A Discrete Sine Transform Approach for Realized Volatility Measurement [*]

Giuseppe Curci [1]

Fulvio Corsi [2]

First Version: March 2003
Current Version: March 2003

Abstract

Realized volatility affords the ex-post empirical measurement of the latent notional volatility. However, the time-varying returns autocorrelation induced by microstructure effects represents a challenging problem for standard volatility measures. In this study, a new nonparametric volatility measures approach based on the *Discrete Sine Transform* (DST) is proposed. We show that the DST exactly diagonalizes the covariance matrix of MA(1) process. This original result provides us an orthonomal basis decomposition of the return process which permits to optimally disentangle the underlying efficient price signal from the time-varying nuisance component contained in tick-by-tick return series. As a result, two nonparametric volatility estimators which fully exploit all the available information contained in high frequency data are constructed. Monte Carlo simulations based on a realistic model for microstructure effects show the superiority of DST estimators, compared to alternative local volatility proxies for every level of the noise to signal ratio and a large class of noise contaminations. These properties make the DST approach a nonparametric method able to cope with time-varying autocorrelation, in a simple and efficient way, providing robust and accurate volatility estimates under a wide set of realistic conditions. Moreover, its computational efficiency makes it well suitable for real-time analysis of high frequency data.

[1] INFN (Istituto Nazionale Fisica Nucleare), Pisa, Via F. Buonarroti 2, 56100 - Pisa, Italy
Department of Physics, University of Pisa, Via F. Buonarroti 2, 56100 - Pisa, Italy
e-mail: curci@df.unipi.it

[2] Institute of Finance, University of Southern Switzerland, Via Buffi 13, CH-6904 - Lugano, Switzerland
e-mail: fulvio.corsi@lu.unisi.ch

1 Introduction

Asset returns volatility is a central feature of many prominent financial problems such as asset allocation, risk management and option pricing. But, despite its key role, it is still an ambiguous term for which there are different concepts and definitions.

In a frictionless continuous-time no arbitrage price process framework, three different conditional volatility concepts can be defined[1]:

(i) the *Notional*, actual or integrated ex-post volatility over a non-vanishing interval,

(ii) the ex-ante *Expected* volatility over a non-vanishing interval,

(iii) the *Instantaneous* volatility.

The notional volatility refers to the ex-post cumulative sample-path return variability over a discrete time interval which under very general conditions corresponds to the increments in the quadratic variation of the return process.

In practice the approaches for empirically quantifying the concept of volatility fall in to two distinct types of categories:

- estimation of *Parametric* models,

- direct *Nonparametric* methods.

So far, most of the studies have focused on the parametric approach considering volatility as an unobservable variable and using a fully specified functional model for the ex-ante expected volatility[2]. Modelling the unobserved conditional variance was one of the most prolific topics in the financial literature which led to all ARCH-GARCH models and stochastic volatility models. In general this kind of models suffer from a twofold weakness: first, they are not able to replicate main empirical features of financial data; second, the estimation procedures required are often rather complex (especially in the case of stochastic volatility models).

This study focus instead on a nonparametric approach to develop ex-post observable proxies for the notional volatility (rather than the expected one) through a new methodology which fully exploits intraday information.

Within the class of nonparametric volatility measurement we can distinguish between the *ARCH Filters or Smoothers* and the *Realized Volatility measures* . The first heavily rely on a continuous sample paths assumption on the price process in order to evaluate the instantaneous volatility. Filters exploit only the information contained in past returns while smoothers also use ex-post future returns (thus they can be seen as two-sided filters). These instantaneous volatility measures require that as the length of the time interval goes to zero the number of observations tends to infinity. However this strong condition (which implies a double limit theory and excludes jumps from both return and volatility processes) are virtually never fulfilled in empirical data, making this approach unfeasible in practice.

On the contrary, realized volatility affords the empirical measurement of the latent notional volatility on the discrete time interval $[t-h, t]$, with h a strictly positive non-vanishing quantity (typically one day). Similarly to the instantaneous volatility measures, realized volatilities may be classified according to whether the estimation of the notional volatility only exploits returns observations falling in the interval $[t-h, t]$, which we call *Local*, or also incorporates returns

---

[1]See Andersen, Bollerslev and Diebold (2002).

[2]Also the Implied Volatility approaches can be included in this category since they are based on a parametric model for the returns together with an option pricing model.

outside $[t-h, t]$. Local measurements have the advantage to be asymptotically unbiased and fast adapting but the disadvantage to neglect potentially useful information contained in adjacent intervals. The most obvious local measure for daily volatility is the daily absolute return. However, as clearly shown by Andersen and Bollerslev (1998) this proxy can be extremely noisy. The inadequacy of volatility proxies obtained with daily observations clearly suggests the use of intraday data to obtain more accurate volatility estimates. In fact, in its standard form realized volatility is nothing more than the sum of squared high-frequency returns over a given time interval $[t-h, t]$, i.e. the second uncentered sample moment of the high-frequency returns. This idea traces back to the seminal work of Merton (1980) who showed that the integrated variance of a Brownian motion can be approximated to an arbitrary precision using the sum of intraday squared returns. More recently a series of papers (Andersen, Bollerslev, Diebold and Labys 2001a,b and Barndoff-Nielsen and Shepard 2001a, 2002a,b,c and Comte and Renault 1998) has formalized and generalized this intuition by applying the quadratic variation theory to the broad class of special (finite mean) semimartingales[3]. In fact, under very general conditions the sum of intraday squared returns converges, as the maximal length of returns go to zero, to the notional volatility over the fixed time interval $[t-h, t]$. Thus, as the sampling frequency from a diffusion is increased, realized volatility provides us, in principle, with a consistent nonparametric measure of the notional volatility.

