DEPARTMENT OF ECONOMICS
UNIVERSITY OF CYPRUS

ROLLING-SAMPLE VOLATILITY ESTIMATORS: SOME NEW THEORETICAL, SIMULATION AND EMPIRICAL RESULTS

Elena Andreou and Eric Ghysels

Discussion Paper 2001-08

P.O. Box 20537, 1678 Nicosia, CYPRUS Tel.: ++357-2-892430, Fax: ++357-2-892432
Web site: http://www.econ.ucy.ac.cy

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

Elena Andreou
Department of Economics,
University of Cyprus, 1678 Nicosia, Cyprus (elena.andreou@ucy.ac.cy).

Eric Ghysels
Departments of Economics and Finance,
University of North Carolina,
Chapel Hill, NC 27599-3305 (eghysels@unc.edu).

September 19, 2001

Abstract

We propose extensions of the continuous record asymptotic analysis for rolling sample variance estimators developed by Foster and Nelson (1996) for estimating the quadratic variation of asset returns, which is also referred to as integrated or realized volatility. The new approach treats integrated volatility as a continuous time stochastic process sampled at high frequencies and suggests rolling sample estimators which share many features with spot volatility estimators. We also discuss asymptotically efficient window lengths and optimal weighting schemes for estimators of the quadratic variation and establish the links between various spot and integrated volatility estimators. Theoretical results are complemented with extensive Monte Carlo simulations and an empirical investigation.

Key Words: Rolling sample estimators, quadratic variation, volatility, efficient filtering, continuous record asymptotics, high-frequency data.

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 (2000) and Andersen, Bollerslev, Diebold and Labys (2001) suggest the sum of squared high frequency intra-day returns as an approximation to the daily volatility defined as the $\int_{t-H}^{t} \sigma_\tau^2 d\tau$, 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 $t+H$. The latter may also incorporate the jump component which in the present analysis we assume to be zero since we follow the setup of Foster and Nelson (1996). The quadratic variation is also referred to as integrated or realized volatility. 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 extend the continuous record asymptotic analysis of Foster and Nelson (1996) for the efficient estimation of the volatility process, Q_t^H, considered as a continuous time process in t for fixed H which is sampled discretely at high frequencies. The advantage of this approach is that the optimal window length and weight design for estimators of Q_t^H can be obtained directly from Foster and Nelson's analysis with the appropriate modifications. This leads to the introduction of a new class of volatility estimators based on the rolling estimation of the quadratic variation with an exponential weighting scheme. The same arguments are also applied to cumulative absolute return measures which also belong to this class of estimators but are found to be more robust to heavy tails and moving average effects. Barndorff-Nielsen and Shephard (2000) show that the noisy nature of realized volatility may benefit from smoothing techniques. Our paper shows explicitly such benefits.

The paper is organized as follows: The first section applies the continuous record asymptotics argument to volatility estimators based on the quadratic variation. The second section provides the details about the Monte Carlo design. The design is built on models which describe FX and equity markets data and takes advantage of the elegant GARCH temporal aggregation results of Drost and Nijman (1993) and Drost and Werker (1996). The simulation results are described in the third section. The fourth section reports empirical comparisons of various quadratic variation filters using different intra-day sampling schemes for FX series to examine the efficiency of h-period ahead forecasts for daily volatility as well as daily S&P 500 data to study the efficiency of forecasting monthly data-driven volatilities. A final section concludes the paper.

1 Continuous record asymptotics for quadratic variation and other volatility processes

The analysis in this section considers the quadratic variation, 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 the quadratic variation of returns. The daily (block sampling) scheme is a particular case of the generic context considered in this study. The advantages of this approach are twofold. First, we can derive the optimal window lengths for alternative sampling frequencies so as to maintain the MSE asymptotic efficiency of volatility filters for Q_t^H, or related volatility processes, across various sampling frequencies m. Second, we can apply both block and rolling sample estimation for these high-frequency volatility filters where for the latter we examine optimal weight designs (e.g. flat versus exponential decaying weights). This section is organized as follows. In a first subsection we define some notation and introduce various volatility processes. In the second subsection we extend the Foster and Nelson analysis to the quadratic variation process. In the final subsection we discuss the link between rolling sample estimators for conditional spot volatility and its quadratic variation.

1.1 Definitions and Various Volatility Processes

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. In particular, let:

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

where W_{1t} and W_{2t} are standard Brownian motions (possibly correlated) and the functions $\mu(\cdot,\cdot)$, $\zeta(\cdot,\cdot)$ are continuous and $\delta(\cdot,\cdot)$ is strictly positive. As 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 (e.g. Drost and Nijman 1993; Drost and Werker 1996).

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. A notation slightly different from that of Foster and Nelson is adopted that is 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. Note that the 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 financial 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.

Foster and Nelson focus on the estimation of σ_t in (1.1) which is usually called the spot volatility. In general, with ever finer sampling intervals, i.e. $m \to \infty$, 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: $Var_{(m),t}(r_{(m),t}) \equiv E[(\Delta_{(m)} M_{(m),t} - \mu_{(m),t})^2 | \mathcal{F}_{(m),t}] = \sigma^2_{(m),t} m^{-1}$ where $\sigma^2_{(m),t}$ measures the conditional variance per unit of time. Various data-driven estimators for instantaneous volatility, $\sigma^2_{(m),t}$, can generically be written as:

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

where $w_{(\tau-t)}$ is a weighting scheme and $\hat{\mu}_{(m),t}$ is a (rolling sample) estimate of the drift.

The analysis of Foster and Nelson can be extended to other notions of volatility, and hence to a variety of other processes related to volatility. Andersen and Bollerslev (1998) considered $\int_{t-H}^{t} \sigma^2_\tau d\tau$, i.e. an integral of instantaneous volatilities over some horizon H which we will refer to as Q_t^H since it relates to the quadratic variation of a process. They suggest the use of 5-minute squared returns on FX data to approximate the integral. French et al. (1987) used a similar measure involving the (monthly) sum of squared daily returns augmented by cross-products to account for correlation. Alternative measures similar to Q_t^H are the cumulative absolute returns. 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. Moreover, 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 Caroll, 1987). Hence, one may examine $\int_{t-H}^{t} |r_t|$ (ignoring the conditional mean correction) as a measure of integrated volatility, which is also called the cumulative absolute returns.

1.2 Asymptotic distribution of Quadratic Variation Filters

We consider the quadratic variation of a process: $Q_t^H = \int_{t-H}^t \sigma_\tau^2 d\tau$ which relates to a measure of volatility. 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 $\{\mathcal{F}_t^Q\}$. Barndorff-Nielsen and Shephard (2000) introduce a class of non-Gaussian Ornstein-Uhlenbeck Lévy-type processes and characterize analytically the quadratic variation.

Throughout this section the benchmark frequency is assumed to be daily. We are interested in discretizations $Q^H_{(m_1,m_2),t}$ 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. The fact that a double index is necessary reveals some outstanding issues in the asymptotic distribution theory. So far the focus has been on $Q^H_{(H,m_2),t}$ for fixed $H = 1$ and asymptotics for $m_2 \to \infty$. Barndorff-Nielsen and Shephard (2000) show that $(Q^1_{(1,m_2),t} - Q^1_t)/\sqrt{2/3 \sum_{j=0}^{m_2} r^4_{(m_2),t-j/m_2}}$ converges to a standard Normal random variable. The convergence rate is root-m_2 and the result shows that the estimators of Andersen and Bollerslev (1998) can be extremely noisy, particularly when volatility is high (since the asymptotic variance depends on the fourth power of returns). We will not tackle the double index asymptotics for $m_1 \neq m_2$ except that we can in principle apply the arguments below to m_2 fixed (as opposed to m_1, hence one takes a fixed finite sum of intra-day observations) and proceed with the asymptotics for $m_1 \to \infty$. Since this may be somewhat awkward we will first look at the $Q^H_{(m,m),t}$ discretization and then examine the Andersen and Bollerslev (1998) scheme $Q^H_{(1,m),t}$. The former also has a Doob-Meyer decomposition (see again Shiryaev, 1999, for details):

$$Q^H_{(m,m),t} = \mu^Q_{(m),t} m^{-1} + M^Q_{(m),t} - M^Q_{(m),t-1/m} \equiv \mu^Q_{(m),t} m^{-1} + \Delta_{(m)} M^Q_{(m),t}$$

