S&P 500 in Table V are almost perfectly in line with the asymptotic predictions. The $RV26$ and $SV3$ daily equity market results are identical, both yielding a relative efficiency of 0.74 as predicted by asymptotic theory.¹⁶ For the monthly spot volatilities MSE (and MAE) ratios, the results also conform with the continuous record asymptotics predictions. Namely, in terms of long-run forecasting a 12-months window of rolling volatility based on monthly data is asymptotically equivalent on MSE (and MAE)

¹⁶In contrast, for FX they are again a bit out of line, namely $RV26$ and $SV1$ are not equivalent, the latter is about 30 % more efficient.

grounds to the 3-month window of block-sample volatility based on daily data. This result is consistent with the theoretical predictions based on Table I.

In an attempt to explain some of the differences between the theoretical and simulation results regarding the window length equivalence in spot volatilities, we must note the following observations: First, as discussed in Section 3.2, the comparison is based on different estimation methods for spot volatilities; RV's are based on rolling-sample estimation whereas the SV's are based on block-sample estimations. A more direct comparison would involve, for instance, a 3-day spot volatility based on rolling-sample (instead of block-sample) estimation of intraday frequency. Second, the continuous record asymptotics may not apply due to the invalidity of certain assumptions. Fortunately, this argument is not the explanation for the difference between the simulation results and the continuous records asymptotic findings. In Table III we report besides the GARCH parameter values also the kurtosis at the various sampling frequencies considered in the Monte Carlo simulations. Comparisons across $\kappa(m)$ for the various values of m indicate that, even for the monthly frequency, the kurtosis does not vary very much neither for equity nor for FX series. Hence, the variation of higher moments is most likely only a minor contributing factor that does not explain some of the deviations from the continuous record asymptotics.

4.1.3 Weighting schemes

The most efficient procedure for all the filtering methods covered in Table IV is one which we did not discuss so far. It is the *RiskMetrics* exponentially weighted moving average filter RM . The MSE (and the MAE) ratios show, with one exception, that RM is more efficient compared to the rolling volatility filters. The efficiency gains between, say RM and $RV26$ (daily) or RM and $RV24$ (monthly), can still be considerable. This observation brings us to the question of filter weights.

In the case of spot volatilities, Foster and Nelson show that the exponential weights perform better than the flat ones for rolling volatilities. In Table IV we observe that for the daily FX, the smallest contemporaneous MSEs are due to the *RiskMetrics* filter. For the DM/US\$ the RM represents a 25 % efficiency gain compared to $RV26$, which is the most efficient among the flat-weighted schemes. For the YN/US\$ the gain is even larger. For equity markets we find the same efficiency gains, comparing $RV26$ and RM , but the $RV52$ filter is more efficient. The relative efficiency of the *RiskMetrics* volatility is due to the fact that it is an EWMA filter which is consistent with the optimal exponential weighting scheme for rolling sample volatilities presented in Foster and Nelson. In the monthly spot volatility category, the RM performs equally well and is also found to be the relative most efficient volatility filter based on monthly frequency. It is interesting to re-evaluate the comparative performance of this filter for long-run forecasting which is relevant for practical applications such as Value-at-Risk (VaR). Indeed, the h -horizon MSE (and MAE) ratios in Table V show that RM is the relatively most efficient estimator for both daily and monthly spot volatilities in both markets and sampling frequencies.

4.2 Integrated Volatilities

It was argued in Section 2 that we could transplant the continuous asymptotic analysis for spot volatility processes to other processes, such as the integrated volatility. By analogy with the discussion of spot volatility we examine the integrated volatility results presented in Table VI for the contemporaneous MSE (and MAE) ratios (3.15). The 1-day and 22-day Integrated Volatilities ($IV1$ and $IV22$) are taken as the benchmarks for the daily and monthly integrated volatilities, respectively. We examine again the lag length of filters and their effect on efficiency first. Subsequently we compare the block-sample versus rolling-sample estimations of integrated volatilities.

4.2.1 Data window length and weighting schemes

We focus first on the comparison of $IV2$ and $IV3$ (relative to $IV1$), and hence examine the window length for IV type estimators. By analogy, we concentrate next on the comparison of $HIVk$ for different values of k . Recall that when the data span increases, namely when $n_R + n_L$ increases, then the first term on the right hand side of equation (2.3) decreases, whereas the third term increases with wider data spans, a result driven by the fact that only local cuts of the data exhibit a relatively stable variance. Consequently, increasing the data span may improve or harm efficiency. In Table IV, for spot volatility filters, we only found efficiency gains when the window lengths increased, though some of the gains were modest. We revisit the lag length issue here, and find in Table VI that indeed more (lags) is not always better. In general we find that for the S&P 500 index there are efficiency gains for moving from $IV1$ to $IV2$ and to $IV3$. Likewise, for $HIVk$ we find gains as k increases and for $EHIVk$ the same is true. The picture is quite the opposite for FX market series. Efficiency is reduced by half when we move from $IV2$ to $IV3$, both for the DM/US\$ and YN/US\$ series. The deterioration is even more dramatic for some of the HIV and $EHIV$ filters. For instance the MSE ratio for the YN/US\$ $EHIV1$ filter is 0.68, while that for the $EHIV3$ filter is 5.94. Clearly lag length selection is a very critical issue in the estimation of IV-type filters, contrary to spot volatility filters. From the asymptotic analysis this seems to indicate that the effect of $\Lambda_{(m),t}^Q$ on $C_{(m),t}^{QF}$ (taking the example of flat weights) can be considerable, suggesting that the variance of the IV process can change considerably, at least for FX markets. In contrast for equity markets and for spot volatility estimators (the later involving $\Lambda_{(m),t}$) the perverse effect of increasing lag lengths on efficiency appears to be negligible, suggesting also that the variance of variance for FX and equity markets is relatively stable for longer horizons.

We further investigate the optimal window length of integrated volatilities based on the MSE efficiency of long-run volatility forecasts. Table VII reports the MSE and MAE ratios from a 20-day ahead forecast applying (3.16) for integrated volatilities. The general picture is dominated by the fact that most MSE ratios are very close to unity for the two FX market series, which implies that the window length does not play an important role for the long-run forecasting performance of integrated volatilities. The exception to this result are the S&P 500 simulated MSE ratios for the 20-day forecasts. In the daily integrated volatilities, a window data length of 3 days seems to improve the long-run volatility forecasts by approximately 60% judged by the MSE criterion. This is the case for the $IV3$, $HIV3$ as well as $EHIV3$, all

compared to $IV1$. This result is consistent with the contemporaneous MSE ratios for the S&P 500 daily integrated volatilities.

The comparison of $HIVk$ and $EHIVk$ also enables us to appraise to what extent exponential weighting of IV-type estimators translates into efficiency gains. The optimal decay rate of 0.99 yields the lowest MSE among an a priori choice of 0.94, 0.96 and 0.99. The results in Table VI show that there are no significant (compare HIV1 and EHIV1 for YN/US\$) gains to be made by changing the weights to exponentially declining when MSE criteria are used. It should be noted however, that the MAE criteria present a different picture, namely at short lag lengths, i.e. when we compare $HIV1$ and $EHIV1$, there are efficiency gains when exponential weights are used for FX series. For the S&P 500 index the difference is negligible.

4.2.2 Block-sample versus rolling-sample volatility filters

A further interesting aspect of the integrated volatility simulation results is that we can compare directly block-sample versus rolling-sample volatility estimators for a given sampling frequency and a given data window length. We noted in Section 2 that we cannot a priori predict the efficiency of IV and HIV type integrated volatility estimators unless we know something about the behavior of higher moments. The daily integrated volatility simulation results in Table VI show that the lowest MSEs for the S&P 500 are the $IV3$, $HIV3$ and $EHIV3$ estimators. Between the first two estimators $HIV3$ is relatively more efficient since the MSE ratio of $HIV3/IV3$ suggests that $HIV3$ is 87% more efficient than $IV3$. Similarly, in the YN case the $EHIV1$ is the relatively most efficient estimator among all those studied. This filter is also defined in terms of rolling sampling techniques. Hence the S&P 500 and YN simulation results suggest that a historically integrated volatility estimator which involves rolling sample estimation is more efficient than block-sample type estimators for integrated volatility. However, for the DM series, the relatively most efficient estimator is the block-sample 2-day IV ($IV2$) which is 87% more efficient than $HIV1$ and 70% more efficient than $EHIV1$. Hence, it seems that for one out of three cases considered the block-sample estimation technique is more optimal than the rolling one for daily integrated volatilities. To be fair, it should be noted that $HIVk$ and $EHIVk$ estimators should be compared with IVm where $m = k + 1$, since the rolling scheme entails one extra day of data in a block-sampling context. Taking this into account we still find that $EHIV1$ is more efficient than $IV2$ for both FX series. It is also interesting to note that $(E)HIV2$ is also more efficient than $IV1$ in all cases (if we consider either weighting scheme for HIV-type estimators).

We further examine these estimation methods in the context of the monthly integrated volatility results (in the lower panel of Table VI). In all simulated cases we find that the 22-day block-sample estimator for monthly volatility ($IV22$) is the relatively least efficient estimator shown by all the MSE ratios being less than unity. On the other hand, the 22-day historical integrated volatility (HIV) appears, in all cases, to be more efficient than $IV22$. In particular, for the S&P 500, $HIV22$ is 35 % more efficient than $IV22$. Similarly, the contemporaneous MSE ratios for the DM and YN show that the $HIV22$ is 37% and 48% more efficient than $IV22$, respectively. Hence, we conclude that for monthly integrated volatilities the rolling sampling used by HIV is more efficient than the block-sampling IV for 22-days. However, a

more direct comparison is between HIV and IV filters both of which employ the same window length of 44 days, given by $IV44$ and $HIV22$. The former is based on a 44-day block-sample and the latter on 44-days rolling-sample estimation. From the simulation results we derive the MSE ratio of $IV44/HIV22$ which is 0.79 for the S&P 500, 0.91 for the DM and 0.91 for the YN. In the FX cases the $IV44$ is only marginally more efficient than $HIV22$. Hence, on average the monthly simulation results show that a window length of 44-day as opposed to a 22-day one is optimal for the estimation of monthly integrated volatilities. In addition, given this window length, FX block-sample monthly integrated volatilities are more efficient than historical ones, on MSE criteria. On the other hand, the S&P 500 results show that the $HIV66$ is the relative most efficient estimate on MSE terms.

4.3 Further evidence for the evaluation of volatility filters

So far we focused on two measures of appraisal, namely the Mean Square Error (along with the Mean Absolute Error), and the out-of-sample forecast performance. We turn now to the three remaining measures of appraisal, namely the coefficient of multiple determination in regression (3.17), the cross-covariances between extraction errors and filtered volatilities and the MSE of the ACFs between the GARCH and data-driven volatilities. Table VIII, top panel, presents the R^2 obtained from the regression in (3.17). The lower panel deals with integrated volatilities.

The daily spot volatility R^2 's show that the interdaily spot volatilities are highly correlated with the simulated volatility process defined by the Normal-GARCH(1,1) process. In all cases the daily RM has on average the highest R^2 being 0.8955, 0.9068 and 0.8058 in the S&P 500, DM and YN, respectively. Similarly, the RV filter for 52 and 26 days have equally high R^2 's (0.9199 and 0.807 for the S&P 500 and DM simulations). These results are consistent with the MSE ratios and provide additional support for the optimality of rolling volatilities with a given window length, as presented in Foster and Nelson. It is interesting to note the difference between the R^2 's presented by the inter- and intra-day spot volatilities. The latter show a relatively lower multiple correlation coefficient which may be the result of their more noisy nature, as opposed to the RM and RV smooth volatility estimates. Similarly, the monthly spot volatilities also show that the RM and $RV12$ have the highest R^2 's as opposed to the SV for 1-3 months. Nevertheless, comparative analysis shows that the level of the R^2 drops significantly for the monthly frequency, especially for the DM and YN cases.