In practice, however, empirical data differs in many ways from the frictionless continuous-time price process assumed in those theoretical studies. Beside the obvious consideration that a continuous record of prices is not available, other reasons prevent the applications of the limit theory necessary to achieve consistency of the realized volatility estimator. In fact, because of market microstructure effects[4] the assumption that log asset prices evolve as a diffusion process becomes less and less realistic as the time scale reduces.

The main sources of microstructure effects are the bid-ask bounce and price discreteness. Studies on the bid-ask spread are largely developed within the framework of quote-driven markets. However, the bid-ask spread is not unique to the dealer markets: Cohen et al. (1981) establish the existence of the bid-ask spread in a limit-order market when investors face transaction costs in assessing information, monitoring the market, and conveying orders to the market; Glosten (1994) shows that limit-order markets have a positive bid-ask spread arising from the possibility of trading on private information. As already noted by Roll (1984) and Blume and Stambaugh (1983), bid-ask spreads produce negative first-order autocovariances in observed price changes. Similarly, if one makes the assumption that observed prices are obtained by rounding underlying true values, Glottlieb and Kalay (1985) and Harris (1990) showed that price discreteness induces negative serial covariance in the observed returns.

Market microstructure generates a transitory effect on the dynamics of the informationally efficient price. This perturbation of the underlying price induces a non-zero autocorrelation in the returns process which makes no longer true that the variance of the sum is the sum of the variances. Thus, the volatility computed with short time intervals becomes a potentially highly biased estimator of the daily volatility. A significant negative autocorrelation induces a bias of positive sign, i.e. the expectation of daily realized volatility computed with high frequency observed returns is systematically larger than the volatility of the true unobservable process. Such a bias increases with the sampling frequency. Therefore, a trade-off arises: on one hand, efficiency considerations suggest a very high number of return observations to reduce the stochastic error of volatility estimation. On the other hand, market microstructure introduces

---

[3]This class encompasses processes used in standard arbitrage-free asset pricing applications, such as Ito diffusions, jump processes, and mixed jump diffusions.

[4]For a good empirically oriented overview of market microstructure effects, see Hasbrouck (1996).

a bias that grows as the sampling frequency increases.

Given such a trade-off between efficiency and bias, a simple approach to overcome this problem is to choose, for each financial instruments, the shortest return interval at which the resulting volatility is still not significantly affected by the bias. This approach exploits the different aggregation properties between the integrated process of the efficient price and the non-scaling behaviour of the pricing error term. Thus, as the aggregation of returns increases the impact of the transitory component on the volatility decreases, reducing the size of the bias. This approach is simple, fully nonparamentric and robust to any source of microstructure effect. On the other hand, in practice this "unbiased return frequency" turns out to be fairly low[5], leaving us with only few return observations per day.

An alternative approach would be to directly try to correct for the microstructure effects at the tick-by-tick level. Along this line are the first order serial covariance correction proposed by French and Roll (1984), Harris (1990) and Zhou (1996) and the exponential moving average (EMA) filtering of Corsi, Zumbach, Müller and Dacorogna (2001). However, the first estimator suffers from the possibility to become negative while the second one is a non-local estimator which adapts only slowly to changes in the properties of the pricing error component. Moreover, both estimators correct only for the bias deriving from the first lag of the return autocorrelation function, while they are very sensitive to non zeros higher lags coefficients.

The presence of significant autocorrelation at lags length greater than one and the possibility that each trading day may be characterized by different autocorrelation structures makes the filtering problem rather complex. In theory, this problem could be tackled by a fully parametric approach where several ARMA models are first estimated every day. Then the best model is chosen on the basis of some loss criteria and finally an estimate for the daily volatility could be obtained from the residuals of the selected model. This approach has been proposed, for instance, by Bollen and Inder (2002) with the so called VARHAC estimator which makes use of a series of AR models selected on the basis of the Schwarz BIC criteria. However, such a parametric approach is, beside being very cumbersome in practice, much dependent on the loss criteria chosen and, relaying on the estimation of (possibly many) parameters, it conveys the estimation errors of those parameters to the volatility estimator, amplifying its variance. Moreover, Bustos and Yohai (1986) show how even very few outlying observations can largely increase the variance of the estimated residuals. What we are looking for is, instead, a nonparametric approach capable to deal with time-varying autocorrelation (with possibly several significant lags) in a simple and efficient way, providing robust and accurate volatility estimates.

In this paper two new alternative realized volatility measures based on the *Discrete Sine Transform* (DST) are presented. The motivation for this approach rests on the ability of the DST to decorrelate signal for data exhibiting MA(1) type of behaviour. To our knowledge this is an original result in signal theory. MA(1) processes arise naturally in microstructure models of tick-by-tick returns. Hence, this nonparametric DST approach, turns out to be very convenient as it provides an orthonormal basis which permits to optimally[6] extract the volatility signal from the noisy tick-by-tick return series. As a result, new nonparametric realized volatility estimators which fully exploit all the available information contained in high frequency data can be constructed. Moreover, we show that this approach provides robust and accurate results also in case of microstructure effects which lead to more general MA(Q) processes for the tick-by-tick returns. It is then robust against a wide class of noise contaminations and model misspecifications. Finally, being simple and computationally fast, it is perfectly suitable for real-time

---

[5]The answer to this question also depends on the interpolation scheme employed when a regular time series is constructed.

[6]In a linear sense.

analysis of high frequency data.