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

$$\hat{Q}^H_{(m_1,m_2),t} = \sum_i w_{(m_1),i} \sum_j r^2_{(m_2),(t-i/m_1)-j/m_2} \equiv \sum_i w_{(m_1),i} QV^H_{(m_2),t-i/m_1} \quad (1.3)$$

where $w_{(m_1),i}$ is a weighting scheme involving discrete (rolling) sampling at frequency $1/m_1$, of the quadratic variation $QV^H_{(m_2),t}$ defined as:

$$QV^H_{(m_2),t} \equiv \textstyle\sum_{j=0}^{Hm_2} r^2_{(m_2),t-j/m_2} \tag{1.4}$$

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 (1.3) is quite 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. For flat weights of leads and lags (denoted by n_R and n_L, respectively) the scheme is characterized as $w_{(m_1),i} = m_1^{-1/2}(n_L+n_R)^{-1}I\{\tau\epsilon[-n_L m_1^{-1/2}, n_R m_1^{-1/2}]\}$. An extension of Theorem 2 (p.148) in Foster and Nelson yields that $m^{1/4}(\hat{Q}^H_{(m,m),t} - Q^H_t) \rightarrow N(0, C^Q_{(m),t})$ as $m \rightarrow \infty$ where the continuous record asymptotic variance for a flat weighting scheme equals:

$$C^{QF}_{(m),t} \equiv \frac{\theta^Q_{(m),t}}{n_R + n_L} + \sqrt{\theta^Q_{(m),t}\Lambda^Q_{(m),t}}\ \rho^Q_{(m),t}\frac{n_R - n_L}{n_R + n_L} + \Lambda^Q_{(m),t}\frac{n_R^3 + n_L^3}{3(n_R + n_L)^2} \tag{1.5}$$

and the superscript QF refers to the equal weighted flat weighting scheme (applicable to the quadratic variation). The formal proof of the result in (1.5) is omitted. It is an extension of Theorem 2 in Foster and Nelson applied to the quadratic variations. It is important to note that the regularity conditions Foster and Nelson impose on the spot volatility process need to be transplanted on the quadratic variation, something which is assumed to hold here. The asymptotic variance of the normalized extraction error $m^{1/4}(\hat{Q}^H_{(m,m),t} - Q^H_t)$ has three components: $\Lambda^Q_{(m),t}$ represents the variance of the quadratic variation, $\theta^Q_{(m),t}$ represents the conditional second moment of the quadratic variation, which is implicitly related to the conditional fourth moment of the process and $\rho^Q_{(m),t}$ measures the correlation between the empirical second moment and the conditional variance. It is also worth noting that the convergence rate is different from the asymptotic results discussed by Barndorff-Nielsen and Shephard (2000). It is worth recalling that in principle we should be dealing with a double index asymptotics, the number of weighted local sums of quadratic variations (m_1) and the number of intra-day components (m_2) potentially having different rates. Barnidoff-Nielsen and Shephard (2000) examine only one dimension of the most general asymptotics and we expand along the other dimension but so far no general asymptotic theory has been established.

1.3 The Relationship between Rolling Quadratic Variation and Spot Volatility Filters

Andersen and Bollerslev (1998) propose an estimator for the quadratic variation 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:

$$plim_{m \longrightarrow \infty} \left(\int_{t-H}^{t} \sigma_\tau^2 d\tau - \sum_{j=1,\ldots,m} r_{(m),t-j/m}^2 \right) = 0$$

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

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

In this subsection we 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 Quadratic Variation denoted by $HQV_{(m),t}^H$, namely:

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

There are two immediate relationships that can be established between HQV-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 integrated volatility of Andersen et al. (2001) denoted as the QV estimator. More interestingly, we can also express the HQV estimator as a QV-type estimator replacing $r_{(m),t-j}^2$ by $\tilde{\sigma}_{(m),t-j}^2$, i.e. replacing squared returns by estimated *instantaneous* volatilities:

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

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-day spot volatility components of the daily quadratic variation and the weights of the HQV estimator, namely $w_{(m_1),i}$ in (1.3). It also follows that the weights for the spot volatility estimators are not particularly optimal weights. Nevertheless, one expects that the HQV estimator will be more efficient since one replaces a noisy $r^2_{(m),t-j}$ by (albeit sub-optimal) $\tilde{\sigma}^2_{(m),t-j}$. The analysis in Andersen and Bollerslev (1998) indeed suggests that replacing $r^2_{(m),t}$ by volatility estimates yields a more efficient scheme.

There is one particular case where the weights of the historical quadratic variation and 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 $\tilde{\sigma}^2_{(m),t-j}$ 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. Note that the HQV estimator involves two days worth of data because the last component $\tilde{\sigma}^2_{(m),t-m}$ involves an entire extra day of data. This will be taken into account when we compare simulated data-driven filters. This particular case is worth pursuing further since it allows us to say more about the comparison of the estimator $QV^H_{(m),t}$ in Andersen and Bollerslev and the HQV estimator proposed here. 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 $HQV^H_{(m),t}$ has $n_L m^{-1/2} = m$ lags. Using (1.5) one can calculate the relative efficiency as the ratio of the asymptotic mean squared error of the HQV estimator relative to that of the QV estimator. This ratio of MSEs is: $[3\theta^Q m^3 - 3\sqrt{\theta^Q \Lambda^Q}\rho^Q m^{7/2} + \Lambda^Q m^4]/[3\theta^Q - 3\sqrt{\theta^Q \Lambda^Q}\rho^Q m^{3/2} + \Lambda^Q m^3]$. This implies that generic efficiency gains and comparisons require the specification of θ^Q, Λ^Q and ρ^Q.

Specification of the optimal weights for extracting integrated volatility when Q^H_t is viewed as a continuous time process leads us to the results of Foster and Nelson. It circumvents the problem of formulating optimal weighting schemes for $\tilde{\sigma}^2_{(m),t-j}$ for all $j = 1, \ldots, Hm$, which becomes even more difficult due to the intra-day cross-correlation between spot volatilities. Approaching the design of optimal filter weights directly via (1.3) 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 (p.155) of Foster and Nelson which covers dominating flat weights (as noted before, using the design of formula (17) (p.155) from Foster and Nelson). The resulting $C^{QD}_{(m),t}$ obviously depends again on $\rho^Q_{(m),t}$, $\theta^Q_{(m),t}$ and $\Lambda^Q_{(m),t}$. Last, but certainly not least, we can consider the optimal exponentially declining weighting scheme considered in

Theorem 5 (p.154) of Foster and Nelson.

For comparison purposes we also investigate other processes in continuous time, such as: $CAR_t^H = \int_{t-H}^{t} |r_\tau| d\tau$ which is the cumulative absolute return (CAR) process that relates to volatility estimators used by Taylor (1986), Hsieh (1991) and Schwert (1989). The Cumulative Absolute Returns $CAR^H_{(m_2),t}$ estimator is:

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

which is based on sampling at frequency $1/m_2$. Similar to $Q^H_{(m,m),t}$ we are particularly interested in $CAR^H_{(m,m),t}$ involving weights $w_{(m),j} = 1/n_L$, which we will call historical cumulative absolute return, $HCAR^H_{(m),t}$, namely:

$$HCAR^H_{(m,m),t} \equiv HCAR^H_{(m),t} = \frac{1}{n_L} \sum_{j=0}^{n_L-1} CAR^H_{(m),t-j/m}. \tag{1.6}$$

This type of estimators are of particular interest when the process features conditional heavy tail behavior. In light of this and the work of Bai, Russell and Tiao (1999), we will also study the properties of these filters in the simulations and empirical analyses. The novelty of the estimators in (1.6) is the use of sliding spans of cumulative absolute returns over fine sampling intervals, similar to the historical quadratic variation estimators. The filters of Taylor (1986), Hsieh (1991) and Schwert (1989) can again be viewed as special cases, similar to QV being a special case of HQV-type filters.

2 Monte Carlo design

The objective of the Monte Carlo study is 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 scope of the Monte Carlo study is also to examine the optimal filter design of quadratic variation estimators. We organize the section in subsections. The first subsection describes the processes we simulate. A second subsection briefly introduces the various filters which will be considered, with details appearing in the Appendix. A third subsection presents the diagnostics.