The correlation coefficients of each data-driven volatility with the extraction error, (3.12), is an additional criterion for assessing the efficiency of the particular volatility filter. Negative correlations are obtained (which are not reported here for economy purposes) for lags 1-5 days and 1-5 months of the extraction error for the daily and monthly spot volatilities, respectively. The daily spot volatilities show that for the three cases, the $RV52$ and RM filters have the lowest correlations with the extraction error. Hence, they appear relatively more efficient than the intra-day volatilities. These results are also consistent with the MSEs conclusions. On the other hand, the monthly results suggest that the spot volatilities based on daily frequency are more efficient than those based on monthly frequency, since as the lag length increases the SV 's have the lowest correlations with the extraction error, for the two FX series.

A final criterion for comparing the data-driven volatility estimates is based on the MSE

of the autocorrelation functions (ACF) between the simulated Normal-GARCH(1,1) volatility and the alternative non-parametric volatility estimates. These results are also not reported for economy purposes but are available from the authors upon request. The lowest MSE of the autocorrelation function (ACF) between the theoretical GARCH(1,1) and different spot volatilities shows that the $RV52$ and the $RV60$ are the relatively most efficient filters for daily and monthly frequencies, respectively. They exhibit the lowest MSEs for a range of lags (5, 30, 60 and 120 minutes in the daily case and 1, 5, 20 days in the monthly case). Although most filters and especially the high frequency spot volatilities also exhibit low MSE for the ACFs, they do not seem to perform so well at longer lags of 120 minutes and 20 days for the intra-day and intra-month frequencies, respectively. In general, these results provide additional supportive evidence of the relative efficiency of the 52-rolling volatility which is in conjunction with the rest of results. Moreover, they present a new dimension for monthly spot volatilities for which we are now able to distinguish a relatively optimal filter, namely the $RV60$ which has on average the lowest ACF MSE (especially for the longer lag of 20 days). Hence, although the contemporaneous and long-run forecast MSEs showed that the monthly rolling volatilities are asymptotically equivalent on efficiency, the analysis of their ACF structure vis-à-vis the theoretical GARCH shows on one hand, that this argument is valid for short lags of 1 and 5 days, but on the other hand, that $RV60$ is more efficient for the longer lag of 20 days.

4.3.1 Integrated Volatilities

The lower panel of Table VIII presents the R^2 obtained from the regression (3.17) for integrated volatilities. In the daily frequency, all types of integrated volatilities seem to have very high R^2 's. These results extend the Andersen and Bollerslev (1998) empirical finding of a relatively high R^2 for the $IV1$ filter. In addition, the daily integrated volatilities R^2 's show that the filters with the highest R^2 are the ones which were found to be the most efficient based on MSEs. In particular, the highest R^2 's for the S&P 500 are the $HIV3$ and $IV3$, for the YN the $EHIV1$ and $IV1$ and for the DM the $IV2$ and $HIV1$. These results consolidate the evidence presented above. A similar picture is sketched by the monthly integrated volatilities which show that the highest R^2 's correspond to the MSE most efficient volatilities.

We also examine the correlation coefficients of each integrated volatility estimate with the extraction error, for lags 1-5 (also not reported for economy purposes). We find that in all cases these correlations have a relatively small size. This observation applies to all integrated volatilities and to both frequencies (daily and monthly) as well as to all three simulated cases. The correlation coefficients of all the IV's with the extraction error, range from 1% to 26%. The daily $IV1$ has the lowest correlation coefficients for all three series, followed by $HIV1$. Similarly, in the monthly case the lowest correlation coefficients apply to $IV22$ followed by $HIV44$. This result also holds for both the equity and FX markets. It is interesting to note is that although the historical integrated volatilities are also rolling regression volatilities, they have very low correlations with their extraction error.

The lowest MSE of the autocorrelation function (ACF) between the theoretical GARCH(1,1) and different types of integrated volatilities shows that among the class of daily block-sample integrated volatilities $IV1$ has the lowest MSE, whereas in the class or the rolling-sample in-

tegrated volatilities $HIV3$ and $EHIV1$ exhibit the lowest MSE. These simulation results hold for the S&P 500 but not for the FX series which show that all MSEs of ACFs are on average equal. Similarly, the MSEs of the ACFs integrated volatilities present the same behaviour. Generally, $IV22$ has the lowest MSE at lag 1 but for longer lags both monthly IVs and HIVs have similar MSEs. This evidence is in conjunction with the contemporaneous MSEs of the extraction error for the daily and monthly integrated volatilities reported above.

4.4 Tail behavior and MA components

The above Monte Carlo analysis is extended to a simulated process given by MA(1)-GARCH(1,1). Bai *et al.* (1999) suggest that intraday autocorrelation patterns in returns affect the MSE of spot volatilities (such as $SV1$ and the French *et al.* (1987) correlation adjusted). In particular they report that the MSE triples and may be even larger if the MA(1) coefficient is large or negative. Similar effects are reported for excess kurtosis which reduces the precision of spot volatility estimators by a half of the kurtosis coefficient. Therefore we perform the above simulation experiment for a range of MA(1) coefficients (0.2, ± 0.5 , 0.85) in the case of the S&P 500 models (reported in Table III). We also find that for a relatively high (0.85) and negative but moderate (-0.5) moving average coefficient, the MSEs ratios of all integrated volatilities rise and are almost equivalent. Hence it becomes difficult to differentiate between them. In contrast, for smaller moving average coefficients the results stay broadly the same with the initial analysis reported above. The existence of MA effects largely depends on the empirical features of the particular data frequency which we examine in the empirical section. In addition, the theoretical analysis can allow for the presence of excess kurtosis given that the fourth conditional moment remains constant for alternative intraday as well as daily vis-à-vis monthly frequencies. Hence, we also study the effects of excess kurtosis in the simulation design and examine the empirical validity of this assumption in the next section.

We first examine the effect of conditional leptokurtosis on spot volatilities. In the top panel of Table IX we report the contemporaneous MSE (and MAE) ratios of daily volatility filters. These are based on a 5-minute sampling frequency for the S&P 500 $t(\nu)$ -GARCH(1,1) model parameters (reported in Table III) and for alternative degrees of freedom, ν . Two interesting results arise: First, the relative efficiency of rolling volatility filters with respect to $SV1$ drops dramatically in the presence of excess kurtosis. This is also the case for the monthly frequency (reported in the lower panel of the table) as well in anticipation of the Foster and Nelson conclusions that the rolling estimator's efficiency drops in the presence of thick-tailed distributions. Yet, it is worth noting that $RV52$ still appears as the relatively most efficient filter. Hence we conclude that in the presence of excess kurtosis the MSE efficiency levels drop especially for rolling volatility estimators. The relative efficiency of intraday spot volatilities appears to be more robust to leptokurtosis. Second, the continuous record asymptotic predictions for the MSE (and MAE) efficiency equivalence are significantly weakened in the presence of a conditional Student's t process. Nevertheless, it is interesting that under conditional t (as opposed to Normal) the spot volatilities (especially at the monthly frequency) show significant variability on efficiency grounds. The daily sampled spot volatilities are more efficient than $RV12$, yet the 24- and 60-month rolling volatilities are more efficient than $RV12$ and the rest

of the daily based block-sample volatilities. The *Riskmetrics* filter consistently appears to be the relative most efficient volatility filter under both conditional Normality and Student's t distributed returns.

We now turn to examine the effect of conditional leptokurtosis on the relative efficiency of integrated volatilities and cumulative absolute returns. In Table X the first two panels report the daily and monthly integrated volatilities MSE (and MAE) ratios. Based on a comparison of the simulation results under conditional Normality and Student's t we can make several observations. First, the rolling-sample estimation method of integrated volatilities results in more efficient daily volatility filters ($HIV3$ is more efficient than $IV3$). Similarly, for the monthly frequency the $HIV66$ is the most efficient filter. It is interesting to point that the efficiency of $HIV66$ is robust to these alternative conditional distributional assumptions. Although the MSE ratio of $HIV66$ falls significantly in the presence of excess kurtosis, the $HIV22/IV22$ appears to maintain approximately similar levels of efficiency under both conditional normality and t distributions. Second, for the monthly frequency we find that block-sample integrated volatilities of alternative window lengths do not present any efficiency gains. This result is consistent under both conditional normality and Student's t distributions as shown by all IVk MSE (and MAE) ratios being close to one.

The last panel of Table X presents the Cumulative Absolute Returns MSE (and MAE) simulation results for the monthly S&P 500 under the assumption of conditional normality and Student's t with various degrees of freedom.¹⁷ The results show that the block-sampling estimation method using either absolute (or squared) returns produces MSE ratios close to unity. Instead, rolling-sample estimation in historical cumulative absolute returns (HCAR) produces relatively more efficient estimators. In fact, the HCARs are relatively more efficient than the HIVs (comparing the last two panels), for alternative degrees of leptokurtosis. Last but not least, the results in the last panel show the ratios of CAR type volatilities under both conditional normality and t . It appears that the efficiency of this type of volatilities remains robust to excess leptokurtosis. Hence the simulation evidence is supportive of the results in Taylor (1986), Davidian and Carroll (1987), Schwert (1990b), and extend the conjecture presented in the conclusions of Foster and Nelson (1996).

5 Empirical Illustration

The efficiency of data-driven volatility filters based on their out-of-sample forecast performance is evaluated for equity and FX series. The empirical results complement the simulation analysis by considering a number of horizons and sampling frequencies. Following Baillie and Bollerslev (1992), the h -period linear projection from the weak GARCH(1,1) model with returns that span $1/m$ day(s) is given in (3.16) where $\sigma_{(m),t}^2$ would be the alternative spot and integrated volatility filters.¹⁸ For daily volatility we consider $h = 1, 5, 20$ days and obtain the MSE and

¹⁷We do not report all the results in Table X. Unlike in Table IX we only focus on ν equal to 8 and 12. Results for $\nu = 10$ were deleted from the table to save space.

¹⁸Recently Bollen and Inder (1999) also undertake an empirical evaluation of various spot volatility estimators and $IV1$ using a number of other criteria.

MAE for each out-of-sample volatility filter forecast. Similarly, for monthly volatility filters we define $h = 1, 6, 12$ months.

The empirical illustration is based on two datasets representing the FX and equity market series, respectively. The five-minute intraday DM/US\$ and YN/US\$ returns cover a ten year period, 1/12/1986 to 30/11/1996, and were obtained from Olsen and Associates.¹⁹ Using this FX dataset we calculate intra-day volatilities for three frequencies (five-, fifteen- and thirty-minute) and compare them with the daily frequency volatility filters. The second dataset utilizes the daily Standard Poor’s composite price index to compare daily and monthly volatility filters which have been traditionally employed in the empirical literature. This dataset extends the sample used by Gallant, Rossi and Tauchen (1992) to cover the period 03/01/1928 - 29/08/1997 producing a sample of 18,571 daily observations. The empirical results are examined in conjunction with the Monte Carlo analysis and are also expected to present further empirical evidence of the theoretical asymptotic predictions analyzed in sections 1 and 2.