The rest of the paper is organized as follows. Section 2 reviews a model for the tick-by-tick observed price process, shows how this process can be diagonalized by the DST approach, describes the volatility estimators based on the DST filter and analyzes their robustness with respect to model misspecification. Section 3 outline the setup for the Monte Carlo simulations and shows the superiority of DST estimators compared to alternative realized volatility estimators. Section 4 concludes.

2 Definitions of the DST volatility estimators

2.1 Price process with microstructure effects

In order to motivate the subsequent definitions of volatility, we briefly review a standard model for the tick-by-tick price process. As described in Hasbrouck (1993, 1996), a general way to model the impact of various sources of microstructure effects is to decompose the observed price in to the sum of two unobservable components: a martingale component representing the informationally efficient price process and a stationary pricing error component expressing the discrepancy between the efficient price and the observed one. The pricing error term impounds the diverse microstructure imperfections which are not explicitly modelled. The price evolution is described in the "intrinsic transaction time"[7] denoted with the integer index n; so that $\Delta t_n = t_n - t_{n-1}$ is the "intertrade duration" process which is also not explicitly modelled here.

The observed logarithmic prices in tick time p_n are then given by

$$p_n = \tilde{p}_n + \eta \omega_n$$

where $\tilde{p}_n$ is the unobserved efficient price which follows a stochastic process with unforecastable changes, and $\eta \omega_n$ represents the pricing error component with η the size of the perturbation. Depending on the structure imposed on the pricing error component, many structural model for microstructure effect could be recovered. Here we take a more statistical perspective assuming ω_n to simply be a (possibly non-spherical) zero mean nuisance component independent of the price process.

The observed (logarithmic) price return r_n at time n can then be decomposed as

$$r_n = \sigma \tilde{\epsilon}_n + \eta \left(\omega_n - \omega_{n-1} \right)$$

where $\tilde{\epsilon}_n$ (the unobserved innovation of the efficient price) and ω_n (the pricing error) are independent stochastic processes such that

$$\tilde{\epsilon}_n \sim \text{IID} \left(0, 1 \right)$$
$$\omega_n \sim \text{IID} \left(0, 1 \right)$$

Hence the r-process is MA (1) with $\mathbb{E} \left(r_n \right) = 0$ and autocovariance function given by

$$\rho \left(k \right) = \mathbb{E} \left(r_{n+k} r_n \right) = \left(\sigma^2 + 2\eta^2 \right) \delta_{k,0} - \eta^2 \left(\delta_{k,+1} + \delta_{k,-1} \right)$$

where δ is the discrete Kronecker delta function.

---

[7]That is a time scale having the number of trades as its directing process, which maps the intrinsic time to the calendar time.

2.2 The Discrete Sine Transform

We now develop a discrete form of a Karhunen-Loéve expansion[8] for the representation of a random vector whose elements are a finite sequence of returns.

Let us consider a vector of length M

$$R(M,n) = \begin{bmatrix} r_n & r_{n-1} & \cdots & r_{n-M+1} \end{bmatrix}^\top$$

and the associated correlation matrix

$$\Omega^{(M)} = \mathbb{E}\left(R(M,n)\, R(M,n)^\top\right)$$

which is a tridiagonal matrix of the form

$$\Omega^{(M)} = \begin{bmatrix} \sigma^2 + 2\eta^2 & -\eta^2 & & \\ -\eta^2 & \sigma^2 + 2\eta^2 & \ddots & \\ & \ddots & \ddots & -\eta^2 \\ & & -\eta^2 & \sigma^2 + 2\eta^2 \end{bmatrix}$$

Let us now solve the eigenvalues equation

$$\Omega^{(M)} \varphi_m^{(M)} = \lambda_m^{(M)} \varphi_m^{(M)} \qquad m = 1, 2, \ldots, M$$

It can be easily shown that the eigenvalues of $\Omega^{(M)}$ are given by

$$\lambda_m^{(M)} = \sigma^2 + 4\eta^2 \sin^2 \frac{\pi m}{2(M+1)} \qquad m = 1, 2, \ldots, M$$

with

$$\lambda_{m+1}^{(M)} - \lambda_m^{(M)} = 4\eta^2 \sin \frac{\pi}{2(M+1)} \sin \frac{\pi(2m+1)}{2(M+1)} > 0$$

and $\lambda_1^{(M)} < \lambda_2^{(M)} < \cdots < \lambda_M^{(M)}$. The corresponding eigenvectors are

$$\varphi_m^{(M)}(k) = \sqrt{\frac{2}{M+1}} \sin \frac{\pi m k}{M+1} \qquad k = 1, 2, \ldots, M$$

The remarkable fact is that, unlike common situations, the eigenvectors $(\varphi_m^{(M)})$ of a MA (1) process are universal as they are given by the orthonormal basis used in the Discrete Sine Transform (DST). Moreover the eigenvalues of the DST components are ordered, separated and all non degenerate. Given that the Karhunen-Loéve expansion represents the optimal solution to a linear filtering problem, this nonparametric property can be very useful for real-time analysis of high frequency return data as it provides an universal basis to optimally decorrelate the price signal.