2.1 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 the YN/US$ exchange rate and S&P 500 stock index. The following continuous time stochastic volatility model is considered which is based on the results of Drost and Nijman (1993) and Drost and Werker (1996):

$$\begin{aligned} dlnY_t &= \sigma_t dW_{pt} \\ d\sigma_t^2 &= \theta(\omega - \sigma_t)dt + (2\lambda\theta)^{1/2} dW_{\sigma t}. \end{aligned} \tag{2.1}$$

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

$$\begin{aligned} \ln Y_t - \ln Y_{t-1/m} &= r_{(m),t} = \sigma_{(m)} z_{(m),t} \\ \sigma^2_{(m),t} &= \phi_{(m)} + \alpha_{(m)} r^2_{(m),t-1/m} + \beta_{(m)} \sigma^2_{(m),t-1/m} \end{aligned} \tag{2.2}$$

where $z_{(m),t}$ is Normal i.i.d. $(0, 1)$ and $r_{(m),t}$ is the returns process sampled at frequency $1/m$. The diffusion parameters of (2.1) and the GARCH parameters of (2.2) 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 FX series we take the results of Andersen and Bollerslev (1998, Table 1, p.886), mainly for comparison purposes. The YN/US$ parameters are $\theta = 0.054$, $\omega = 0.476$ and $\lambda = 0.480$ for the daily sample period 01/10/87 to 30/09/92 (and the corresponding daily GARCH parameters are reported in Table I). 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, 5-minute and 1-minute frequencies are based on $m = 1, 288$ 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 sampling frequencies encountered in practice. The computations of the GARCH parameters for alternative m are reported in Table I and were obtained using the software available from Drost and Nijman (1993). Following the same paradigm we consider an analogous example for the equity market with only 6.5 hours of trading as opposed to the 24 hours trading in FX markets. Therefore a GARCH(1,1) model is estimated for daily S&P 500 returns and using the above disaggregation results we consider the equivalent intra-day frequencies for $m = 1, 78$ and 390. The results in Table I show the daily Normal-GARCH(1,1) estimated parameters for the S&P 500 for the period 02/04/86-29/08/97 ($T = 2884$). In 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) to obtain the monthly GARCH parameters in Table I. It is interesting to note that the reported kurtosis coefficient $\kappa(m)$ for each asset at different sampling frequencies, m, remains approximately constant which is evidence of the constancy of higher-order moments required by a modification of the regularity conditions of Foster and Nelson for comparing filters with different sampling frequencies (see Andreou and Ghysels (1999) and Andreou, Chernov and Ghysels (2001) for further details).

2.2 Data-Driven Volatility Filters

We examine the properties of volatility estimators using the simulation design described above. The details of the three classes of volatility filters based on the intra-day quadratic variation of returns appear in the Appendix. Note that the main simulation analysis does not assume conditional mean effects for the return process. Nevertheless, the study addresses the MA effects in intra-day returns for the simulated filters.

Following the analysis of various types of data-driven quadratic variation measures discussed in Section 1 we will first consider the *one-day Quadratic Variation* $\hat{\sigma}_t^{QV1}$ (defined by Andersen and Bollerslev, 1998) as the sum of squared returns $r_{(m),t}$ for different values of m, to produce the daily volatility measure (see also (A.1) in the Appendix) where for the 5-minute sampling frequency the lag length takes 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 QVk with k equal to one. Andersen et al. (2001) 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 log of monthly QV estimates the skewness effects. We do not consider this transformation as we are primarily concerned with the MSE of extraction filters. Quadratic variation filters with an extended window length of 2 and 3 days denoted by $QV2$ and $QV3$ respectively, are also considered. Next we will consider the *one-day Historical Quadratic Variation*, which was discussed at length in Section 1, namely (for $H = 1$ and $n_L = m$ as appearing in (A.3)). In the remainder of the paper we will refer to this filter as $HQVk$ with k equal to one. The $HQV2$ and $HQV3$ filters will also be considered. In order to appraise the efficiency gains of using optimal weighting schemes to extract Q_t^1 we consider exponential declining weights in (A.5) where we select the decay rate a, equal to a range of values (0.940, 0.960, 0.999) for daily filters. These filters will be denoted $EHQVk$ with k equal to $1, 2$ and 3 days. It must be noted that due to the rolling estimation scheme of the HQV type filters an extra day of data is used in a block-sampling context. By analogy the monthly integrated volatilities and historical quadratic variation filters of $k = 1, 2$ and 3 months are based on an approximation of $22, 44$ and 66 trading days, respectively. A selection of CAR filters is also considered. The *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, produces the daily volatility measure defined in (A.2). This filter will be called $CAR1$. The $CAR2$ and $CAR3$ filters will also be considered and are defined in analogous manner to QVk for $k = 2, 3$ days. In addition, the *one-day Historical Cumulative Absolute Return*, discussed in Section 1, and defined in (A.4) is also extended to 2 and 3 trading days, similar to $HQVk$. Similarly the monthly CAR and historically CAR filters of $k = 1, 2$ and 3 months are based on $22, 44$ and 66 days, respectively.

2.3 Statistical Criteria of Appraisal

For the daily benchmark frequency case 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 respectively, in Table I. This is considered to be the highest intra-day sampling frequency which represents the true generating process. The sample sizes for the Monte Carlo experiments are summarized in the Appendix. Next, we apply the GARCH dynamics to obtain the quadratic variation filters. Since we take $H = 1$ we will denote the quadratic variation by Q_t^1.

Then we can write the extraction error as:

$$\varepsilon_t^i = Q_t^1 - Q_{(m,m),t}^i \tag{2.3}$$

for $i = QVk, HQVk$ and $EHQVk$, $CARk$, $HCARk$ and $k = 1, 2, 3$. We take some liberty in (2.3) with regards to notation by using $Q_{(m,m),t}^i$ to facilitate the definition of the extraction process. For the 'true' daily quadratic variation 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 other intra-day sampling frequencies such as 30-minutes which gave rise to similar results and hence for conciseness are not reported). The same analysis is followed in the monthly Monte Carlo design where we simulate the daily GARCH models and aggregate the daily GARCH variance to obtain the 'true' monthly quadratic variation. The monthly volatility filters will be denoted by $(H)QVk$, $k = 22, 44, 66$ days. The simulated samples are summarized in the Appendix. The behavior of the extraction error, ε_t^i, is examined according to the following four dimensions: (a) The efficiency of filters is examined using the Mean Square Error (MSE) (the criterion relevant for the asymptotic results in Section 1) which we compare across different filters and sample sizes. Note that we also obtain the Mean Absolute Error (MAE) which is more robust than the RMSE that is argued to be susceptible to outliers (Andersen, Bollerslev and Lange, 1999). The relative efficiency of one filter vis-à-vis another is studied by computing ratios of MSEs (and MAEs) using the corresponding MSE (and MAE) of the one-day quadratic variation as the benchmark:

$$MSE^i/MSE^{QV1} \qquad or \qquad MAE^i/MAE^{QV1} \tag{2.4}$$

where i refers to the MSEs obtained from $Q_{(m_1,m_2),t}^i$ for $i = QV2$, $QV3$, $HQVk$ and $EHQVk$, $k = 1$ to 3. This analysis extends to the monthly frequency where the benchmark is the one-month quadratic variation, denoted as $QV22$. (b) The out-of-sample forecast performance is studied in an analogous manner to Andersen et al. (1999) for the family of the above quadratic variation filters for which we compare their relative efficiency. 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:

$$\begin{aligned} P_{(m),t}(r^2_{(1/h),t+h}) &= m \cdot h \cdot \sigma^2_{(m)} + (\alpha_{(m)} + \beta_{(m)}) \cdot [1 - (\alpha_{(m)} + \beta_{(m)})^{m \cdot h}] \\ &\cdot [1 - \alpha_{(m)} - \beta_{(m)}]^{-1} \cdot (\sigma^2_{(m),t} - \sigma^2_{(m)}) \end{aligned} \quad (2.5)$$

where $\sigma^2_{(m)} \equiv \phi_{(m)} \cdot (1 - \alpha_{(m)} - \beta_{(m)})^{-1}$ and $\sigma^2_{(m),t}$ would be the alternative quadratic variation filters discussed in the previous section. We shall consider $h = 20$ days as in Andersen et al. (1999) and $h = 1$ day and obtain the MSE and MAE for each $h = 1$ and 20 days out-of-sample volatility filter forecast. For the monthly experiment $h = 12$ months. It is interesting to note that if we ignore parameter uncertainty, which is the case for our Monte Carlo simulations, we can view (2.5) as a functional transformation of $\sigma^2_{(m),t}$, 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 MSEs of forecasts in empirical applications, whereas MSEs 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:

$$Q^i_{(m_1,m_2),t} = a + b.Q^1_t + u_t, \quad (2.6)$$

where $i = QVk$, $HQVk$ and $EHQVk$. (d) Finally the inefficiency of filters is examined using 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. We also examine the MSE of the autocorrelation functions between the theoretical GARCH(1,1) and the autocorrelation coefficients of the alternative data-driven volatilities.