The estimation of AR(1)-GARCH(1,1) models for the S&P 500 and DM/US\$ data are reported in Table XI.²⁰ Although the estimation results are not directly comparable, we may observe how the sampling frequency affects higher moments. For the monthly and daily S&P 500 frequencies, it seems that the volatility persistence and the higher moments are broadly constant. On the other hand, the intraday frequencies show that there is a large variability in the kurtosis coefficient. As noted in Foster and Nelson (1996), Bai *et al.* (1999) and in Section 1 this is anticipated to affect the MSE’s of volatility estimators examined in the context of the empirical results below.

5.1 Spot Volatilities

We first consider the relative efficiency of empirical intra- and inter-daily volatility measures. The h -period ahead MSE and MAE ratios of spot volatilities derived from the linear projection formula in (3.16) are presented in Table XII for the short-run forecast horizon of one day ($h = 1$) and the longer-run horizons of one week and one month ($h = 5$ and 20 , respectively). Following the Monte Carlo analysis, the benchmark MSEs and MAEs, in the ratios (3.14) and (3.15), refer to the 1 day spot volatility forecasts. The analysis is performed for the two FX series and for three intra-day sampling frequencies (five-, fifteen- and thirty-minutes).

The results show that for both exchange rates the relatively most efficient daily volatility filter is the 52-day rolling volatility. This result holds for all the sampling frequencies and forecasting horizons considered. Hence, the empirical FX results are in broad agreement with the Monte Carlo simulations which show that the relatively most efficient daily filters are inter-day rolling volatilities (of 26 and 52 day window). Similarly, the *RiskMetrics* filter (RM) is also found empirically efficient for all h ’s, relative to all SV’s and $RV26$. This is also in

¹⁹The original sample is 1,052,064 five-minute return observations (2,653 days \cdot 288 five-minute intervals per day). The returns for some days were removed from the sample to avoid having regular and predictable market closures which affect the characterization of the volatility dynamics. For the description of the data removed refer Andersen, Bollerslev and Lange (1998) and Andersen, Bollerslev, Diebold and Labys (1999). The final sample includes 705,312 five-minute returns reflecting 2,449 trading days.

²⁰Similar results were obtained for the YN/US\$ series, and although we do not report the AR(1)-GARCH(1,1) estimations we, nevertheless, evaluate the efficiency of volatility forecasts using the YN/US\$ series.

harmony with the Monte Carlo results as well as the theoretically optimal exponential-type weights suggested in Foster and Nelson. The robustness of the MSE efficiency for the RM in both simulation and empirical paradigms reinforces its use for VaR.

We further compare the efficiency of spot volatilities using intra- and inter-daily frequencies. Recall that according to the FX theoretical results, a 26-day window of interdaily data is asymptotically equivalent in terms of efficiency with the 1.5-day window of 5-minute intra-day data, and with a 2.7-day window of 15-minute data, as well as with a 3.8-day window of 30-minute data. Similarly, a 52-day window of daily data is MSE equivalent to a window of 3.1, 5.3 and 7.5 days of intra-day 5-, 15-, 30-minute sampling frequency, respectively. The MSE ratios in Table XII show that the window length for spot volatilities plays a significant role. The intra-day 2 and 3 day spot volatilities ($SV2$ and $SV3$) are more efficient than the respective 1 day spot volatility. This observation is based on MSE and MAE criteria and applies to short-run ($h = 1$ day) and long-run ($h = 5, 20$ days) forecasting horizons. In addition, it holds for both the DM/US\$ and the YN/US\$ series and it can also be generalized to the three intra-day frequencies (5-, 15- and 30-minute data). In fact, the $SV3$ is the relatively most efficient intra-day spot volatility filter which is on average 65% more efficient than $SV1$. This result is consistent with the theoretical results, just mentioned, that for the intraday frequencies of spot volatilities a window length beyond that of a day and at least three days, is required to reach the asymptotic efficiency of a 26- and 52-day rolling volatility. The results in Table XII also show that the higher the sampling frequency (e.g. for 5-minute frequency as opposed to 15- and 30-minute), the more efficient is $SV3$ than $SV1$ and the MSE of $SV3/SV1$ for 20-day ahead forecasts decreases. The latter presents empirical support of the theoretical result, that the higher the frequency, the more efficient is $SV3$ for the long-run forecasting of spot volatility.

We observe that the MSE ratios decrease as we move from the 5-minute sampling frequency to the 15-minute. Note that this effect does not apply to the MAE which is regarded as a more robust criterion. Andersen, Bollerslev and Lange (1998) also documented this result (for $IV1$) and suggest that it may be due to microstructure effects.

This analysis is also performed for the traditional frequencies of volatility filters derived from daily and monthly data. The relevant MSE and MAE ratios in Table XIII use the 12-month rolling volatility (suggested by Officer (1974) and Merton (1980)) as a benchmark. The forecast MSEs for $h = 6, 12$ months show that for the S&P 500 the 12-month rolling volatility is more efficient than the 24- or 60-months rolling volatilities. On the contrary, the MAE ratios suggest that all spot volatilities based on monthly frequency ($RV12$, $RV24$, $RV60$ and RM) are approximately equivalent on MAE efficiency grounds. The latter result is in conjunction with the Monte Carlo simulations in Table VI. In addition, as mentioned, the MAE is more robust to outliers than the MSE. Hence, we present theoretical, simulation and empirical justification for the relative equivalence on efficiency grounds of alternative rolling volatilities in the empirical literature (e.g. Chan et al., 1999, Fleming *et al.*, 2000, Campbell *et al.*, 2000). The results in Table XIII also show that the monthly spot volatility based on the daily frequency for a 3-month window is a relatively efficient volatility filter in MSE and MAE terms. This is also in agreement with the simulation results.

5.2 Integrated Volatilities

We now examine the empirical efficiency of integrated volatility filters. First, we study the efficiency of the daily filters obtained from alternative h period ahead forecasts for the DM/US\$ and YN/US\$ intra-day sampling frequencies. These results are reported in Table XIII.

The Monte Carlo analysis showed that in the 5-minute sampling frequency, the MSE and MAE ratios for the 20-day ahead forecast of FX series is approximately one, implying that all these volatility measures have the same long-run forecasting performance. The notable exception was the daily S&P 500 simulation results which showed that all integrated volatilities have a lower MSE than $IV1$ (e.g. Andersen and Bollerslev, 1998). In particular, the 5-minute sample simulations showed that the $IV3$ and $HIV3$ were the relatively most efficient filters for the 20-day ahead forecast. Turning now to the empirical results in Table XIV we are able to make a direct comparison and obtain additional evidence using the 5-minute as well as other intra-day sampling frequencies. Both the DM/US\$ and YN/US\$ empirical MSE (and MAE) ratios are less than one, suggesting that $IV1$ is the relatively least efficient filter among the integrated filters considered. Therefore, the empirical results suggest that a window length beyond that of one day improves the empirical MSE volatility efficiency for both exchange rates, for short and long forecast horizons and for alternative intra-day sampling frequencies. This result is in conjunction with the simulation results of the extraction (rather than the forecast) error MSEs. Based on all the MSE ratios in Table XIV we conclude that the 3-day data window length is the relatively most efficient intra-day volatility filter for daily volatility forecasting. Consequently, the integrated and spot volatility results present a coherent picture in the sense that for the three intra-day frequencies considered a 3-day data window improves the MSE volatility efficiency for alternative forecast horizons.

The second dimension of the analysis refers to the estimation method of integrated volatility filters. For the given optimal window length of three days, we compare the MSEs of the block-sample filter, $IV3$, with the rolling-sample filter, $HIV3$. Both estimation methods suggest almost equivalent MSE efficiency levels, with $HIV3$ being only marginally more efficient than $IV3$. Extending this observation to other window lengths and sampling frequencies we conclude that both block and rolling-sample estimation methods are empirically optimal and that on average the rolling-sample is only marginally more efficient. This conclusion is in conjunction with the Monte Carlo simulations for the FX 5-minute sampling frequency in the context of which both IV 's and HIV 's were found to be asymptotically efficient. A related aspect of the estimation method refers to the choice of the weighting scheme. The exponential weighting scheme for a decay factor equal to 0.99 results in the lowest MSEs (a result also consistent and guided from the Monte Carlo analysis). Comparing the exponential and flat HIV 's, they perform approximately equally on MSE grounds, especially for the 15- and 30-minute sampling frequencies.

Andersen, Bollerslev and Lange (1998) also evaluate the efficiency of $IV1$ for alternative sampling frequencies and find that as the sampling frequency increases the efficiency of the filter drops. This result is not in agreement with the concept that the higher the frequency, the more efficient estimators are obtained. They find that the latter is empirically valid for hourly (and not 5- and 10-minute) sample frequency. We examine whether the results in Table XIV

also present a similar picture for the three sample frequencies. The 5-, 10- and 15-minute MAE ratios for the most optimal filters ($IV3$ and $HIV3$) do not signal any dramatic change in the efficiency lost from these high frequency samples. On the contrary, they seem quite similar. This observation generalizes to almost all integrated volatilities. The hourly frequency has also been examined (yet not reported) and it seems that those MSE/MAE ratios are only insignificantly lower, which is in agreement with the Andersen, Bollerslev and Lange findings.

Following the above analysis we also examine the empirical results for monthly S&P 500 integrated volatility filters based on daily frequency, reported in Table XV (top panel). Although the MSE ratios suggest that a window length of 3 months (approximately 66 days) is more optimal for both IV's and HIV's, the MAE ratios show that there are no significant gains in efficiency. The latter result is also consistent with the simulations of MSEs and MAEs obtained for the 12-month ahead forecasting horizon. Given the overall picture, we conclude that the window length choice does not seem to play an important role for forecasting monthly integrated volatilities based on daily data.

In an attempt to compare and consolidate our empirical results we extend the empirical analysis to other types of integrated volatilities, presented in the literature French *et al.* (1987) provide a correction for serial correlation in the estimation of block sample integrated volatilities. They add a second term to the IV which is the sum of the cross product of returns at t and $t+1$. Following their methodology we compute this estimator (denoted, $IV22 + SC2$) which is used as the benchmark. By analogy we extend this measure to allow for longer window lengths of 2 and 3 months (or approximately 44 and 66 days, respectively). The empirical MSE and MAE ratios presented in the second panel of Table XV present two interesting results: First, the MAEs show that alternative window lengths do not improve significantly the efficiency gains. However, on MSE grounds the 3 month window length for block sample integrated volatilities adjusted for serial correlation appears more efficient. These results are consistent with those obtained for the integrated volatilities in the first panel. The comparison of the MSE/MAE ratios, in the two panels, suggests that the correction for serial correlation by French *et al.* improves the efficiency of monthly S&P 500 data-driven volatility estimators using daily data. Second, the relatively most efficient estimator is the one month rolling sample volatility which corrects for serial correlation within that month ($HIV22+SC2$). It appears twice as efficient as the respective block sample estimator ($IV22+SC2$). Hence, the combination of the rolling estimation and the correction for serial correlation has shortened the optimal data-window length and produced the relatively most efficient estimator in terms of MSE.

Another data-driven volatility filter which is in the spirit of the cumulative quadratic variation of returns is the Cumulative Absolute Returns (CAR) filter (e.g. Hsieh, 1991). The CAR is also computed for 1-3 months and the MSE ratios are based on the 1-month (or approximately 22 days) CAR benchmark. The results are in broad agreement with the IV type estimators and therefore strengthen our empirical conclusions. Moreover, they provide additional support, as shown by the range of the MSEs being close to unity, that there are no significant gains from alternative daily data windows for estimating monthly cumulative quadratic variations of returns.