According to the Karhunen-Loéve expansion, the simple and computationally fast DST of the returns

$$c_m^{(M)}(n) = \sum_{k=1}^{M} \varphi_m^{(M)}(k)\, r_{n-k+1}$$

---

[8]Also known as Hotelling transformation or Principal Component Analysis.

acts as a projector of the signal into its principal components. The autocovariance functions of the DST components are directly the eigenvalues of the correlation matrix:

$$\mathbb{E}\left(c_m^{(M)}(n)\, c_m^{(M)}(n)\right) = \left(\varphi_m^{(M)}\right)^T \Omega^{(M)} \varphi_m^{(M)} = \lambda_m^{(M)} = \sigma^2 + 4\eta^2 \sin^2 \frac{\pi m}{2(M+1)}$$

Since we are interested in the permanent component of the volatility the idea is to consider the projection of the returns on the minimal principal component which is the one less contaminated by the transient volatility coming from the microstructure noise. Therefore, an asymptotically unbiased estimator of the instantaneous volatility σ^2 is given by the mean value of the square of the DST component associated with the minimal eigenvalue of the correlation matrix $(c_1^{(M)})$ in the limit of a large window M:

$$\sigma_M^2 = \mathbb{E}\left(c_1^{(M)}(n)\, c_1^{(M)}(n)\right) = \sigma^2 + 4\eta^2 \sin^2 \frac{\pi}{2(M+1)} \tag{1}$$

since for large M the effect of the price error vanishes as

$$\sigma_M^2 \simeq \sigma^2 + \eta^2 \frac{\pi^2}{M^2}$$

The last equation clearly shows how the aggregation on the minimal component decreases the impact of the pricing error at a much higher speed compared with the standard aggregation of returns. In fact in this second case the bias is reduced at the rate M while on the minimal DST component the bias is cut down at rate M^2, allowing to substantially increase the "unbiased return frequency" and then improving the precision of the volatility estimation. We term this volatility measure *Minimal DST estimator*.

Another way to construct a simple volatility estimator from the DST decomposition, is to evaluate σ_M^2 for different values of M and then perform a simple linear regression on the equation (1). Then the intercept is an unbiased (not only asymptotically but also in finite sample) estimator of the instantaneous volatility. We call this measure *Fitted DST estimator*. This approach would be particularly useful when the number of observations is not very high and thus sufficiently large values of M are not attainable.

2.3   Stability and robustness

To judge the stability of the DST filter with respect to misspecification of the underlying model, let us consider an $\mathrm{MA}\,(Q)$ process

$$r_n = \sigma \tilde{\epsilon}_n + \sum_{q=1}^{Q} \eta_q \left(\omega_n^{(q)} - \omega_{n-q}^{(q)}\right)$$

where

$$\tilde{\epsilon}_n \sim \mathrm{IID}\,(0,1)$$
$$\omega_n^{(q)} \sim \mathrm{IID}\,(0,1) \quad q = 1, \ldots, Q$$

It can be shown that the autocorrelation of the Minimal DST component is given by

$$\sigma_M^2 = \mathbb{E}\left(c_1^{(M)}(n)\, c_1^{(M)}(n)\right) = \sigma^2 + \sum_{q=1}^{Q} \eta_q^2 F\,(M,q)$$

6

where

$$F\left(M, q\right) = \frac{2}{M+1}\left[M + 1 - \left(M + 1 - q\right)\cos\frac{\pi q}{M+1} - \cot\frac{\pi}{M+1}\sin\frac{\pi q}{M+1}\right]$$

Because

$$F\left(M, q\right) = \pi^2\frac{q^2}{M^2} - 2\pi^2\frac{q^3}{M^3}\left(1 + \frac{3}{q} - \frac{1}{q^2}\right) + O\left(\frac{q^4}{M^4}\right)$$

for $M/Q \to \infty$, we obtain

$$\sigma_M^2 \simeq \sigma^2 + \frac{\pi^2}{M^2}\sum_{q=1}^{Q}\left(q\,\eta_q\right)^2$$

which makes clear that also the bias coming from higher order autocorrelations is cut down at the same rate M^2, guaranteeing the robustness of the DST estimators respect to a wide class of model misspecifications.

3   Monte Carlo Simulations

In this section we analyze the model proposed by Hasbrouck (1999) and recently employed in Alizadeh et al. (2002). The model views the discrete bid and ask quotes as arising from the efficient price plus the quote-exposure costs (information and processing costs). Then the bid price is the efficient price less the bid cost rounded down to the next tick; the ask quote is the efficient price plus the ask cost rounded up to the next tick. As in Alizadeh et al. (2002) we simplify the model by assuming that the bid cost and the ask cost are both equal to the minimum tick size. Then assuming a gaussian logarithmic random walk for the log price

$$\tilde{p}_n = \tilde{p}_{n-1} + \sigma\epsilon_n$$

with $u_n \sim \text{NID}\left(0, 1\right)$, the bid and ask prices are

$$B_n = \Delta\left\lfloor \tilde{P}_n/\Delta - 1\right\rfloor$$
$$A_n = \Delta\left\lceil \tilde{P}_n/\Delta + 1\right\rceil$$

where Δ represents the ticksize, $\lfloor x \rfloor$ is the floor function, $\lceil x \rceil$ the ceiling one and the unobserved efficient price is $\tilde{P}_n = e^{\tilde{p}_n}$. Then the observed price is given by the following bid-ask model

$$P_n = B_n q_n + A_n\left(1 - q_n\right)$$

with $q_n \sim \text{Bernoulli}\left(1/2\right)$. Hence the observed logarithmic return can be written

$$r_n = \ln\frac{P_n}{P_{n-1}} = \ln\frac{\lceil P_n/\Delta + 1\rceil}{\lceil P_{n-1}/\Delta + 1\rceil} + q_n\ln\frac{\lfloor P_n/\Delta - 1\rfloor}{\lceil P_n/\Delta + 1\rceil} - q_{n-1}\ln\frac{\lfloor P_{n-1}/\Delta - 1\rfloor}{\lceil P_{n-1}/\Delta + 1\rceil}$$

which is an MA (1) process.