3 Monte Carlo results

The Monte Carlo simulation results are analyzed in view of the theoretical results in Section 1 for the quadratic variation of returns. We examine whether for this 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. In addition, we evaluate the optimality of alternative weighting schemes for data-driven volatilities using the above statistical criteria. The simulation results presented refer to the five year daily sample. Similar results apply to the ten year sample period.

3.1 Data window length, estimation method and weighting schemes

We focus first on the comparison of $QV2$ and $QV3$ (relative to $QV1$), and hence examine the effect of the window length for QV type estimators. By analogy, we concentrate next on the comparison of $(E)HQVk$ for different values of k and alternative weighting schemes.

The simulation results in panel (A) of Table II show the contemporaneous MSE (and MAE) ratios defined in (2.4) obtained from the extraction error defined in (2.3). We find that among the block sample intra-day quadratic variation filters $QV3$ is the most efficient filter in the equity representative series. In fact, for the S&P500 $QV3$ is 57% more efficient on MSE grounds than $QV1$. A window of 3 days is also employed for the $(E)HQV2$ filters which is equally efficient for both flat and exponential weights in the S&P500. Recall that $HQVk$ and $EHQVk$ estimators should be compared with QVn where $n = k + 1$, since the rolling scheme entails one extra day of data in a block-sampling context. In the FX representative series the YN/US$ contemporaneous MSE results show that in general a shorter window of 2 days provides the most efficient quadratic variation filters as shown by $QV2$ and $(E)HQV1$. $QV2$ appears 68% more efficient than $QV1$.

The comparison of $HQVk$ and $EHQVk$ enables us to appraise to what extent exponential weighting of QV-type estimators translates into efficiency gains. The optimal decay rate of 0.999 yields the lowest MSE among an a priori choice of 0.940, 0.960 and 0.999. The MSE results in Table II show that there are marginal gains (compare $HQV1$ and $EHQV1$ for YN/US$ and $HQV2$ and $EHQV2$ for S&P500) to be made by changing the weights to exponentially declining. Nevertheless, it is worth noting the efficiency gains of an exponentially weighted scheme are only achieved conditional on the identification of the optimal window length (as opposed to adopting exponential weights with inefficient windows that yield higher simulated MSEs).

A further interesting aspect of the quadratic variation 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 1 that we cannot a priori predict the efficiency of QV and HQV type volatility estimators unless we know something about the behavior of higher moments. The daily integrated volatility simulation results in Table II show that the lowest MSEs for the S&P 500 are achieved by $QV3$, $HQV2$ and $EHQV2$ estimators. Between the first two estimators $QV3$ is more efficient than $HQV2$. Similarly, in the YN case the $QV2$ is the relatively most efficient estimator among. Hence the S&P 500 and YN/US$ simulation results suggest than block sample estimation is marginally more efficient that rolling sample estimation of the quadratic variation of returns. The above results extend to the daily MSE (MAE) for short-run volatility forecasting ($h = 1$) as shown in Panel (B) of Table II and the long-run ($h = 20$) days with S&P 500 volatility forecasts. Similar results apply to the daily DM/US$ case found in Andreou and Ghysels (1999).

We further examine these estimation methods for the monthly integrated volatilities (in Panel (A) of Table II). In all simulated cases we find that the 22-day block-sample estimator for monthly volatility ($QV22$) is the relatively least efficient estimator shown by all the MSE ratios being less than unity. On the other hand, the 22-day historical quadratic variation (HQV) appears to be more efficient than $QV22$. A more direct comparison is between HQV and QV filters both of which employ the same window length of 44 days, given by $QV44$ and $HQV22$. 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 $QV44/HQV22$ is 0.79 for the S&P 500 and 0.91 for the YN. Hence, on average the FX simulation results show that a 44-day window length as opposed to a 22-day one is optimal for the estimation of monthly integrated volatilities as shown by the MSEs of $QV44$ and $HQV22$. Note, however, that the S&P 500 results show that a longer window is required equal to 66 days as shown by the relatively most efficient estimates being $QV66$ and $HQV66$ on MSE grounds. Nevertheless, these results do not extend to long-run forecasting of 12-months where we observe in the last panel of Table II that all MSE ratios are close to unity.

These volatility filters are further evaluated using the coefficient of multiple determination in regression (2.6), the cross-covariances between extraction errors and filtered volatilities and the MSE of the ACFs between the GARCH and data-driven volatilities. In this section we briefly discuss the main results. The Mincer and Zarnowitz (1969) simulation results extend the Andersen and Bollerslev (1998) empirical finding of a relatively high R^2 for $QV1$. The simulation results show that the daily filters with the highest R^2 are the ones found to be the most efficient on MSEs grounds. In particular, the highest R^2 for the S&P 500 is due to $QV3$ ($R^2 = 0.938$) and for the YN to $EHQV1$ ($R^2 = 0.987$) and $QV1$ ($R^2 = 0.981$). A similar picture is sketched by the monthly QVs which show that the highest $R^{2\prime}s$ correspond to the MSE most efficient volatilities. For S&P 500 this is $QV66$ ($R^2 = 0.965$) and $HQV44$ ($R^2 = 0.974$) and for YN/US\$ $QV44$ ($R^2 = 0.911$) and $HQV22$ ($R^2 = 0.977$).

Examination of the correlation coefficients of each estimate of the quadratic variation with the extraction error, for lags 1-5 reveals that these correlations have a relatively small size. This observation applies to all volatilities and to both frequencies (daily and monthly) as well as to the two simulated cases. The correlation coefficients of all the QVs with the extraction error, range from 1% to 26%. The daily $QV1$ has the lowest correlation coefficients for all three series, followed by $HQV1$. Similarly, in the monthly case the lowest correlation coefficients apply to $QV22$ followed by $HQV44$. This result also holds for both the equity and FX markets. It is interesting to note is that although the historical quadratic variations 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 $QV1$ has the lowest MSE, whereas in the class or the rolling-sample $HQV3$ and $EHQV1$ exhibit the lowest MSE. These simulation results hold for the S&P 500 but not for the FX series for which all MSEs of ACFs are on average equal.

3.2 Tail behavior and MA components

The Monte Carlo analysis is extended to a simulated MA(1)-GARCH(1,1) process for a range of MA(1) coefficients $(0.2, \pm 0.5, 0.85)$ for S&P 500. We find that for a relatively high (0.85) and negative but moderate (-0.5) moving average coefficient, the MSE ratios of all quadratic variation filters rise and are almost equivalent as shown in Table III. In contrast, for smaller moving average coefficients the results stay broadly the same with the initial analysis reported above.

The theoretical analysis can also allow for the presence of excess kurtosis given that the fourth conditional moment remains constant for alternative frequencies. Based on the t-GARCH for alternative degrees of freedom (and the GARCH parameters reported in Table I) we first report, in Table III, the contemporaneous MSE (and MAE) ratios of daily volatility filters based on a 5-minute sampling frequency for the S&P 500. The following observations arise: First, the rolling-sample estimation method of QVs is more efficient for daily volatility filters (e.g. $HQV3$ is more efficient than $QV3$). Similarly, for the monthly frequency the $HQV66$ is the most efficient filter. It is interesting to point that the efficiency of $HQV66$ is robust to these conditional distributional assumptions. Note, however, that the MSE ratio of $HQV66$ falls significantly in the presence of excess kurtosis. Second, for the monthly frequency we find that block-sample QVs of alternative window lengths do not present any efficiency gains. This result is consistent under both conditional Normality and t since all QVk MSE (and MAE) ratios are close to one.

The lower panel of Table III presents the Cumulative Absolute Returns MSEs (and MAEs) for the monthly S&P 500 under the assumption of conditional Normality and Student's t with various degrees of freedom. 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 found to be relatively more efficient than the HQVs, a result that applies for alternative degrees of excess kurtosis. Moreover, comparing the MSE ratios of CAR type volatilities under both conditional Normality and Student's t we find 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) and Davidian and Carroll (1987) and extend the conjecture presented in the conclusions of Foster and Nelson (1996).