6 Conclusions

The paper presents a comprehensive theoretical, simulation and empirical analysis of the efficiency of spot and integrated types of volatilities. The theoretical results extend the Foster and Nelson (1996) analysis in the following ways. First, equally efficient spot volatility estimators are derived for alternative sampling frequencies based on the continuous record asymptotics. Second, the continuous record asymptotics arguments are applied to the new class of volatility estimators defined in terms of the cumulative quadratic variation of returns or cumulative absolute returns yielding various types of historical integrated volatilities. Third, the block- and rolling-sample estimation methods and alternative weighting schemes for both spot and integrated types of volatilities are discussed. The theoretical results are examined by an extensive Monte Carlo study and complemented by an empirical illustration.

The exponential weighting scheme is found to be the most optimal in terms of efficiency gains for spot volatilities. An example of that is the *RiskMetrics* filter which is an exponentially weighted moving average filter and seems to perform consistently well for both daily and monthly frequencies and for both markets in short-run and long-run volatility forecasting. These results support the theoretical argument that exponential weights for spot volatilities are expected to be more optimal. Examining the same argument for integrated volatilities we present empirical and simulation evidence that (for a priori) exponential weights, there are not substantial efficiency gains among the the flat-weighted and exponentially-weighted integrated volatilities.

Comparing the daily and monthly spot volatility simulation results on a number of criteria, we draw three broad conclusions. First, daily spot volatilities for 1 to 3 days based on 5-minute intra-day frequency are less efficient than inter-day volatilities, such as the rolling volatilities for 26- and 52-days. On the other hand, we find that monthly block-sample volatilities (based on daily frequency) are only marginally more efficient than the 60-month rolling volatility filter. Second, the data window length for rolling or block sample volatility estimation plays a more significant role in *daily* spot volatilities. The 52-day window is the relatively most optimal one for the S&P 500 (contrary to the much shorter filter which Foster and Nelson found to be optimal) and the 26-day window for the DM and YN. In contrast within the monthly spot volatilities the windows of 12, 24 and 60 months perform asymptotically equivalently on MSE grounds, for all the simulated cases with some evidence that the 60-month rolling volatility performs marginally better. Thirdly, the *RiskMetrics* filter which is consistently found to be the most efficient filter performs comparatively better for daily rather than monthly frequency.

Turning now to the comparison of the daily and monthly integrated volatility simulation and empirical results, we draw the following broad conclusions. First, the optimal window length for integrated volatilities extends beyond one day and one month. In particular, the optimal window length is 2 days for daily integrated volatilities and 2 months (equivalently 44 days) for monthly *IV*'s in the FX market simulations. This result also holds in the S&P 500 since the 3-day and 3-month windows are optimal for estimating daily and monthly integrated volatilities, respectively. Second, the weighting schemes of exponential and triangular weights appear to be equally optimal for the daily integrated volatilities. In general, historical integrated volatility estimates are found to be relatively efficient filters. Third, the integrated volatility filter attains

optimality for certain of the simulated cases, once the estimation window is extended to 2-days and 2-months for the daily and monthly integrated volatilities, respectively.

The effects of excess kurtosis on the relative efficiency of the above estimators are examined. We find that the MSE efficiency of spot estimators (and especially rolling volatilities) drops dramatically in the presence of excess kurtosis. In contrast, high frequency spot volatilities appear to be less susceptible to excess kurtosis. Nevertheless, the conclusions discussed above (found in the presence of conditional normality) are still generally valid. The effect of leptokurtosis on integrated volatilities show that the historical integrated volatilities are relatively more robust. It is interesting that a comparative analysis suggests that the efficiency of *CAR*-type estimators is more robust to alternative distributional assumptions (such as conditional normality and Student's t).

The broad conclusion from our results is the fact that window length, data frequency, weighting scheme and estimation methods of volatility filters play a relatively more important role for high frequency intra-day filters that are used to extract daily volatilities. The typical asset pricing applications involving monthly sampling frequencies are fairly insensitive to filter designs. In contrast, high-frequency data filters are very much dependent on judicious choices of the filtering scheme. For instance, we found that intra-day high frequency spot and integrated volatilities are non-optimal filter designs and do considerably worse than inter-daily filters such as rolling volatilities and the *RiskMetrics* filter. Moreover, we introduce *HIV* and *HCAR* filters which aim to improve the efficiency of high-frequency cumulative volatility filters. We also found that *CARs* and *HCARs* may be an improvement in the presence of excess kurtosis.

References

- [1] Andersen, T. and T. Bollerslev (1998) “Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts”, *International Economic Review*, 39, 885-905.
- [2] Andersen, T., T. Bollerslev, F. X. Diebold and P. Labys (1999), “The Distribution of Exchange Rate Volatility”, Discussion Paper NBER.
- [3] Bai X., J. Russell and G. Tiao (1999), “Beyond Merton’s Utopia: effects of non-normality and dependence on the precision of variance estimates using high-frequency financial data”, Discussion Paper, University of Chicago.
- [4] Baillie, R. T. and T. Bollerslev (1989), “The Message in Daily Exchange Rates: A Conditional-Variance Tale” *Journal of Business and Economic Statistics*, 7, 297-305.
- [5] Barndorff-Nielsen, O. E. and N. Shephard (1999), “Non-Gaussian OU Based Models and some of their Uses in Financial Economics”, Discussion Paper, Nuffield College, Oxford University.
- [6] Bollen B. and B. Inder (1999), “Ex Post, Unconditional Estimators of Daily Volatility”, Discussion Paper, Monash University.
- [7] Bollerslev T. and I. Domowitz (1993), “Trading Patterns and Prices in the Interbank Foreign Exchange Market”, *The Journal of Finance*, Vol. XLVIII, No.4.
- [8] Bollerslev, T., R. F. Engle and D. B. Nelson (1994), “ARCH Models”, in Robert F. Engle and Dan McFadden (eds.), *Handbook of Econometrics*, Volume IV, 2959-3038. Amsterdam: North-Holland.
- [9] Campbell, J. Y. and M. Lettau (1999), “Dispersion and Volatility in Stock Returns: An Empirical Investigation”, Discussion Paper NBER.
- [10] Campbell, J. Y., M. Lettau, B. G. Malkiel and Y. Xu (2000), “Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk”, Discussion paper Harvard University and Princeton University.
- [11] Chan, L. K.C., J. Karceski and J. Lakonishok (1999), “On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model”, *Review of Financial Studies*, 12, 937-974.
- [12] Chin, K., K.C. Chan and G. A. Karolyi (1991), “Intraday Volatility in the Stock Index Futures Markets”, *The Review of Financial Studies*, 4, 637-684.
- [13] Davidian M. and R.J. Carroll (1987), “Variance Function Estimation”, *Journal of the American Statistical Association*, 82, 1079-1091.
- [14] Drost, F. C. and T. Nijman (1993), “Temporal Aggregation of GARCH Processes”, *Econometrica*, 61, 909-727.

- [15] Drost, F. C. and B. M.J. Werker (1996), "Closing the GARCH Gap: Continuous Time GARCH Modeling", *Journal of Econometrics*, 74, 31-57.
- [16] Engle R. F., T. Ito and W-L. Lin (1990), "Meteor Showers or Heat Waves? Heteroskedastic Intra-Daily Volatility in the Foreign Exchange Market", *Econometrica*, 58, 525-542.
- [17] Fama, E. and J. D. MacBeth (1973) "Risk, Return, and Equilibrium: Empirical Tests", *Journal of Political Economy*, 81, 607-636.
- [18] Fleming, J., C. Kirby and B. Ostdiek (2000) "The Economic Value of Volatility Timing", *Journal of Finance*, (forthcoming).
- [19] Foster, D. and D. Nelson (1996) "Continuous Record Asymptotics for Rolling Sample Estimators", *Econometrica*, 64, 139-174.
- [20] French, K.R., G.W. Schwert and R.F. Stambaugh (1987), "Expected Stock Returns and Volatility", *Journal of Financial Economics*, 19, 3-29.
- [21] Gallant, A.R., P. Rossi and G. Tauchen (1992), "Stock Prices and Volume", *Review of Financial Studies*, 5, 199-242.
- [22] Hsieh, D.A. (1991), "Chaos and Nonlinear Dynamics: Application to Financial Markets", *Journal of Finance*, 46, 1839-1877.
- [23] Malkiel, B .G. and Y. Xu (1999), "The Structure of Stock Market Volatility", Discussion Paper Princeton University.
- [24] Meddahi, N. and E. Renault (1997) "Aggregation and Marginalization of GARCH and Stochastic Volatility Models", Discussion Paper GREMAQ, Toulouse.
- [25] Merton, R. C. (1980), "On Estimating the Expected Return on the Market: An Exploratory Investigation", *Journal of Financial Economics* 8, 323-361.
- [26] Officer, R.R. (1973), "The Variability of the Market Factor of the New York Stock Exchange", *Journal of Business*, 46, 434-453.
- [27] Poterba, J.M. and L. H. Summers, (1986), "The Persistence of Volatility and Stock Market Fluctuations", *American Economic Review*, 76, 1142-1151.
- [28] *Riskmetrics* Manual, (1995), J.P. Morgan Corporation, New York.
- [29] Schwert, G.W. (1989), "Why does Stock Market Volatility Change Over Time?", *Journal of Finance*, 44, 1115-1154.
- [30] Schwert, G.W. (1990a), "Stock Volatility and the Crash of '87", *Review of Financial Studies*, 3, 77-102.
- [31] Schwert, G.W. (1990b), "Stock Market Volatility", *Financial Analysts Journal*, 46, 23-34.

- [32] Schwert, G.W. and P.J. Seguin (1990), "Heteroskedasticity in Stock Returns", *Journal of Finance*, 45, 1129-1155.
- [33] Shiryaev, A. N. (1999), *Essentials of stochastic Finance: Facts, Models, Theory*, World Scientific Publisher.
- [34] Taylor, S. (1986), *Modeling Financial Time Series*, New York: Wiley.

Table I: Asymptotically Equivalent One-Sided Equal Weighting Schemes
for Volatility Filters

Frequency	m	Lags	Days FX	Days Eq.
Panel A: Equivalence to 22-day filter				
Half-daily	2	32	16	16
Hourly FX	24	108	4.5	-
Half-hourly Equity	13	80	-	6.1
Five-min. FX	288	374	1.3	-
Five-min. Equity	78	195	-	2.5
One-min. FX	1440	835	0.6	-
One-min. Equity	390	435	-	1.1
Panel B: Equivalence to 26-day filter				
Half-daily	2	37	18.5	18.5
Hourly FX	24	128	5.3	-
Half-hourly Equity	13	94	-	7.2
Five-min. FX	288	442	1.5	-
Five-min. Equity	78	230	-	2.9
One-min. FX	1440	987	0.7	-
One-min. Equity	390	514	-	1.3
Panel C: Equivalence to 30-day filter				
Half-daily	2	43	21.5	21.5
Hourly FX	24	147	6.1	-
Half-hourly Equity	13	109	-	8.3
Five-min. FX	288	510	1.8	-
Five-min. Equity	78	265	-	3.4
One-min. FX	1440	1139	0.8	-
One-min. Equity	390	593	-	1.5
Panel D: Equivalence to 60-month filter				
Daily	22	282	13	
Half-daily	44	398	9	
Hourly FX	528	1379	2.6	
Half-hourly Equity	286	1015	3.5	
Panel E: Equivalence to 12-month filter				
Daily	22	57	2.56	
Half-daily	44	80	1.81	
Hourly FX	528	276	0.5	
Half-hourly Equity	286	203	0.7	

Notes : The entries to the table report numerical calculations based on equation (1.4) using $C_{(1),t}^F$ evaluated at $n_R = 0$ and $n_L = 22$, $n_L = 26$ and $n_L = 30$ for the daily filters, and $n_L = 60$ and $n_L = 12$ for the monthly, as fixed. All asymptotically equivalent $C_{(m),t}^F$ filters require $n_L m^{-1/2}$ lags, e.g. $22 m^{-1/2}$.