Following Hasbrouck and Alizadeh et al. we first take $P_0 = 25$, $\Delta = 1/16$ and $\sigma = 0.0011$, which implies an annualized 30% volatility, and simulate one-day sample paths of five-minute log price (which means 288 returns observations per day) for 10,000 days. This choice of the
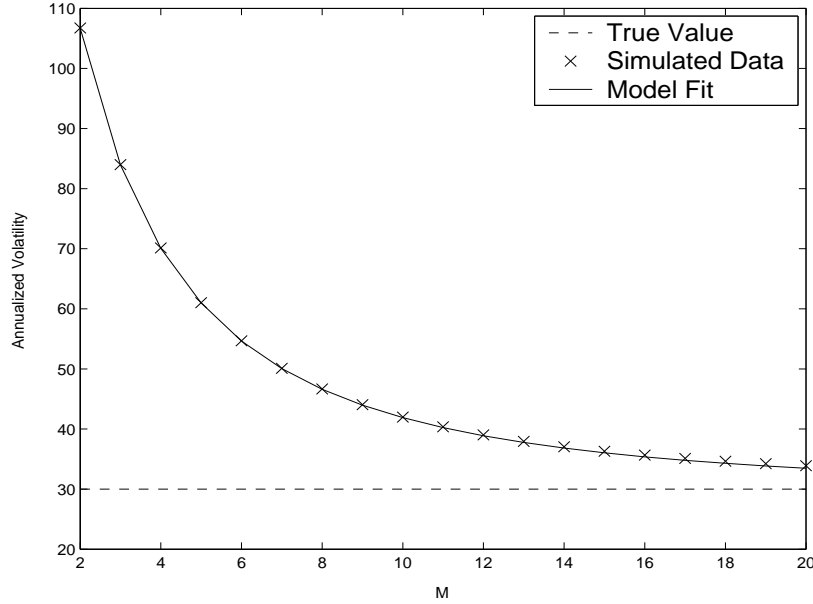
Figure 1: Comparison of the minimal DST annualized volatility predicted by the DST model for different value of M and the ones obtained by simulations $\left(\sqrt{\hat{\sigma}_M^2 \cdot 288 \cdot 250} \right)$.

parameters implies a strong price fluctuation between bids and asks, inducing a highly negative first order-autocorrelation $\rho(1) \approx 48\%$ for the process r_t and a noise to signal ratio $\left(\frac{\eta}{\sigma} \right)$ of about 3.5. This values reflect a microstructure impact on the return process which is remarkably large and rarely observed on real data. However, such an extreme setting provides a useful stress test for realized volatility measures and harden the competition versus range-based estimators which are favourite under this circumstances.

We compute the Minimal DST estimator for $M = 20$. While, for the Fitted DST we first construct a series of minimal eigenvalues of the Karhunen-Loéve expansion, using a sequence of DST windows from $M = 2$ to $M = 20$ and then fit (by simple OLS) the eigenvalues to the equation (1) which we report here:

$$\sigma_M^2 = \sigma^2 + \eta^2 4 \sin^2 \frac{\pi}{2(M+1)}$$

figure 1 shows the perfect fit of the model predictions to the simulated data.

For comparison purposes we compute the daily absolute return, the daily range[9] and two alternative realized volatility measures: the one obtained with a local EMA filter (i.e. calibrated on a single day) which then simply corresponds to a daily MA(1) filter and the one calculated with 30 minutes returns as suggested by Andersen et al. (2001b).

Table 1 reports the mean, standard deviation and Root Mean Square Error (RMSE) of the five volatility estimators and figure 2 shows their probability density functions. Given the high level of noise and the relatively small number of observations per day, the estimation of the first order autocorrelation required to calibrate the EMA filter, is very noisy and does not always satisfy the theoretical bound for MA(1) process $| \rho(1) | < 1/2$ (in the 30% of the cases), leading to a complex MA(1) coefficient θ. In such cases the EMA filter would fail and we are then

---

[9]Developed by Parkinson (1980) and recently advocated by Alizadeh et al. (2002) in the contest of stochastic volatility models estimation.
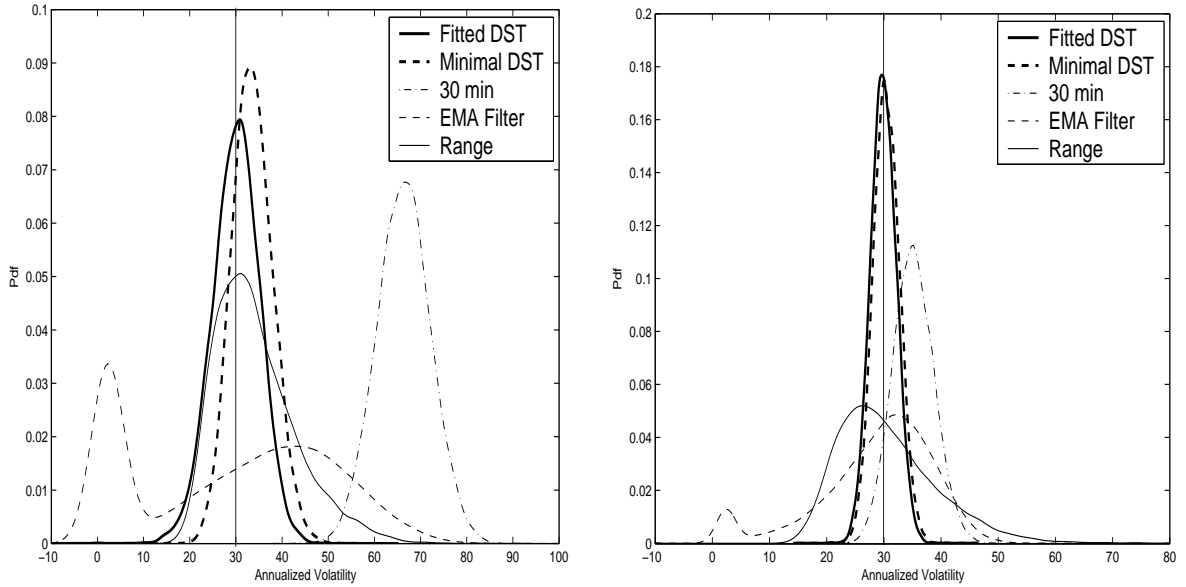
Figure 2: Comparison of the pdf of the different volatility estimators obtained at 5 minutes (left panel) and 30 seconds (right panel) frequency with $\frac{\eta}{\sigma} = 3.5$ (true value for the annualized volatility: 30%).