4 Empirical illustration

The efficiency of the above data-driven volatility filters is evaluated based on their out-of-sample forecasting performance for equity and FX series. The empirical results complement the simulation analysis by considering a number of horizons as well as 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 (2.5) where $\sigma^2_{(m),t}$ would denote the quadratic variation filters. For daily volatility we consider $h = 1, 5, 20$ days and obtain the MSE and 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 for FX and equity market series. 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. The original sample is 1,052,064 five-minute return observations (2,653 days · 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. A description of how these days are removed is found in Andersen et al. (2001). The final sample includes 705,024 five-minute returns reflecting 2,448 trading days. Using this FX dataset we calculate intra-day volatilities for three frequencies (5-, 15- and 30-minutes) 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 with a sample of 18,571 daily observations. The empirical results are examined in conjuction with the simulation analysis and present further evidence of the asymptotic predictions in Section 1.

We first examine the empirical MSE and MAE efficiency of intra-day volatility filters obtained from alternative h period ahead forecasts for the DM/US$ and YN/US$ and different sampling frequencies, m. These results are reported in Table IV. Recall that in the Monte Carlo analysis for the 5-minute sampling frequency, the MSE and MAE ratios for the 20-day ahead forecast of FX series are 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 measures of quadratic variation have a lower MSE than $QV1$. In particular, the 5-minute sample simulations showed that $QV3$ and $(E)HQV3$ were the relatively most efficient filters for the 20-day ahead forecast. Turning now to the empirical results in Table IV we are able to make a direct comparison and obtain additional evidence.

With respect to the empirical efficiency of h-step ahead volatility forecasts there are some interesting results. Both the DM/US$ and YN/US$ empirical MSE (and MAE) ratios are less than one, suggesting that $QV1$

is the relatively least efficient filter among this family of volatility filters. Therefore, a window length beyond that of one day improves the empirical MSE volatility efficiency for both exchange rates, for short and long forecast horizons ($h = 1, 5, 20$) and for alternative intra-day sampling frequencies (from 5 to 30 minutes). This result is consistent with the simulation results of the extraction (and less so with the forecast) error MSEs. It is worth emphasizing that in general the empirical efficiency measures for different filters appear robust across h, m, and the two FX series. The MSEs (and MAEs) tend to fall marginally as the sampling frequency increases from 5 to 30 minutes. Using the empirical MSEs we conclude that the 3-day data window length yields the relatively most efficient intra-day volatility filters for daily volatility forecasting. Compare, for instance the MSEs of $QV3$ and $HQV3$ across different FX, m and h. Given the optimal window length of three days, we compare the MSEs of the block-sample filter, $QV3$, with the rolling-sample filter, $HQV3$. Both estimation methods suggest almost equivalent MSE efficiency levels, with $HQV3$ being only marginally more efficient than $QV3$. Extending this observation to shorter window lengths and sampling frequencies we conclude that both block and rolling-sample estimation methods for $k > 1$ are empirically optimal for $m = 5, 15, 30$. In addition, on average the rolling-sample is only marginally more efficient. Similarly the Monte Carlo simulations for the FX 5-minute sampling frequency both QVs and $HQVs$ are asymptotically efficient. A related aspect of the rolling estimation method refers to the choice of the weighting scheme. The exponential weighting scheme for a 0.999 decay factor yields the lowest MSEs (a result also consistent and guided from the Monte Carlo analysis). Comparing the exponential and flat $HQVs$ we find that there are significant MSE gains in adopting a *combination* of optimal window length and exponential weighting schemes. The efficiency of $EHQV3$ is valid across FX series and forecast horizons.

Following the above analysis we also examine the empirical results for monthly S&P 500 volatility filters based on daily frequency, reported in Table V (top panel). Although the MSE ratios suggest that a window length of 3 months (approximately 66 days) is more optimal for both QVs and $HQVs$, the MAE ratios show that there are no significant gains in efficiency. The latter result also verifies the simulations of MSEs (and MAEs) that for the long-run 12-month ahead forecasting horizon there are no efficiency gains among volatility filters. Given the overall picture, we conclude that the choice of a 66-day window length seems to matter for the short- and medium-run forecasting ($h = 1, 6$) of the monthly quadratic variation based on daily data. In an attempt to compare and consolidate the empirical results we extend the empirical analysis to other types quadratic variation filters, 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 QV which is the sum of the cross product of returns at t and $t + 1$. Following their methodology we compute this estimator (denoted, $QV22$ + 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 V show 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. The latter result is consistent with that obtained for volatility filters based only on the quadratic variation, shown in the top panel of Table V. 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 ($HQV22$+SC22). It appears twice as efficient as the respective block sample estimator ($QV22$+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 MSE terms. 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. Given the range of MSEs being close to unity, we find that at the monthly frequency there are no significant gains from alternative data windows for estimating monthly cumulative quadratic variations of daily returns. The picture of monthly empirical results shows that for the short- and medium-run ($h = 1, 6$) forecasting of monthly volatilities, a filter of the quadratic variation based on 22 and 44 days with block and rolling sample estimation, respectively, with (or without) augmentation by the serial correlation moment estimator, has the lowest MSE forecast. Summarizing, the empirical results provide support for the theoretical analysis in section

1.

5 Conclusion

The paper applies the continuous record asymptotics arguments in Foster and Nelson (1996) 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 quadratic variation filters. The efficiency of block- and rolling-sample estimation methods and alternative weighting schemes for high frequency data-driven volatility estimators is examined. The theoretical continuous record asymptotic results are examined by a simulation study and complemented by an empirical investigation. We present empirical and simulation evidence that for a combination of optimal window length and (a priori) exponential weights, there are efficiency gains among the flat-weighted intra-day volatilities. Similarly, for monthly volatilities augmenting the monthly historical quadratic variation (of 22 days) by the serial correlation estimate improves the MSE efficiency.

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 3 days for the daily quadratic variation filters and 2 months (equivalently 44 days) for monthly QVs 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 exponential weighting schemes appear to be relatively more optimal than triangular weights for the daily volatilities when combined with the optimal window. In general, rolling quadratic variation filters are found to be marginally more efficient. Third, the quadratic variation filter attains optimality for certain of the simulated cases, once the estimation window is extended to 3-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. In the presence of excess kurtosis the relative efficiency of historical quadratic variations appears relatively more robust than that of block sampling QVs. It is interesting that a comparative simulation analysis suggests that the efficiency of CAR-type estimators is more robust to alternative elliptical distributional assumptions (such as conditional Normality and Student's t). Moreover, the HQV and $HCAR$ filters are introduced which aim to improve the efficiency of high-frequency volatility filters where the latter may be an improvement in the presence of excess kurtosis.

The broad conclusion is that the window length, data frequency, weighting scheme and estimation methods of volatility filters play an 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 also found sensitive to filter designs but to a far lesser extend than high frequency data filters.

ACKNOWLEDGMENTS

The authors would like to thank the hospitality and financial support of CIRANO, Montreal and the University of Manchester. We would like to thank Tim Bollerslev, Denise Osborn and Neil Shephard for comments on an earlier draft of this paper as well as an Associate Editor, Referee and the Editor, Jeff Wooldridge. We would also like to thank Michel Dacorogna of Olsen and Associates for providing us the data.

APPENDIX: DETAILS OF MONTE CARLO DESIGN

In this Appendix we provide details about filters and sample sizes used in the Monte Carlo design.

A Filters

Following the analysis of various types of data-driven quadratic variation measures discussed in Section 1 we will consider the following selection of filters: The *one-day Quadratic Variation* $\hat{\sigma}_t^{QV1}$ is defined as the sum of squared returns $r_{(m),t}$ for different values of m, to produce the daily volatility measure:

$$Q^1_{(1,m),t} = \sum_{j=1}^{m} r^2_{(m),t+1-j/m} \qquad t = 1, ..., n_{days}. \tag{A.1}$$

where for the 5-minute sampling frequency the lag length takes 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 QVk with k equal to one. Similarly, the *one-day Cumulative Absolute Return* is based on the sum of absolute returns $|r_{(m),t}|$ for different values of m:

$$CAR^1_{(m),t} = \sum_{j=1}^{m} |r_{(m),t+1-j/m}| \qquad t = 1, ..., n_{days}. \tag{A.2}$$

The *one-day Historical Quadratic Variation*, is specified as:

$$HQV^1_{(m),t} = \frac{1}{m} \sum_{j=1}^{m} QV1_{(m),t+1-j/m}. \qquad t = 1, ..., n_{days}. \tag{A.3}$$

Similarly the *one-day Historical Cumulative Absolute Return*, is defined as:

$$HCAR^1_{(m),t} = \frac{1}{m} \sum_{j=1}^{m} CAR^1_{(m),t+1-j/m}. \qquad t = 1, ..., n_{days}. \tag{A.4}$$

The *Exponentially weighted Historical Integrated Volatility*, which involves exponential declining weights, is:

$$EHQV^1_{(m),t} = A^{-1} \sum_{j=1}^{km} a^{-j} QV^1_{(m),t-1-j/m}. \qquad t = 1, ..., n_{days}. \tag{A.5}$$

where we select the decay rate a, equal to a range of values (0.940, 0.960, 0.999) for daily filters. Finally, A is a scaling constant to guarantee that the filter weights sum to one. These filters will be denoted $EHQVk$ with k equal to $1, 2$ and 3 days.