Table II: Equivalent Historical and Benchmark Frequency Volatility Filters

Panel A: Daily Benchmark Frequency

Frequency	m	Historical Volatility Equivalent Number of days
Hourly FX	24	5
Half-hourly Equity	13	4
Five-min. FX	288	17
Five-min. Equity	78	9
One-min. FX	1440	38
One-min. Equity	390	20

Panel B: Monthly Benchmark Frequency

Frequency	m	Historical Volatility Equivalent Number of months
Daily	22	5
Hourly FX	528	23
Half-hourly Equity	268	17
Five-min. Equity	1716	41
Five-min. FX	6336	80

Notes : The entries to the table report numerical calculation based on equation (1.4) using $n_L = \sqrt{m}$ as the number of daily observations in a one-sided historical volatility filter which is asymptotically equivalent to a one-day (top panel) or one month (lower panel) volatility filter with sampling frequency $1/m$. All results are rounded off to the next integer.

Table III: GARCH(1,1) Models Used in Simulation Design

Parameters	Daily frequency			
	$m = 1$		$m = 1$	
	DM/US\$	YN/US\$	S&P 500 1986-97	S&P 500 1988-97
$\phi_{(m)}$	0.022	0.026	0.033130	0.282641
$\alpha_{(m)}$	0.068	0.104	0.028523	0.022832
$\beta_{(m)}$	0.898	0.844	0.967347	0.972683
$\kappa_{(m)}$	3	3	3	3
$v_{(m)}$	0.647	0.250	8.021792	62.39316
$\kappa_{(m)} \times (v_{(m)})^2$	1.256	0.750	193.0474	11678.72

Parameters	One-minute frequency			
	$m = 1440$		$m = 390$	
	$\phi_{(m)}$	0.0000155	0.0000185	0.0000851
$\alpha_{(m)}$	0.0026574	0.0041994	0.00165600	0.0013715
$\beta_{(m)}$	0.9973186	0.9957635	0.9983334	0.9986170
$\kappa_{(m)}$	2.6412370	2.4754561	2.8461378	2.8758057
$v_{(m)}$	0.645833	0.498652	7.736364	6.315652
$\kappa_{(m)} \times (v_{(m)})^2$	1.101663	0.615532	170.3451	111468.80

Parameters	Five-minute frequency			
	$m = 288$		$m = 78$	
	$\phi_{(m)}$	0.033130	0.282641	0.004260
$\alpha_{(m)}$	0.028523	0.022832	0.0036702	0.0030328
$\beta_{(m)}$	0.967347	0.972683	0.9962767	0.9969096
$\kappa_{(m)}$	3	3	2.8568004	2.8848233
$v_{(m)}$	8.021792	62.39316	8.022599	63.05208
$\kappa_{(m)} \times (v_{(m)})^2$	193.0474	11678.72	183.8695	11468.800

Table III Continued

Parameters	Thirty-minute frequency			
	$m = 2$		$m = 2$	
	DM/US\$	YN/US\$	S&P 500 1986-97	S&P 500 1988-97
$\phi_{(m)}$	0.007456	0.008891	0.0011064	0.009441
$\alpha_{(m)}$	0.0173169	0.027265	0.0058594	0.0048296
$\beta_{(m)}$	0.9815309	0.9709569	0.9940027	0.9950206
$\kappa_{(m)}$	2.7076174	2.5653969	2.971635	2.8946778
$v_{(m)}$	6.472222	5.000281	8.023205	63.02403
$\kappa_{(m)} \times (v_{(m)})^2$	113.4212	64.14213	191.28960	11497.74

Parameters	Monthly Frequency			
	$n = 22$ days		$n = 22$ days	
	$\phi_{(m)}$	0.34475	0.34556	0.69811
$\alpha_{(m)}$	0.05917	0.05944	0.07704	0.05526
$\beta_{(m)}$	0.40803	0.29313	0.83594	0.85058
$\kappa_{(m)}$	3.44359	3.77329	3.41591	3.26471
$v_{(m)}$	0.64707	0.49999	8.02241	63.01947
$\kappa_{(m)} \times (v_{(m)})^2$	1.44184	0.94331	219.844	12965.64

Note: The GARCH model parameters are, $\phi_{(m)}$, $\alpha_{(m)}$ and $\beta_{(m)}$, as defined in (3.2). The kurtosis parameter is $\kappa_{(m)}$. The unconditional variance is $v_{(m)} = \phi_{(m)} / (1 - \alpha_{(m)} - \beta_{(m)})$. The daily parameters (on the top panel) for the DM/US \$ and YN/US \$ were obtained from Andersen and Bollerslev (1998) and cover the period 01/10/87 - 30/09/92. The S&P 500 estimated parameters cover the daily samples 04/01/86 - 29/08/97 (T=2884 observations) and the post-1987 crash period 04/01/88 - 29/08/97 (T=2443). The disaggregated GARCH parameters, in the panels that follow, refer to the one-, five- and thirty-minute frequencies, which are computed via the software available from Drost and Nijman (1993). The final panel in the table refers to the aggregated daily GARCH parameters for monthly frequency using the approximation of 22 trading days per month.

Table IV: Monte Carlo Simulated Contemporaneous MSE and MAE Ratios of Spot Volatilities

Daily Frequency, Benchmark: $SV1$

	$SV2$	$SV3$	RM	$RV26$	$RV52$	$RV26 - 2$
S&P 500						
MSE	0.8422	0.7815	0.3914	0.5338	0.2884	0.7968
MAE	0.9275	0.8981	0.5883	0.7110	0.5509	0.8266
DM/US\$						
MSE	0.9477	0.9097	0.4508	0.5792	0.7497	0.6313
MAE	0.9812	0.9676	0.8208	0.8290	0.9286	0.8700
YN/US\$						
MSE	0.9182	0.8555	0.3844	0.5641	0.7447	0.6069
MAE	0.9689	0.9434	0.7601	0.8116	0.9544	0.9018

Monthly Frequency, Benchmark: $RV12$

	$SV1$	$SV2$	$SV3$	RM	$RV24$	$RV60$
S&P 500						
MSE	1.3200	1.1310	1.0500	0.7872	0.9474	0.9281
MAE	1.0030	0.9677	0.9525	0.9896	1.0070	1.0140
DM/US\$						
MSE	1.2210	1.0930	1.0360	0.8541	0.9091	0.8600
MAE	1.0240	1.0230	1.0230	1.0050	1.0010	1.0050
YN/US\$						
MSE	1.3140	1.3160	1.0560	0.8459	0.9025	0.8495
MAE	1.0240	1.0230	1.0230	1.0050	1.0010	1.0050

Note: The MSE and MAE are the mean square error and mean absolute error ratios, respectively, defined in (3.14), and obtained from the extraction error (3.12). The daily spot volatilities are: $SV1$, $SV2$ and $SV3$ are the 1-day, 2- and 3-days Spot Volatilities, respectively. RM is the textitRiskMetrics, $RV26$ and $RV52$ are the 26-days and 52-days one-sided Rolling Volatilities, respectively. $RV26 - 2$ is the two sided 26-days Rolling Volatility. The monthly spot volatilities are: $SV1$, $SV2$ and $SV3$ are the 1-month, 2-months and 3-months Spot Volatilities, respectively. $RV12$, $RV24$ and $RV60$ are the 12-, 24- and 60-months Rolling Volatilities, respectively. These daily simulation results refer to the 5 year sample and the 5-minute intraday frequency as reported in section 3.4. Similar results apply to 30-minute intraday frequency as well as the 10 year sample. The monthly simulation results refer to a 30 year sample size.

Table V: Monte Carlo Simulated 20-day and 12-month ahead MSE and MAE Ratios of Spot Volatilities

Daily Frequency, Benchmark: <i>SV1</i>						
	<i>SV2</i>	<i>SV3</i>	<i>RM</i>	<i>RV26</i>	<i>RV52</i>	
S&P 500						
MSE	0.8071	0.7357	0.5141	0.7374	0.4345	
MAE	0.9066	0.8702	0.6912	0.7899	0.6275	
DM/US\$						
MSE	0.9443	0.9158	0.9406	1.3500	1.6180	
MAE	0.9765	0.9654	0.9949	1.1550	1.2950	
YN/US\$						
MSE	0.9758	0.9714	1.249	1.408	1.607	
MAE	0.9872	0.9828	1.066	1.161	1.240	
Monthly Frequency, Benchmark: <i>RV12</i>						
	<i>SV1</i>	<i>SV2</i>	<i>SV3</i>	<i>RM</i>	<i>RV24</i>	<i>RV60</i>
S&P 500						
MSE	1.1390	0.9176	0.8338	0.6893	0.9283	0.9391
MAE	1.0010	0.9123	0.8798	0.9388	0.9992	1.0700
DM/US\$						
MSE	1.1050	0.9844	0.9362	0.8254	0.9121	0.8606
MAE	1.0280	0.9662	0.9397	0.8773	0.9412	0.8948
YN/US\$						
MSE	1.0150	0.9786	0.9650	0.9488	0.9749	0.9621
MAE	1.0110	0.9855	0.9743	0.9543	0.9745	0.9596

Note: The h -period ahead forecast MSE and MAE are the mean square error and mean absolute error ratios, respectively, defined in (3.14), and obtained from the extraction error in (3.12), using the h -period linear projection GARCH(1,1) equation in (3.16). The daily spot volatilities are: $SV1$, $SV2$ and $SV3$ are the 1-day, 2- and 3-days Spot Volatilities, respectively. RM is the textitRiskMetrics, $RV26$ and $RV52$ are the 26-days and 52-days one-sided Rolling Volatilities, respectively. The monthly spot volatilities are: $SV1$, $SV2$ and $SV3$ are the 1-month, 2-months and 3-months Spot Volatilities, respectively. $RV12$, $RV24$ and $RV60$ are the 12-, 24- and 60-months Rolling Volatilities, respectively. These daily simulation results refer to the 5 year sample and the 5-minute intraday frequency as reported in section 3.4. Similar results apply to 30-minute intraday frequency as well as the 10 year sample. The monthly simulation results refer to a 30 year sample size.