forced to impose an artificial floor to $\rho(1)$. But besides its arbitrariness, this procedure induces unreasonably low volatility estimates (responsible for the left bump close to zero presents in the EMA estimator pdf of figure 2). Moreover, under these conditions, the variance of the estimator is extremely large. For the 30 minutes realized volatility, the fact that the aggregation from 5 minute to half an hour returns is not able to eliminate all the negative autocorrelation, makes this estimator strongly upward biased (with an annualized mean value of 66% against a true value of 30%). In the case of the Minimal DST estimator instead, the aggregation works much better but, due to a relatively low $M = 20$, a small upward bias of about 12% is still present. Even the daily range[10] suffers a bias of a similar size (17%) but with a much larger variance (two times respect to the Minimal DST one). Under this extreme setting, the only measure which is still able to be unbiased and sufficiently precise is the Fitted DST estimator, which has in fact a much lower RMSE. Summarizing, even in the most unfavorable setting for the realized volatilities, the DST estimators are still quite accurate and posses a RMSE 42% to 47% smaller than that of the daily range.

Keeping the same level of noise we repeat the simulation at 30 seconds frequency (which means 2880 observations per day). With ten times more data the realized volatility measures are much more precise: the local EMA filter has less failings (5%) and lower variance, while the 30 min realized volatility (thanks to the longer aggregation period) has a smaller bias. DST estimators are now both unbiased and equally very accurate, remaining by far the best choices among the estimators considered.

But real financial time series present a noise to signal ratio at tick-by-tick level usually comprises between 1 and 2. However, even with such a moderate level of noise, a naive high frequency realized volatility measure would be from two to three times the actual one. We then repeat the simulation with a more realistic noise to signal ratio of 1.5 for both frequency. Table 2 and figure 3 summarize the results. At the five minute frequency the daily range and the local

---

[10]Evaluated on the price process observed at the highest frequency (i.e. five minute in this case).

9

<div align="center">VOLATILITIES ESTIMATES WITH $\frac{\eta}{\sigma}=3.5$</div>

	5 min			30 sec		
	mean	std	RMSE	mean	std	RMSE
Fitted DST	30.239	5.055	5.057	29.932	2.223	2.227
Minimal DST	33.745	4.381	5.598	30.467	2.271	2.303
EMA filter	35.200	20.811	20.871	30.147	9.995	10.114
30 minutes	66.251	5.798	36.482	35.220	3.561	6.177
Daily Range	35.354	8.800	9.778	31.675	8.582	8.597
Daily $\mid r \mid$	31.217	18.970	19.665	30.221	18.306	19.246

Table 1: Comparison of daily volatility estimation performances of the DST estimators (Fitted DST and Minimal DST), the EMA filter and 30 minutes realized volatility, the daily range and the daily absolute return for the 5 minutes and 30 seconds frequency when the true value of the annualized volatility is 30% and the noise to signal ratio $\frac{\eta}{\sigma}=3.5$.

EMA filter are unbiased but quite inaccurate while the realized volatility based on 30 minutes has a large bias. Again the DST estimators are the more accurate, with the Fitted DST having a slightly smaller variance than the Minimal DST. At the 30 seconds frequency, with a moderate level of noise and a large number of data, the EMA filter start to become competitive compares to the DST ones; while the 30 minutes measure (even if less biased) has a much larger variance.

Empirical studies on the autocorrelation of tick-by-tick data often show significative values not only for the first order but also for the second order lag (though of much smaller amplitude). A possible explanation, and way to model it, is by assuming a correlation in the sequence on which ask and bid prices arrive[11]. Hence, instead of having an "unbiased" Bernoulli$(1/2)$ for the q_n process, we construct a Bernoulli process which produces an autocorrelation in q_n. This "biased" Bernoulli is obtained by taking $q_n = $ Bernoulli$(1/2 + b)$ if $q_{n-1} = 1$ and $q_n = $ Bernoulli$(1/2 - b)$ if $q_{n-1} = 0$. We choose $b = -0.10$ which induces a second order autocorrelation of about -6%. Now the local EMA filter, which was unbiased and very precise at 30 seconds frequency, becomes highly biased at both frequency (see figure 4). On the contrary DST estimators remain unbiased and accurate as before, clearly showing their robustness against model mispecifications (as analytically described in the previous section).