B Sample sizes

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

		24hrs FX Market			6.5hrs Equity Market		
n_{years}	n_{days}	$n_{30\,\text{min.}}$	$n_{5\,\text{min.}}$	$n_{1\,\text{min.}}$	$n_{30\,\text{min.}}$	$n_{5\,\text{min.}}$	$n_{1\,\text{min.}}$
5	1,250	60,000	360,000	1,800,000	16,250	97,000	487,500
10	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. With $n_{years} = 30$, we have $n_{months} = 360$ and $n_{days} = 4,320$. Each experiment is performed with 500 replications. Note that for the one-sided rolling estimates we create sufficient data (equivalent to one year) before the effective sample.

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 (2001), "The Distribution of Exchange Rate Volatility", *Journal of American Statistical Association*, 96, 42-55.

[3] Andersen, T., T. Bollerslev, S. Lange (1999), "Forecasting Financial Market Volatility: Sample Frequency vis-a-vis Forecast Horizon", *Journal of Empirical Finance*, 6, 5, 457-477.

[4] Andreou, E., M. Chernov and E. Ghysels (2001) "Variations on the Theme of Spot Volatility Efficiency", Discussion Paper, UNC Chapel Hill.

[5] Andreou, E. and E. Ghysels (1999), "Rolling Sample Volatility Estimators: Some New Theoretical, Simulation and Empirical Results", CIRANO Working Paper 2000-s19 `http://www.unc.edu/ eghysels/papers/eaegr.pdf`

[6] 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, Graduate School of Business, University of Chicago.

[7] Baillie, R. T. and T. Bollerslev (1992), "Prediction in dynamic models with time dependent conditional variances" *Journal of Econometrics*, 52, 91-113.

[8] Barndorff-Nielsen, O. E. and N. Shephard (2000), "Econometric analysis of realised volatility and its use in estimating stochastic volatility models", Discussion Paper, Nuffield College, University of Oxford.

[9] 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.

[10] Davidian M. and R.J. Carroll (1987), "Variance Function Estimation", *Journal of the American Statistical Association*, 82, 1079-1091.

[11] Drost, F. C. and T. Nijman (1993), "Temporal Aggregation of GARCH Processes", *Econometrica*, 61, 909-727.

[12] Drost, F. C. and B. M.J. Werker (1996), "Closing the GARCH Gap: Continuous Time GARCH Modeling", *Journal of Econometrics*, 74, 31-57.

[13] Foster, D. and D. Nelson (1996), "Continuous Record Asymptotics for Rolling Sample Estimators", *Econometrica*, 64, 139-174.

[14] French, K.R., G.W. Schwert and R.F. Stambaugh (1987), "Expected Stock Returns and Volatility", *Journal of Financial Economics*, 19, 3-29.

[15] Gallant, A.R., P. Rossi and G. Tauchen (1992), "Stock Prices and Volume", *Review of Financial Studies*, 5, 199-242.

[16] Hsieh, D.A. (1991), "Chaos and Nonlinear Dynamics: Application to Financial Markets", *Journal of Finance*, 46, 1839-1877.

[17] Mincer J. and V. Zarnowitz (1969) "The Evaluation of Economic Forecasts and Expectations", in Economic Forecasts and Expectations (J. Mincer, ed.) New York: National Bureau of Economic Research.

[18] Merton, R. C. (1980), "On Estimating the Expected Return on the Market: An Exploratory Investigation", *Journal of Financial Economics* 8, 323-361.

[19] Poterba, J.M. and L. H. Summers, (1986), "The Persistence of Volatility and Stock Market Fluctuations", *American Economic Review*, 76, 1142-1151.

[20] Schwert, G.W. (1989), "Why does Stock Market Volatility Change Over Time?", *Journal of Finance*, 44, 1115-1154.

[21] Shiryaev, A. N. (1999), *Essentials of stochastic Finance: Facts, Models, Theory*, World Scientific Publisher.

[22] Taylor, S. (1986), *Modeling Financial Time Series*, New York: Wiley.

Table I: GARCH(1,1) parameters obtained using the Drost and Nijman (1993) (dis)aggregation results used in the simulation design

Assets:	YN/US$	S&P 500	YN/US$	S&P 500
	Sampling Frequency and GARCH Parameters			
	Daily frequency		Monthly frequency	
	$m = 1$	$m = 1$	$n = 22$ days	$n = 22$ days
$\phi_{(m)}$	0.026	0.033130	0.34556	0.69811
$\alpha_{(m)}$	0.104	0.028523	0.05944	0.07704
$\beta_{(m)}$	0.844	0.967347	0.24943	0.83594
$\kappa(m)$	3	3	3.77329	3.41591
$v(m)$	0.500	8.021792	0.49999	8.02241
$\kappa(m) \times (v(m))^2$	0.750	193.0474	0.94329	219.844
	One-minute frequency		Five-minute frequency	
	$m = 1440$	$m = 390$	$m = 288$	$m = 78$
$\phi_{(m)}$	0.0000185	0.0000851	0.000093	0.000426
$\alpha_{(m)}$	0.0041994	0.0016560	0.009267	0.003670
$\beta_{(m)}$	0.9957635	0.9983334	0.990550	0.996276
$\kappa(m)$	2.4754561	2.8461378	2.494272	2.856800
$v(m)$	0.498652	8.0283019	0.508197	7.888889
$\kappa(m) \times (v(m))^2$	0.615532	183.44392	0.644180	117.79171

NOTE: The GARCH parameters, $\phi_{(m)}, \alpha_{(m)}, \beta_{(m)}$, are defined in (2.2). The kurtosis parameter is $\kappa(m)$. The unconditional variance is $v(m) = \phi_{(m)}/(1 - \alpha_{(m)} - \beta_{(m)})$. The daily parameters for the YN/US$ are obtained from Andersen and Bollerslev (1998). The S&P 500 estimated Normal GARCH(1,1) parameters cover the daily samples 04/01/86 - 29/08/97.

Table II: Monte Carlo Simulated MSE and MAE Ratios of Volatilities based on the Quadratic Variation of Returns

(A): Contemporaneous MSE and MAE Ratios

Daily Frequency Benchmark: $QV1$		$QV2$	$QV3$	$HQV1$	$HQV2$	$HQV3$	$EHQV1$	$EHQV2$	$EHQV3$
S&P 500	MSE	0.5903	0.5749	0.6735	0.6213	0.6356	0.6680	0.6123	0.6373
	MAE	0.7769	0.7693	0.8295	0.7928	0.8072	0.8253	0.7889	0.8061
YN/US$	MSE	0.6753	0.6783	0.7446	0.8329	0.8096	0.7220	0.8283	0.8219
	MAE	0.8144	0.8266	0.8578	0.9069	0.8974	0.8449	0.8973	0.9042

Monthly Frequency Benchmark: $QV22$		$QV44$	$QV66$	$HQV22$	$HQV44$	$HQV66$
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
YN/US$	MSE	0.4362	0.5630	0.4770	0.6360	0.8188
	MAE	0.7242	0.7959	0.7507	0.8363	0.9155

(B): 1-day, 20-day and 12-month ahead MSE and MAE Ratios

Daily Frequency Benchmark: $QV1$			$QV2$	$QV3$	$HQV1$	$HQV2$	$HQV3$	$EHQV1$	$EHQV2$	$EHQV3$
		h=1								
S&P 500	MSE		0.7279	0.6489	0.9467	0.7881	0.7082	0.9403	0.7893	0.7213
	MAE		0.8840	0.8401	0.9771	0.8840	0.8401	0.9744	0.8828	0.8460
YN/US$	MSE		0.7096	0.6056	0.9179	0.7456	0.6489	0.9123	0.7435	0.6612
	MAE		0.8585	0.8057	0.9487	0.8585	0.8057	0.9450	0.8591	0.8101
		h=20								
S&P 500	MSE		0.6238	0.4847	0.7616	0.5605	0.4595	0.7621	0.5659	0.4655
	MAE		0.7509	0.6802	0.8677	0.7509	0.6802	0.8683	0.7545	0.6865
YN/US$	MSE		0.9892	0.9857	0.9931	0.9876	0.9851	0.9931	0.9884	0.9876
	MAE		0.9999	1.000	0.9999	0.9999	0.9999	0.9899	0.9899	0.9999

Monthly Frequency Benchmark: $QV22$			$QV44$	$QV66$	$HQV22$	$HQV44$	$HQV66$
		h=12					
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
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 obtained from the extraction error in (2.3), using the linear projection GARCH(1,1) equation in (2.5). The volatility filters are discussed in section 2.2. These simulation results refer to the 5 year sample and the 5-minute intraday frequency as reported in section 3.4.