Table VI: Monte Carlo Simulated Contemporaneous MSE and MAE Ratios of Integrated

		Daily Frequency, Benchmark: <i>IV1</i>							
		<i>IV2</i>	<i>IV3</i>	<i>HIV1</i>	<i>HIV2</i>	<i>HIV3</i>	<i>EHIV1</i>	<i>EHIV2</i>	<i>EHIV3</i>
Volatilities		S&P 500							
	MSE	0.4517	0.2868	0.5999	0.3543	0.2534	0.6088	0.3897	0.3203
	MAE	0.6723	0.5359	0.7741	0.5956	0.5034	0.7801	0.6243	0.5678
		DM/US\$							
	MSE	0.6143	1.1210	0.7052	1.2560	2.0290	0.5860	1.427	2.635
	MAE	0.7825	1.054	0.8365	1.1160	1.4170	0.7712	1.132	1.516
		YN/US\$							
	MSE	1.091	2.643	1.190	2.964	4.947	0.6757	3.049	5.939
	MAE	1.040	1.615	1.085	1.713	2.220	0.8233	1.565	2.211

		Monthly Frequency, Benchmark: <i>IV22</i>				
		<i>IV44</i>	<i>IV66</i>	<i>HIV22</i>	<i>HIV44</i>	<i>HIV66</i>
		S&P 500				
MSE		0.2762	0.1265	0.3507	0.1372	0.1234
MAE		0.5270	0.3623	0.5870	0.3687	0.3509
		DM/US\$				
MSE		0.3346	0.4691	0.3673	0.5487	0.7718
MAE		0.5962	0.6989	0.6166	0.7513	0.8937
		YN/US\$				
MSE		0.4362	0.5630	0.4770	0.6360	0.8188
MAE		0.7242	0.7959	0.7507	0.8363	0.9155

Note: MSE and MAE are the mean square error and mean absolute error ratios, respectively, defined in (3.15), and obtained from the extraction error, (3.13). The daily integrated volatilities are: *IV1*, *IV2* and *IV3* are the 1-day, 2- and 3-days Integrated Volatilities, respectively. *HIV1*, *HIV2* and *HIV3* are the 1-day, 2- and 3-day Historical Integrated Volatilities, respectively. *EHIV1*, *EHIV2* and *EHIV3* are the 1-day, 2- and 3-days Exponentially Historical Integrated Volatilities, respectively. The decay rate in the exponential weighting scheme that minimizes the MSE/MAE is 0.99. The monthly integrated volatilities are: *IV22*, *IV44* and *IV66* are the 22-, 44-, and 66-days Integrated Volatilities, respectively, and, *HIV22*, *HIV44* and *HIV66* are the 22-, 44- and 66-day Historical Integrated Volatilities, respectively. These daily simulation results refer to the 5 year sample and the 5-minute intraday frequency as reported in section 3.4. Similar results apply to 30-minute intraday frequency as well as the 10 year sample. The monthly simulation results refer to a 30 year sample size.

Table VII: Monte Carlo Simulated 20-day and 12-month ahead MSE and MAE Ratios of Integrated Volatilities

Daily Frequency, Benchmark: *IV1*

	<i>IV2</i>	<i>IV3</i>	<i>HIV1</i>	<i>HIV2</i>	<i>HIV3</i>	<i>EHIV1</i>	<i>EHIV2</i>	<i>EHIV3</i>
S&P 500								
MSE	0.6893	0.5959	0.8051	0.6587	0.5964	0.8140	0.6949	0.6549
MAE	0.8304	0.7721	0.8980	0.8120	0.7720	0.9028	0.8327	0.8084
DM/US\$								
MSE	1.011	1.029	1.022	1.040	1.060	1.009	1.036	1.072
MAE	1.013	1.019	1.007	1.013	1.019	1.003	1.013	1.028
YN/US\$								
MSE	1.014	1.030	1.023	1.041	1.058	1.012	1.038	1.070
MAE	1.014	1.019	1.008	1.014	1.019	1.004	1.015	1.028

Monthly Frequency, Benchmark: *IV22*

	<i>IV44</i>	<i>IV66</i>	<i>HIV22</i>	<i>HIV44</i>	<i>HIV66</i>
S&P 500					
MSE	0.9521	0.9474	0.9804	0.9690	0.9714
MAE	0.9781	0.9748	0.9902	0.9824	0.9805
DM/US\$					
MSE	1.000	1.001	1.001	1.001	1.001
MAE	1.000	1.000	1.001	1.001	1.001
YN/US\$					
MSE	1.000	1.001	1.001	1.001	1.001
MAE	1.000	1.000	1.001	1.001	1.001

Note: The h-period ahead forecast MSE and MAE are the mean square error and mean absolute error ratios, respectively, defined in (3.15), and obtained from the extraction error in (3.13), using the h-period linear projection GARCH(1,1) equation in (3.16). The daily integrated volatilities are: *IV1*, *IV2* and *IV3* are the 1-day, 2- and 3-days Integrated Volatilities, respectively. *HIV1*, *HIV2* and *HIV3* are the 1-day, 2- and 3-day Historical Integrated Volatilities, respectively. *EHIV1*, *EHIV2* and *EHIV3* are the 1-day, 2- and 3-days Exponentially Historical Integrated Volatilities, respectively. The decay rate that minimizes the MSE/MAE is 0.99. The monthly integrated volatilities are: *IV22*, *IV44* and *IV66* are the 22-, 44-, and 66-days Integrated Volatilities, respectively and, *HIV22*, *HIV44* and *HIV66* are the 22-, 44- and 66-day Historical Integrated Volatilities, respectively. These daily simulation results refer to the 5 year sample and the 5-minute intraday frequency as reported in section 3.4. Similar results apply to 30-minute intraday frequency as well as the 10 year sample. The monthly simulation results refer to a 30 year sample size.

Table VIII: Monte Carlo Simulated R^2 of the Mincer-Zarnovitz Regression

Spot Volatilities							
Daily	$SV1$	$SV2$	$SV3$	RM	$RV26$	$RV52$	
S&P 500	0.4671	0.5210	0.5460	0.8955	0.8464	0.9199	
DM/US\$	0.4240	0.4542	0.4886	0.9068	0.8070	0.5244	
YN/US\$	0.4416	0.4848	0.5209	0.8058	0.6577	0.3498	
Monthly	$SV1$	$SV2$	$SV3$	RM	$RV12$	$RV24$	$RV60$
S&P 500	0.1913	0.2859	0.3403	0.6101	0.7752	0.5370	0.2140
DM/US\$	0.0325	0.1126	0.1410	0.1765	0.1919	0.0936	0.0334
YN/US\$	0.0248	0.1171	0.1327	0.1278	0.1343	0.0652	0.0231
Integrated Volatilities							
Daily	$IV1$	$IV2$	$IV3$	$HIV1$	$HIV2$	$HIV3$	
S&P 500	0.8140	0.9062	0.9379	0.8805	0.9251	0.9449	
DM/US\$	0.9716	0.9821	0.9670	0.9797	0.9634	0.9407	
YN/US\$	0.9809	0.9786	0.9479	0.9770	0.9421	0.9033	
Daily	$EHIV1$	$EHIV2$	$EHIV3$				
S&P 500	0.8838	0.9220	0.9342				
DM/US\$	0.9828	0.9578	0.9226				
YN/US\$	0.9870	0.9411	0.8867				
Monthly	$IV22$	$IV44$	$IV66$	$HIV22$	$HIV44$	$HIV66$	
S&P 500	0.6940	0.9169	0.9653	0.9070	0.9738	0.9636	
DM/US\$	0.7686	0.9164	0.9409	0.9794	0.7582	0.5514	
YN/US\$	0.7478	0.9110	0.7021	0.9770	0.7192	0.4915	

Note: The Mincer-Zarnovitz regression is defined in (3.17). The daily integrated volatilities are: $IV1$, $IV2$ and $IV3$ are the 1-day, 2- and 3-days Integrated Volatilities, respectively. $HIV1$, $HIV2$ and $HIV3$ are the 1-day, 2- and 3-day Historical Integrated Volatilities, respectively. $EHIV1$, $EHIV2$ and $EHIV3$ are the 1-day, 2- and 3-days Exponentially Historical Integrated Volatilities, respectively. The monthly integrated volatilities are: $IV22$, $IV44$ and $IV66$ are the 22-, 44-, and 66-days Integrated Volatilities, respectively and, $HIV22$, $HIV44$ and $HIV66$ are the 22-, 44- and 66-day Historical Integrated Volatilities, respectively. The daily volatilities refer to a five year sample and the monthly volatilities to a thirty year sample.

Table IX: Monte Carlo Simulated Contemporaneous MSE and MAE Ratios of Spot Volatilities under conditional Student's t (ν)

S&P500						
Daily Frequency, Benchmark: SV1						
	SV2	SV3	RM	RV26	RV52	
$\nu = 8$						
MSE	0.7915	0.6363	0.0405	0.0172	0.0014	
MAE	0.8957	0.8042	0.1654	0.1101	0.0262	
$\nu = 10$						
MSE	0.8217	0.6861	0.0556	0.0277	0.0016	
MAE	0.9148	0.8381	0.2018	0.1474	0.0303	
$\nu = 12$						
MSE	0.8507	0.7327	0.0668	0.0391	0.0018	
MAE	0.9303	0.8655	0.2228	0.1778	0.0312	
Monthly Frequency, Benchmark: RV12						
	SV1	SV2	SV3	RM	RV24	RV60
$\nu = 4$						
MSE	0.7852	0.7089	0.7193	0.1747	0.1613	0.4964
MAE	0.7659	0.7564	0.7737	0.3915	0.3736	0.6381
$\nu = 6$						
MSE	0.7165	0.7216	0.7364	0.1047	0.1726	0.3343
MAE	0.7920	0.8013	0.8136	0.2743	0.3592	0.5029
$\nu = 8$						
MSE	0.6702	0.6702	0.6748	0.1394	0.3624	0.2162
MAE	0.7893	0.7933	0.7981	0.3291	0.5506	0.3986

Note: The MSE and MAE are the mean square error and mean absolute error ratios, respectively, defined in (3.14), and obtained from the extraction error (3.12). The simulated process refers to a GARCH(1,1) with conditional Student's t distribution and degrees of freedom ν , where ν is defined as 8, 10 and 12 for the daily frequency and 4, 6 and 8 for monthly frequency. The parameters of the GARCH(1,1) model for the one-minute S&P 500 are defined in Table III. The daily spot volatilities are: SV1, SV2 and SV3 are the 1-day, 2- and 3-days Spot Volatilities, respectively. RM is the RiskMetrics, RV26 and RV52 are the 26-days and 52-days one-sided Rolling Volatilities, respectively. The monthly spot volatilities are: SV1, SV2 and SV3 are the 1-month, 2-months and 3-months Spot Volatilities, respectively. RV12, RV24 and RV60 are the 12-, 24- and 60-months Rolling Volatilities, respectively. These simulation results refer to the 5 year sample and the 5-minute intraday frequency as reported in section 3.4.

Table X: Monte Carlo Simulated Contemporaneous MSE and MAE Ratios of Integrated Volatilities under conditional Student's t (ν) for S&P 500

Daily Frequency, Benchmark: SV1								
	IV2	IV3	HIV1	HIV2	HIV3	EHIV1	EHIV2	EHIV3
$\nu = 8$								
MSE	0.3661	0.1114	0.4024	0.1178	0.0620	0.4521	0.1677	0.1106
MAE	0.6154	0.3235	0.6230	0.3244	0.2080	0.6664	0.3818	0.2980
$\nu = 12$								
MSE	0.3982	0.1616	0.4741	0.1872	0.1061	0.5124	0.2339	0.1616
MAE	0.6396	0.3985	0.6786	0.4203	0.2991	0.7092	0.4672	0.3792
Monthly Frequency, Benchmark: IV22								
	IV44	IV66	HIV22	HIV44	HIV66			
$\nu = 4$								
MSE	1.027	1.048	0.3286	0.076	0.025			
MAE	1.014	1.024	0.4932	0.2511	0.1296			
$\nu = 8$								
MSE	1.003	1.008	0.1168	0.0496	0.0339			
MAE	1.002	1.004	0.2680	0.1894	0.1531			
Monthly Frequency, Benchmark: CAR22								
	CAR44	CAR66	HCAR22	HCAR44	HCAR66			
Cond. Normality								
MSE	0.999	0.999	0.0692	0.0660	0.0658			
MAE	1.000	1.000	0.2561	0.2554	0.2554			
$\nu = 4$								
MSE	1.012	1.023	0.1468	0.0543	0.0193			
MAE	1.006	1.011	0.3412	0.2215	0.1226			
$\nu = 8$								
MSE	1.002	1.004	0.0538	0.0335	0.0257			
MAE	1.001	1.002	0.1976	0.1705	0.1485			

Note: The simulated process refers to a GARCH(1,1) with conditional Student's t distribution and degrees of freedom ν , where ν is defined as 8 and 12 for the daily frequency and 4 and 8 for monthly frequency. The parameters of the GARCH(1,1) model for the daily S&P 500 are defined in Table III. MSE and MAE are the mean square error and mean absolute error ratios, respectively, defined in (3.15), and obtained from the extraction error, (3.13). The daily integrated volatilities are: IVk , for k -day Integrated Volatilities. $HIVk$ are the k -day Historical Integrated Volatilities. $EHIVk$ are the k -day Exponentially Historical Integrated Volatilities. The decay rate in the exponential weighting scheme that minimizes the MSE/MAE is 0.99. The monthly Cumulative Absolute Returns (CAR) refer to the $CARk$ defined as the k -day Cumulative Absolute Returns and $HCARk$ the k -day Historical Cumulative Absolute Returns based on rolling samples of CAR. These simulation results refer to the 5 year sample and the 5-minute intraday frequency as reported in section 3.4.