---

[11]Hasbrouck and Ho (1987) suggest that positive autocorrelation at lag lengths greater than one may be the results of traders working an order: "a trader may distribute purchases or sales over time". However also significantly negative autocorrelation at lag two are often observed.
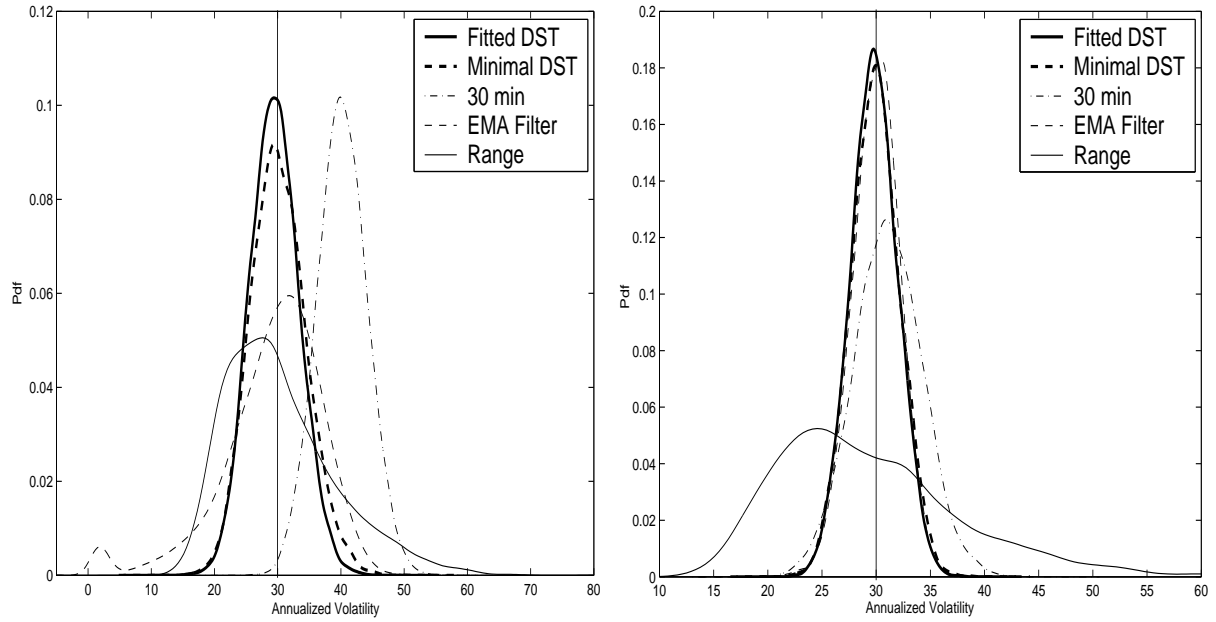
Figure 3: Comparison of the pdf of the different volatility estimators obtained at 5 minutes (left panel) and 30 seconds (right panel) frequency with $\frac{\eta}{\sigma} = 1.5$ (true value for the annualized volatility: 30%).
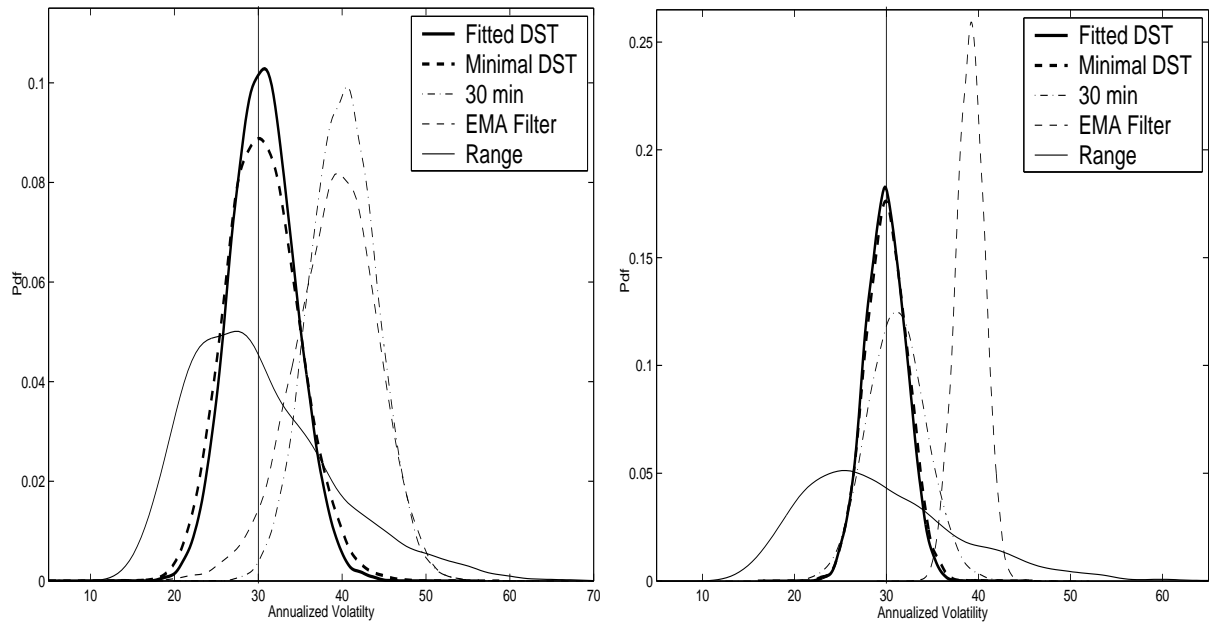

Figure 4: Comparison of the pdf of the different volatility estimators obtained at 5 minutes (left panel) and 30 seconds (right panel) frequency with $\frac{\eta}{\sigma} = 1.5$ and biased Bernoulli process (true value for the annualized volatility: 30%).