Table III: Monte Carlo Simulated Contemporaneous MSE and MAE Ratios of the Quadratic Variation under conditional Student's t (ν) for S&P 500

Daily Frequency Benchmark: $QV1$	$QV2$	$QV3$	$HQV1$	$HQV2$	$HQV3$	$EHQV1$	$EHQV2$	$EHQV3$
				Student's t with $\nu = 8$				
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

Monthly Frequency Benchmark: $QV22$	$QV44$	$QV66$	$HQV22$	$HQV44$	$HQV66$
		Student's t with $\nu = 8$			
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
		Conditional Normality			
MSE	0.999	0.999	0.0692	0.0660	0.0658
MAE	1.000	1.000	0.2561	0.2554	0.2554
		Student's t with $\nu = 8$			
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. The parameters of the GARCH(1,1) model for the daily S&P 500 are defined in Table I. MSE and MAE are obtained from the extraction error, (2.3).

Table IV: Empirical Results of h-day ahead MSE and MAE Ratios of Daily Volatilities (Benchmark: $QV1$ in MSE, MAE ratios)

	$QV2$	$QV3$	$HQV1$	$HQV2$	$HQV3$	$EHQV1$	$EHQV2$	$EHQV3$
				Five-minute Sampling Frequency DM/US$				
$h = 1$ day								
MSE	0.7341	0.6333	0.8966	0.7017	0.6078	0.8945	0.7219	0.5589
MAE	0.9131	0.8669	0.9538	0.9002	0.8585	0.9514	0.9025	0.8577
$h = 5$ days								
MSE	0.7207	0.6164	0.8932	0.6890	0.5913	0.8908	0.7094	0.5420
MAE	0.9074	0.8693	0.9525	0.8955	0.8523	0.9500	0.8977	0.8506
$h = 20$ days								
MSE	0.6860	0.5727	0.8850	0.6565	0.5506	0.8818	0.6770	0.5053
MAE	0.9030	0.8537	0.9567	0.8921	0.8474	0.9536	0.8939	0.8434
				YN/US$				
$h = 1$ day								
MSE	0.7597	0.6423	0.9060	0.7272	0.6349	0.9060	0.7035	0.6514
MAE	0.9307	0.8928	0.9648	0.9193	0.8846	0.9640	0.9154	0.8889
$h = 5$ days								
MSE	0.7366	0.6212	0.8991	0.7039	0.6056	0.8990	0.6789	0.6227
MAE	0.9225	0.8832	0.9624	0.9139	0.8785	0.9615	0.9098	0.8836
$h = 20$ days								
MSE	0.7457	0.6380	0.9111	0.7235	0.6308	0.9103	0.7012	0.6478
MAE	0.9369	0.9082	0.9779	0.9403	0.9148	0.9762	0.9345	0.9191
				Fifteen-minute sampling frequency DM/US$				
$h = 1$ day								
MSE	0.6807	0.5659	0.8221	0.6297	0.5348	0.8222	0.6373	0.5145
MAE	0.9099	0.8633	0.9512	0.8970	0.8547	0.9504	0.8973	0.8521
$h = 5$ days								
MSE	0.6674	0.5493	0.8154	0.6159	0.5184	0.8155	0.6236	0.4979
MAE	0.9057	0.8557	0.9513	0.8940	0.8489	0.9504	0.8941	0.8464
$h = 20$ days								
MSE	0.6377	0.5119	0.8013	0.5855	0.4822	0.8014	0.5933	0.4629
MAE	0.9002	0.8447	0.9521	0.8852	0.8391	0.9512	0.8852	0.8363

Table IV continued

	YN/US$							
$h = 1$ day								
MSE	0.7418	0.6279	0.8967	0.7066	0.6099	0.8958	0.6998	0.6050
MAE	0.9247	0.8828	0.9643	0.9116	0.8760	0.9638	0.9105	0.8744
$h = 5$ days								
MSE	0.7206	0.5998	0.8907	0.6859	0.5836	0.8897	0.6786	0.5786
MAE	0.9141	0.8716	0.9613	0.9030	0.8653	0.9607	0.9019	0.8642
$h = 20$ days								
MSE	0.7314	0.6189	0.9014	0.7059	0.6106	0.9002	0.6997	0.6063
MAE	0.9248	0.8909	0.9702	0.9243	0.8949	0.9696	0.9229	0.8942
	Thirty-minute sampling frequency DM/US$							
$h = 1$ day								
MSE	0.8316	0.7732	0.9454	0.8188	0.7673	0.9454	0.8196	0.7645
MAE	0.9673	0.9525	0.9924	0.9649	0.9503	0.9924	0.9648	0.9479
$h = 5$ days								
MSE	0.8073	0.7409	0.9375	0.7929	0.7346	0.9376	0.7939	0.7319
MAE	0.9629	0.9457	0.9912	0.9602	0.9433	0.9913	0.9602	0.9408
$h = 20$ days								
MSE	0.6176	0.4866	0.8759	0.5893	0.4743	0.8762	0.5913	0.4733
MAE	0.9079	0.8641	0.9672	0.8999	0.8602	0.9674	0.9001	0.8579
	YN/US$							
$h = 1$ day								
MSE	0.7933	0.7185	0.9532	0.7791	0.7714	0.9538	0.7839	0.7115
MAE	0.9711	0.9521	0.9952	0.9685	0.9690	0.9952	0.9689	0.9475
$h = 5$ days								
MSE	0.7544	0.6655	0.9457	0.7387	0.6579	0.9464	0.7443	0.6585
MAE	0.9658	0.9431	0.9937	0.9638	0.9399	0.9937	0.9643	0.9381
$h = 20$ days								
MSE	0.5886	0.4431	0.9111	0.5652	0.4332	0.9123	0.5746	0.4381
MAE	0.9347	0.9009	0.9765	0.9338	0.9021	0.9767	0.9353	0.9034

NOTE: The h-day ahead forecast MSE and MAE are obtained from the extraction error (2.3) using the h-day linear projection GARCH(1,1) equation in (2.5). The daily GARCH(1,1) parameters are estimated for the 10 year sample period to be for the YN and DM the vectors, (0.0511, 0.919, 0.0028), (0.0548, 0.9240, 0.0112), respectively. The decay parameter in the exponential weights is set equal to 0.999.

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

	Benchmark: 1-month IV ($QV22$)				
	$QV44$	$QV66$	$HQV22$	$HQV44$	$HQV66$
$h = 1$ 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 ($QV22$ + SC22)				
	$QV44+$ SC44	$QV66+$ SC66	$HQV22+$ SC22	$HQV44+$ SC44	$HQV66+$ SC66
$h = 1$ month					
MSE	0.6404	0.5093	0.4964	0.5929	0.6063
MAE	0.9165	0.8884	0.8904	0.9142	1.2361
$h = 6$ 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

	Benchmark: 1-month CAR (CAR22)				
	CAR44	CAR66	HCAR22	HCAR44	HCAR66
$h = 1$ month					
MSE	0.9329	0.9040	0.9620	0.9282	0.9034
MAE	0.9998	0.9995	0.9999	0.9996	0.9994
$h = 6$ months					
MSE	0.7827	0.6903	0.8891	0.7799	0.7016
MAE	0.9466	0.9181	0.9741	0.9489	0.9303
$h = 12$ 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 obtained using the h-month linear projection GARCH(1,1) equation in (2.5). The monthly GARCH parameters for the S&P 500 were obtained from Table I mainly for consistency purposes with the simulation evidence.