Table XI: Empirical GARCH(1,1) Model Estimates

Parameters	Monthly	Daily		Intra-day Frequency	
	S&P 500	S&P 500	DM/US\$	15-min. DM/US\$	30-min. DM/US\$
$\gamma_{(m)}$	0.1659 (0.0348)	0.0389 (0.0059)	-0.0108 (0.0185)	0.0011 (0.0015)	0.0007 (0.0023)
$\delta_{(m)}$	-0.0679 (0.0394)	0.1477 (0.0084)	-0.0019 (0.0212)	-0.0655 (0.0026)	-0.0406 (0.0038)
$\phi_{(m)}$	0.0289 (0.0147)	0.0089 (0.0012)	0.0161 (0.0055)	0.0515 (0.0009)	0.08633 (0.0018)
$\alpha_{(m)}$	0.0720 (0.0219)	0.0728 (0.0042)	0.0462 (0.0075)	0.2032 (0.0029)	0.2221 (0.0039)
$\beta_{(m)}$	0.9000 (0.0273)	0.9201 (0.0046)	0.9378 (0.0111)	0.7836 (0.0027)	0.7335 (0.0037)
$\kappa_{(m)}$	4.6047	7.08292	4.6309	138.0416	70.5397
$v_{(m)}$	1.0321	1.2535	1.0312	3.8772	1.9444
$\kappa_{(m)} \times (v_{(m)})^2$	4.9055	11.1295	4.6599	2057.099	226.6804

Notes: The results refer to the QMLE of and AR(1)-GARCH(1,1) model. The $\gamma_{(m)}$ and $\delta_{(m)}$ are the constant and autoregressive coefficients in the AR(1). The GARCH model parameters are, $\phi_{(m)}$, $\alpha_{(m)}$ and $\beta_{(m)}$, as defined in (3.2). The estimated kurtosis parameter is $\kappa_{(m)}$. The unconditional variance is $v_{(m)} = \phi_{(m)} / (1 - \alpha_{(m)} - \beta_{(m)})$. The S&P 500 is the daily Standard Poor's composite returns covering the daily sample period of 03/01/1928 - 29/08/1997, T=18,751 daily observations. These were aggregated for every 22 trading days to approximate the number of trading days within a month. The DM/US \$ returns covers the 5-minute sample period of ten years, 01/12/86 - 30/11/96, (T=2,653 days), obtained from Olsen and Associates. The final sample removes a number of days as described in section 5 and in Andersen, Bollerslev and Lange (1998). The final sample includes 2,449 days or 705,312 5-minutes. These were aggregated to obtain the 15- and 30-minutes samples of 235,104 and 117,552 observations, respectively.

Table XII: Empirical Results of h -day ahead MSE and MAE Ratios of Spot Volatilities

	Benchmark: $SV1$				
	$SV2$	$SV3$	RM	$RV26$	$RV52$
Five-minute Sampling Frequency					
DM/US\$					
$h = 1$ day					
MSE	0.8353	0.7728	0.6182	0.6541	0.5834
MAE	0.9585	0.9382	0.8685	0.8950	0.8355
$h = 5$ days					
MSE	0.8203	0.7526	0.5848	0.6239	0.5513
MAE	0.9542	0.9328	0.8577	0.8898	0.8238
$h = 20$ days					
MSE	0.7473	0.6543	0.4410	0.4960	0.4225
MAE	0.9295	0.8946	0.8063	0.8625	0.7913
YN/US\$					
$h = 1$ day					
MSE	0.8520	0.7858	0.7387	0.7832	0.6984
MAE	0.9599	0.9380	0.9466	0.9766	0.9129
$h = 5$ days					
MSE	0.8372	0.7653	0.7109	0.7596	0.6719
MAE	0.9560	0.9325	0.9416	0.9766	0.9102
$h = 20$ days					
MSE	0.7741	0.6756	0.5673	0.6281	0.5433
MAE	0.9348	0.9014	0.9156	0.9664	0.8998
Fifteen-minute Sampling Frequency					
DM/US\$					
$h = 1$ day					
MSE	0.8288	0.7646	0.5616	0.5924	0.5298
MAE	0.9568	0.9369	0.8171	0.8423	0.7859
$h = 5$ days					
MSE	0.8129	0.7433	0.5277	0.5634	0.4971
MAE	0.9532	0.9318	0.8042	0.8347	0.7720
$h = 20$ days					
MSE	0.7339	0.6367	0.3813	0.4315	0.3639
MAE	0.9289	0.8918	0.7400	0.7941	0.7249

Table XII continued

		YN/US\$				
<i>h</i> = 1 day						
MSE		0.8544	0.7903	0.6521	0.6916	0.6164
MAE		0.9591	0.9389	0.8736	0.9015	0.8424
<i>h</i> = 5 days						
MSE		0.8399	0.7701	0.6230	0.6662	0.5884
MAE		0.9555	0.9329	0.8646	0.8975	0.8357
<i>h</i> = 20 days						
MSE		0.7713	0.6728	0.4771	0.5323	0.4555
MAE		0.9353	0.9030	0.8191	0.8702	0.8074
		Thirty-minute Sampling Frequency				
		DM/US\$				
<i>h</i> = 1 day						
MSE		0.7918	0.7225	0.4817	0.5098	0.4544
MAE		0.9537	0.9305	0.7797	0.8040	0.7500
<i>h</i> = 5 days						
MSE		0.7746	0.6999	0.4468	0.4773	0.4207
MAE		0.9497	0.9252	0.7641	0.7935	0.7338
<i>h</i> = 20 days						
MSE		0.6912	0.5905	0.3013	0.3428	0.2869
MAE		0.9246	0.8848	0.6895	0.7407	0.6751
		YN/US\$				
<i>h</i> = 1 day						
MSE		0.8223	0.7531	0.5754	0.6103	0.5438
MAE		0.9607	0.9401	0.8401	0.8671	0.8100
<i>h</i> = 5 days						
MSE		0.5434	0.8061	0.7310	0.5814	0.5128
MAE		0.9577	0.9350	0.8299	0.8619	0.8016
<i>h</i> = 20 days						
MSE		0.7292	0.6258	0.3916	0.4394	0.3702
MAE		0.9370	0.9062	0.7751	0.8293	0.7639

Notes: The *h*-day ahead forecast MSE and MAE are the mean square error and mean absolute error ratios, respectively, defined in (3.14) and obtained from the extraction error (3.12) using the *h*-day linear projection GARCH(1,1) equation in (3.16). The forecast horizons of *h*=1,5,20 days represent the one day, week and month, respectively. The daily GARCH(1,1) parameters for each series were obtained in Table III (top panel) mainly for consistency purposes with the simulation results and for comparison with other studies. The DM/US \$ and YN/US \$ returns cover the ten year period, 01/12/86 - 20/11/96 (removing some days as described in section 5). The final daily sample of 2,449 days covers 705,312 5-minute observations. These were aggregated to obtain the 15- and 30-minutes samples of 235,104 and 117,552 observations, respectively. The data source is Olsen and Associates. The daily spot volatilities based on intraday frequency are the 1-, 2- and 3-days spot volatilities, *SV*1, *SV*2 and *SV*3, respectively. The spot volatilities based on daily frequency are the textitRiskMetrics (*RM*) and the one-sided 26- and 52-days Rolling Volatilities (*RV*26 and *RV*52, respectively).

Table XIII: Empirical Results of h -month ahead MSE and MAE Ratios of Spot Volatilities

	Monthly Frequency, Benchmark: RV_{12}					
	$SV1$	$SV2$	$SV3$	RM	RV_{24}	RV_{60}
	S&P 500					
$h = 1$ month						
MSE	1.2519	0.9176	0.8163	0.7915	0.9069	0.9830
MAE	0.7970	0.7990	0.8036	1.1124	1.0388	1.1856
$h = 6$ months						
MSE	1.5649	0.9812	0.8331	0.9208	1.1292	1.4024
MAE	0.8857	0.8226	0.8093	1.1591	1.1234	1.3523
$h = 12$ months						
MSE	1.3327	1.1262	1.0838	1.2025	1.3019	1.4613
MAE	1.1297	1.0759	1.0568	1.0063	1.0632	1.1286

Note: The h -month ahead forecast MSE and MAE are the mean square error and mean absolute error ratios, respectively, defined in (3.14) obtained from the extraction error in (3.12) using the h -month linear projection GARCH(1,1) equation in (3.16). The forecast horizons are $h=1,6,12$ months. The monthly GARCH parameters for the S&P 500 were obtained from Table III (last panel), mainly for consistency purposes with the simulation evidence. The equivalent monthly sample was approximated by aggregating every 22 days to obtain a sample of 844 months. The monthly spot volatilities based on daily frequency for 1-, 2- and 3-months are $SV1$, $SV2$ and $SV3$, respectively. The spot volatilities based on monthly frequency are the textitRiskMetrics (RM) and the Rolling Volatilities for 12-, 24- and 60-months, RV_{12} , RV_{24} and RV_{60} , respectively.

Table XIV: Empirical Results of h -day ahead MSE and MAE Ratios of Daily Integrated

		Benchmark: $IV1$								
		$IV2$	$IV3$	$HIV1$	$HIV2$	$HIV3$	$EHIV1$	$EHIV2$	$EHIV3$	
		Five-minute Sampling Frequency DM/US\$								
Volatilities	$h = 1$ day									
	MSE	0.7338	0.6344	0.9015	0.7049	0.6106	0.9046	0.7218	0.6286	
	MAE	0.9180	0.8742	0.9602	0.9075	0.8684	0.9531	0.9342	0.9214	
	$h = 5$ days									
	MSE	0.7115	0.6058	0.8958	0.6828	0.5825	0.8975	0.7101	0.6152	
	MAE	0.9101	0.8605	0.9581	0.9010	0.8560	0.9483	0.9309	0.9208	
	$h = 20$ days									
	MSE	0.7409	0.6543	0.9173	0.7277	0.6455	0.9122	0.7900	0.7452	
	MAE	0.9342	0.9044	0.9768	0.9359	0.9102	0.9707	0.9675	0.9707	
			YN/US\$							
	$h = 1$ day									
	MSE	0.7575	0.6494	0.9065	0.7264	0.6330	0.9231	1.0556	0.6869	
	MAE	0.9299	0.8910	0.9706	0.9196	0.8865	0.9714	0.9813	0.9300	
	$h = 5$ days									
	MSE	0.7320	0.6163	0.8984	0.7005	0.6015	0.9157	1.0565	0.6712	
	MAE	0.9230	0.8807	0.9663	0.9139	0.8758	0.9680	0.9828	0.9337	
	$h = 20$ days									
	MSE	0.7758	0.6835	0.9249	0.7602	0.6805	0.9337	1.0489	0.7724	
MAE	0.9541	0.9361	0.9906	0.9614	0.9438	0.9859	1.0166	0.9882		
		Fifteen-minute sampling frequency DM/US\$								
$h = 1$ day										
MSE	0.7179	0.6136	0.8899	0.6847	0.5896	0.8802	0.6826	0.5781		
MAE	0.9217	0.8767	0.9612	0.9087	0.8713	0.9561	0.9153	0.8804		
$h = 5$ days										
MSE	0.6930	0.5814	0.8819	0.6592	0.5572	0.8708	0.6602	0.5485		
MAE	0.9103	0.8609	0.9559	0.8985	0.8552	0.9495	0.9071	0.8682		
$h = 20$ days										
MSE	0.7180	0.6224	0.8994	0.6967	0.6096	0.8869	0.7088	0.6211		
MAE	0.9300	0.8947	0.9749	0.9270	0.8960	0.9675	0.9328	0.9054		

Table XIV continued

		YN/US\$							
<i>h</i> = 1 day									
MSE	0.7533	0.6453	0.9033	0.7221	0.6297	0.8972	0.8196	0.6310	
MAE	0.9278	0.8931	0.9707	0.9197	0.8888	0.9688	0.9352	0.8941	
<i>h</i> = 5 days									
MSE	0.7273	0.6107	0.8954	0.6956	0.5964	0.8879	0.8012	0.6031	
MAE	0.9198	0.8788	0.9671	0.9117	0.8746	0.9643	0.9299	0.8859	
<i>h</i> = 20 days									
MSE	0.7631	0.6662	0.9169	0.7447	0.6622	0.9079	0.8285	0.6852	
MAE	0.9465	0.9241	0.9816	0.9492	0.9301	0.9784	0.9639	0.9429	
		Thirty-minute sampling frequency							
		DM/US\$							
<i>h</i> = 1 day									
MSE	0.6743	0.5675	0.8659	0.6443	0.5429	0.8579	0.6358	0.5283	
MAE	0.9177	0.8750	0.9567	0.9042	0.8679	0.9534	0.9036	0.8674	
<i>h</i> = 5 days									
MSE	0.6507	0.5380	0.8571	0.6203	0.5130	0.8483	0.6128	0.4991	
MAE	0.9068	0.8597	0.9516	0.8931	0.8514	0.9478	0.8939	0.8539	
<i>h</i> = 20 days									
MSE	0.6738	0.5748	0.8707	0.6521	0.5582	0.8616	0.6491	0.5229	
MAE	0.9273	0.8929	0.9694	0.9242	0.8947	0.9653	0.9239	0.8991	
		YN/US\$							
<i>h</i> = 1 day									
MSE	0.7093	0.5969	0.8686	0.6718	0.5767	0.8546	0.7198	0.5739	
MAE	0.9284	0.8940	0.9695	0.9200	0.8886	0.9681	0.9275	0.8875	
<i>h</i> = 5 days									
MSE	0.6837	0.5634	0.8553	0.6448	0.5733	0.8429	0.6964	0.5411	
MAE	0.9186	0.8778	0.9652	0.9091	0.8738	0.9631	0.9179	0.8738	
<i>h</i> = 20 days									
MSE	0.7143	0.6107	0.8756	0.6876	0.6003	0.8630	0.7315	0.6013	
MAE	0.9470	0.9251	0.9874	0.9483	0.9305	0.9858	0.9574	0.9331	

Note: The *h*-day ahead forecast MSE and MAE are the mean square error and mean absolute error ratios, respectively, defined in (3.15) and obtained from the extraction error (3.13) using the *h*-day linear projection GARCH(1,1) equation in (3.16). The forecast horizons of *h*=1,5,20 days represent the one day, week and month, respectively. The daily GARCH(1,1) parameters for each series were obtained in Table III (top panel) mainly for consistency purposes with the simulation results and for comparison with other studies. The DM/US \$ and YN/US \$ returns cover the ten year period, 01/12/86 - 20/11/96 (removing some days as described in section 5). The final daily sample of 2,449 days covers 705,312 5-minute observations. These were aggregated to obtain the 15- and 30-minutes samples of 235,104 and 117,552 observations, respectively. The data source is Olsen and Associates. The block-sample integrated volatilities for 1-, 2- and 3-days are the *IV*1, *IV*2 and *IV*3, respectively. The Historical Integrated Volatilities for 1-, 2- and 3-days are *HIV*1, *HIV*2 and *HIV*3, respectively. The Exponential HIV for 1-, 2- and 3-days are the *EHIV*1, *EHIV*2 and *EHIV*3, respectively. The decay factor in the exponential weights that yielded the minimum MSE was 0.990 (among the choice of 0.96 and 0.94).

Table XV: Empirical Results of h -month ahead MSE and MAE Ratios of S&P 500 Monthly

		Benchmark: 1-month IV ($IV22$)					
		$IV44$	$IV66$	$HIV22$	$HIV44$	$HIV66$	
Integrated Volatilities	$h = 1$ month						
	MSE	0.7255	0.6273	0.8028	0.6900	0.6142	
	MAE	0.9991	1.0005	1.0024	1.0044	1.0042	
	$h = 6$ months						
	MSE	0.6414	0.5118	0.7491	0.5994	0.5015	
	MAE	0.9177	0.8892	0.9546	0.9192	0.8943	
	$h = 12$ months						
	MSE	0.9029	0.8719	0.9336	0.8969	0.8752	
	MAE	0.9697	0.9520	0.9807	0.9630	0.9518	
			Benchmark: 1-month IV + SC ($IV22 + SC22$)				
			$IV44+$ $SC44$	$IV66+$ $SC66$	$HIV22+$ $SC22$	$HIV44+$ $SC44$	$HIV66+$ $SC66$
		$h = 1$ month					
	MSE	0.6404	0.5093	0.4964	0.5929	0.6063	
	MAE	0.9165	0.8884	0.8904	0.9142	0.2361	
	$h = 6$ months						
	MSE	0.6404	0.5093	0.4664	0.5929	0.6063	
	MAE	0.9165	0.8884	0.8904	0.9142	1.2361	
	$h = 12$ months						
	MSE	0.9015	0.8694	0.8717	0.8936	1.6015	
	MAE	0.9691	0.9511	0.9498	0.9612	1.3914	

Table XV continued
 Benchmark: 1-month CAR (CAR22)

	CAR44	CAR66	HCAR22	HCAR44	HCAR66
<i>h</i> = 1 month					
MSE	0.9329	0.9040	0.9620	0.9282	0.9034
MAE	0.9998	0.9995	0.9999	0.9996	0.9994
<i>h</i> = 6 months					
MSE	0.7827	0.6903	0.8891	0.7799	0.7016
MAE	0.9466	0.9181	0.9741	0.9489	0.9303
<i>h</i> = 12 months					
MSE	0.9868	0.9835	0.9951	0.9903	0.9881
MAE	0.9937	0.9917	0.9976	0.9946	0.9927

Note: The *h*-month ahead forecast MSE and MAE are the mean square error and mean absolute error ratios, respectively, defined in (3.14) obtained from the extraction error in (3.12) using the *h*-month linear projection GARCH(1,1) equation in (3.16). The forecast horizons are *h*=1,6,12 months. The monthly GARCH parameters for the S&P 500 were obtained from Table III (last panel), mainly for consistency purposes with the simulation evidence. The equivalent monthly sample was approximated by aggregating every 22 days to obtain a sample of 844 months. The monthly block sample integrated volatilities for 22-, 44- and 66-days are *IV*22, *IV*44 and *IV*66, respectively. The Historical Integrated Volatilities for 22-, 44- and 66-days are the *HIV*22, *HIV*44 and *HIV*66, respectively. The second panel presents a serial correlation correction to the integrated volatilities; The *IV*22+*SC*22, *IV*44+*SC*44 and *IV*66+*SC*66 are the 22-, 44- and 66-day block sampled Integrated Volatilities (IV) corrected for serial correlation by adding 2 times the sum of the cross-product terms of returns with the first lag of returns (SC). The *HIV*+*SC* refer to respective Historical Integrated Volatility plus the cross-product term. The last panel presents a different definition of cumulative quadratic variation. The CAR22, CAR44, CAR66 are the 22-, 44- and 66-day Cumulative Absolute Returns, respectively. HCAR22, HCAR44, HCAR66 are the 22-, 44- and 66-day Historical Cumulative Absolute Returns based on rolling samples of CAR.

Liste des publications au CIRANO *

Cahiers CIRANO / *CIRANO Papers* (ISSN 1198-8169)

- 99c-1 Les Expos, l'OSM, les universités, les hôpitaux : Le coût d'un déficit de 400 000 emplois au Québec — Expos, Montréal Symphony Orchestra, Universities, Hospitals: The Cost of a 400,000-Job Shortfall in Québec / Marcel Boyer
- 96c-1 Peut-on créer des emplois en réglementant le temps de travail? / Robert Lacroix
- 95c-2 Anomalies de marché et sélection des titres au Canada / Richard Guay, Jean-François L'Her et Jean-Marc Suret
- 95c-1 La réglementation incitative / Marcel Boyer
- 94c-3 L'importance relative des gouvernements : causes, conséquences et organisations alternative / Claude Montmarquette
- 94c-2 Commercial Bankruptcy and Financial Reorganization in Canada / Jocelyn Martel
- 94c-1 Faire ou faire faire : La perspective de l'économie des organisations / Michel Patry

Série Scientifique / *Scientific Series* (ISSN 1198-8177)

- 2000s-18 Confidence Regions for Calibrated Parameters in Computable General Equilibrium Models / Touhami Abdelkhalek et Jean-Marie Dufour
- 2000s-17 Markovian Processes, Two-Sided Autoregressions and Finite-Sample Inference for Stationary and Nonstationary Autoregressive Processes / Jean-Marie Dufour et Oliver Torrès
- 2000s-16 Exact Tests for Contemporaneous Correlation of Disturbances in Seemingly Unrelated Regressions / Jean-Marie Dufour et Lynda Khalaf
- 2000s-15 Simulation Based Finite and Large Sample Tests in Multivariate Regressions / Jean-Marie Dufour et Lynda Khalaf
- 2000s-14 Les obstacles à l'innovation dans les industries de services au Canada / Pierre Mohnen et Julio Rosa
- 2000s-13 Finite Sample Inference Methods for Simultaneous Equations and Models with Unobserved and Generated Regressors / Jean-Marie Dufour et Joanna Jasiak
- 2000s-12 Protection, Lobbying, and Market Structures / Arye Hillman, Ngo Van Long et Antoine Soubeyrand
- 2000s-11 The Asian Financial Crisis: The Role of Derivative Securities Trading and Foreign Investors / Eric Ghysels et Junghoon Seon
- 2000s-10 Productivity Trends and Employment across Industries in Canada / Pierre Mohnen et Thijs ten Raa
- 2000s-09 Nouvelle économie, nouvelle organisation et technologies de l'information / Suzanne Rivard

* Vous pouvez consulter la liste complète des publications du CIRANO et les publications elles-mêmes sur notre site Internet à l'adresse suivante :