SELECTED RECENT PUBLICATIONS

Bera A. K. and Yannis Bilias, Rao´s Score, Neyman´s C (α) and Silvey´s LM Tests: An Essay on Historical Developments and Some New Results, *Journal of Statistical Planning and Inference,* forthcoming.

Bertaut C. and M. Haliassos, Precautionary Portfolio Behavior from a Life - Cycle Perspective, *Journal of Economic Dynamics and Control,* 21, 1511-1542, 1997.

Bilias Y., Minggao Gu and Zhiliang Ying, Towards a General Asymptotic Theory for the Cox model with Staggered Entry, *The Annals of Statistics*, 25, 662-682, 1997.

Blundell R., P. Pashardes and G. Weber, What Do We Learn About Consumer Demand Patterns From Micro-Data?, *American Economic Revie*w, 83, 570-597, 1993.

Bougheas S., P. Demetriades and T. P. Mamouneas, Infrastructure, Specialization and Economic Growth, *Canadian Journal of Economics,* forthcoming.

Caporale W., C. Hassapis and N. Pittis, Unit Roots and Long Run Causality: Investigating the Relationship between Output, Money and Interest Rates, *Economic Modeling*, 15(1), 91-112, January 1998.

Caporale G. and N. Pittis, Efficient estimation of cointegrated vectors and testing for causality in vector autoregressions: A survey of the theoretical literature, *Journal of Economic Surveys*, forthcoming.

Caporale G. and N. Pittis, Unit root testing using covariates: Some theory and evidence, *Oxford Bulletin of Economics and Statistics*, forthcoming.

Caporale G. and N. Pittis, Causality and Forecasting in Incomplete Systems, *Journal of Forecasting,* 16, 6, 425-437, 1997.

Clerides K. S., Lach S. and J.R. Tybout, Is Learning-by-Exporting Important? Micro-Dynamic Evidence from Colombia, Morocco, and Mexico, *Quarterly Journal of Economics* 113(3), 903- 947, August 1998.

Cukierman A., P. Kalaitzidakis, L. Summers and S. Webb, Central Bank Independence, Growth, Investment, and Real Rates", Reprinted in Sylvester Eijffinger (ed), Independent Central Banks and Economic Performance, Edward Elgar, 416-461, 1997.

Dickens R., V. Fry and P. Pashardes, Non-Linearities and Equivalence Scales, *The Economic Journa*l, 103, 359-368, 1993.

Demetriades P. and T. P. Mamuneas, Intertemporal Output and Employment Effects of Public Infrastructure Capital: Evidence from 12 OECD Economies, *Economic Journal,* July 2000.

Eicher Th. and P. Kalaitzidakis, The Human Capital Dimension to Foreign Direct Investment: Training, Adverse Selection and Firm Location”. In Bjarne Jensen and Kar-yiu

Wong (eds), Dynamics,Economic Growth, and International Trade, The University of Michigan Press, 337-364, 1997.

Fry V. and P. Pashardes, Abstention and Aggregation in Consumer Demand, *Oxford Economic Papers*, 46, 502-518, 1994.

Gatsios K., P. Hatzipanayotou and M. S. Michael, International Migration, the Provision of Public Good and Welfare, *Journal of Development Economics*, 60/2, 561-577, 1999.

Haliassos M. and C. Hassapis, Non-expected Utility, saving, and Portfolios, *The Economic Journal*, 110, 1-35, January 2001.

Haliassos M. and J. Tobin, The Macroeconomics of Government Finance, reprinted in J.Tobin, Essays in Economics, vol. 4, Cambridge: MIT Press, 1996.

Haliassos M. and C. Bertaut, Why Do So Few Hold Stocks?, *The Economic Journal*, 105, 1110- 1129, 1995.

Haliassos M., On Perfect Foresight Models of a Stochastic World, *Economic Journal*, 104, 477-491, 1994.

Hassapis C., N. Pittis and K. Prodromidis, Unit Roots and Granger Causality in the EMS Interest Rates: The German Dominance Hypothesis Revisited, *Journal of International Money and Finance*, 18(1), 47-73, 1999.

Hassapis C., S. Kalyvitis and N. Pittis, Cointegration and Joint Efficiency of International Commodity Markets", *The Quarterly Review of Economics and Finance,* 39, 213-231, 1999.

Hassapis C., N. Pittis and K. Prodromides, EMS Interest Rates: The German Dominance Hypothesis or Else?" in European Union at the Crossroads: A Critical Analysis of Monetary Union and Enlargement, Aldershot, UK., Chapter 3, 32-54, 1998. Edward Elgar Publishing Limited.

Hatzipanayotou P., and M. S. Michael, General Equilibrium Effects of Import Constraints Under Variable Labor Supply, Public Goods and Income Taxes, *Economica*, 66, 389-401, 1999.

Hatzipanayotou, P. and M.S. Michael, Public Good Production, Nontraded Goods and Trade Restriction, *Southern Economic Journal,* 63, 4, 1100-1107, 1997.

Hatzipanayotou, P. and M. S. Michael, Real Exchange Rate Effects of Fiscal Expansion Under Trade Restrictions, *Canadian Journal of Economics,* 30-1, 42-56, 1997.

Kalaitzidakis P., T. P. Mamuneas and Th. Stengos, A Nonlinear Sensitivity Analysis of Cross-Country Growth Regressions, *Canadian Journal of Economics,* forthcoming.

Kalaitzidakis P., T. P. Mamuneas and Th. Stengos, European Economics: An Analysis Based on Publications in Core Journals, *European Economic Review*, 43, 1150-1168, 1999.

Kalaitzidakis P., On-the-job Training Under Firm-Specific Innovations and Worker Heterogeneity, *Industrial Relations*, 36, 371-390, July 1997.

Ludvigson S. and A. Michaelides, Does Buffer Stock Saving Explain the Smoothness and Excess Sensitivity of Consumption?, *American Economic Review*, forthcoming

Lyssiotou Panayiota, Dynamic Analysis of British Demand for Tourism Abroad, *Empirical Economics*, forthcoming, 2000.

Lyssiotou P., P. Pashardes and Th. Stengos, Testing the Rank of Engel Curves with Endogenous Expenditure, *Economics Letters*, 64, 61-65, 1999.

Lyssiotou P., P. Pashardes and Th. Stengos, Preference Heterogeneity and the Rank of Demand Systems, *Journal of Business and Economic Statistics*, 17 (2), 248-252, April 1999.

Lyssiotou Panayiota, Comparison of Alternative Tax and Transfer Treatment of Children using Adult Equivalence Scales, *Review of Income and Wealth*, 43 (1), 105-117, March 1997.

Mamuneas, Theofanis P., Spillovers from Publicly – Financed R&D Capital in High-Tech Industries, *International Journal of Industrial Organization*, 17(2), 215-239, 1999.

Mamuneas, T. P. and Nadiri M. I., R&D Tax Incentives and Manufacturing-Sector R&D Expenditures, in *Borderline Case: International Tax Policy, Corporate Research and Development, and Investment*, James Poterba (ed.), National Academy Press, Washington D.C., 1997. Reprinted in *Chemtech*, 28(9), 1998.

Mamuneas, T. P. and Nadiri M. I., Public R&D Policies and Cost Behavior of the US Manufacturing Industries, *Journal of Public Economics*, 63, 57-81, 1996.

Michaelides A. and Ng, S., Estimating the Rational Expectations Model of Speculative Storage: A Monte Carlo Comparison of three Simulation Estimators, *Journal of Econometrics*, 96(2), 231-266, 2000.

Pashardes Panos, Equivalence Scales in a Rank-3 Demand System, *Journal of Public Economics*, 58, 143-158, 1995.

Pashardes Panos, Bias in Estimating Equivalence Scales from Grouped Data, *Journal of Income Distribution*, Special Issue: Symposium on Equivalence Scales, 4, 253-264,1995.

Pashardes Panos., Bias in Estimation of the Almost Ideal Demand System with the Stone Index Approximation, *Economic Journal*, 103, 908-916, 1993.

Spanos Aris, Revisiting Date Mining: ´Hunting´ With or Without a License, *Journal of Methodology*, July 2000.

Spanos Aris, On Normality and the Linear Regression Model, *Econometric Reviews*, 14,195-203, 1995.

Spanos Aris, On Theory Testing in Econometrics: Modeling with nonexperimental Data, *Journal of Econometrics*, 67, 189-226, 1995.

Spanos Aris, On Modeling Heteroscedasticity: The Student's *t* and Elliptical Linear Regression Models, *Econometric Theory*, 10, 286-315, 1994.