11

VOLATILITIES ESTIMATES WITH $\frac{\eta}{\sigma} = 1.5$

	5 min			30 sec		
	mean	std	RMSE	mean	std	RMSE
Unbiased Bernoulli						
Fitted DST	29.678	3.843	3.880	29.832	2.146	2.160
Minimal DST	30.336	4.353	4.353	29.936	2.238	2.243
EMA filter	30.078	7.970	8.027	29.986	2.157	2.159
30 minutes	40.289	3.897	10.863	31.153	3.151	3.303
Daily Range	31.443	8.715	8.718	30.514	8.627	8.657
Daily $\mid r \mid$	30.784	18.314	19.042	30.312	18.303	19.211
Biased Bernoulli						
Fitted DST	30.735	3.781	3.815	29.931	2.175	2.179
Minimal DST	30.736	4.319	4.342	30.070	2.270	2.270
EMA filter	39.1621	5.026	10.176	39.079	1.542	9.176
30 minutes	40.237	3.932	10.796	31.203	3.138	3.307
Daily Range	31.336	8.743	8.743	30.522	8.567	8.596
Daily $\mid r \mid$	30.599	18.422	19.241	30.170	18.173	19.111

Table 2: Comparison of daily volatility estimation performances of the DST estimators (Fitted DST and Minimal DST), the EMA filter and 30 minutes realized volatility, the daily range and the daily absolute return for the 5 minutes and 30 seconds frequency when the true value of the annualized volatility is 30% and the noise to signal ratio $\frac{\eta}{\sigma} = 1.5$.

4   Conclusions

The autocorrelation induced by microstructure effects represents a challenging problem for realized volatility measures. It makes the naive realized volatility computed at short time interval highly biased. While the time-varying behavior of the autocorrelation makes non-local filters misspecified. Local filters based on first order covariance correction are also prone to misspecification (due to the occasional significance of higher order lags) and suffer from the unpleasant possibility to become negative.

In this study, a new nonparametric volatility measures approach based on the *Discrete Sine Transform* (DST) is proposed. This approach is justified by the new theoretical result on the ability of the DST to exactly diagonalize an MA(1) process. Hence, we utilize the DST orthonomal basis decomposition to optimally disentangle the underlying efficient price signal from the time-varying nuisance component contained in tick-by-tick return series. As a results two nonparametric realized volatility estimators which fully exploit all the available information contained in high frequency data are constructed.

Monte Carlo simulations based on a realistic model for microstructure effects, show the superiority of DST estimators compared to alternative local volatility proxies such as the daily range, the EMA filter and 30 minutes realized volatility and the daily absolute return. DST estimators results to be the most accurate daily volatility measures for every level of the noise to signal ratio, and highly robust against the presence of significant autocorrelation at lags greater than one in the return process.

Those properties make the DST approach a nonparametric method able to cope with time-varying autocorrelation, (with possibly several significant lags) in a simple and efficient way, providing robust and accurate volatility estimates under a very wide set of conditions. Moreover, its computational efficiency, makes it well suitable for real-time analysis of high frequency data.

An obvious direction for future research would be to empirically analyse the forecasting performance of parametric and nonparametric volatility models fed by the DST volatility estimates. A thorough investigation of the performance of the DST estimators applied to various type of financial time-series and their use in conjunction with adaptive Kalman filter models, are currently pursed by the authors.

References

Alizadeh, S., Brandt, M. W. and Diebold, F. X.: 2002, Range-based estimation of stochastic volatility models, *Journal of Finance* **LVII**(3), 1047–90.

Andersen, T. G. and Bollerslev, T.: 1998, Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, *International Economic Review* **39**, 885–905.

Andersen, T. G., Bollerslev, T. and Diebold, F. X.: 2002, Parametric and nonparametric volatility measurement, *in* Y. Ait-Sahalia and L. P. Hansen (eds), *Handbook of Financial Econometrics*, North Holland, Amsterdam.

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P.: 2001a, The distribution of exchange rate volatility, *Journal of the American Statistical Association* **96**, 42–55.

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P.: 2001b, Modeling and forecasting realized volatility, *NBER Working Paper 8160. The National Bureau of Economic Research, Cambridge, MA.* .

Barndoff-Nielsen, O. E. and Shepard, N.: 2001, Non-gaussian ornstein-uhlembech-based models and some of their uses in financial economics, *Journal of the Royal Statistical Society* **63**, 167–241.

Barndoff-Nielsen, O. E. and Shepard, N.: 2002a, Econometric analysis of realized volatility and its use in estimating stochastic volatility models, *Journal of the Royal Statistical Society* **64**, 253–280.

Barndoff-Nielsen, O. E. and Shepard, N.: 2002b, Estimating quadratic variation using realized variance, *Forthcoming in Journal of Applied Econometrics* .

Barndoff-Nielsen, O. E. and Shepard, N.: 2002c, How accurate is the asymptotic approximation to the distribution of realized volatility?, *in* D. W. F. Andrews, J. L. Powell, P. A. Ruud and J. H. Stock (eds), *Identification and Inference for Econometric Models*, forthcoming.

Blume, M. and Stambaugh, S.: 1983, Biases in computed returns, *Journal of Financial Economics* **12**, 387–404.

Bollen, B. and Inder, B.: 2002, Estimating daily volatility in financial markets utilizing intraday data, *Journal of Empirical Finance* **186**.

Bustos, O. and Yohai, V. J.: 1986, Robust estimates for arma models, *Journal of the American Statistical Association* **81**, 155–68.

Cohen, K., S., M., R., S. and D., W.: 1981, Transaction costs, order placement strategy, and existence of the bid-ask spread, *Journal of Political Economy* **89**, 287–305.

Comte, F. and Renault, E.: 2001, Long memory in continuous time stochastic volatility models, *Mathematical Finance* **8**, 291–323.

Corsi, F., Zumbach, G., Müller, U. A. and Dacorogna, M.: 2001, Consistent high-precision volatility from high-frequency data, *Economic Notes* **30**(2), 183–204.

French, K. R. and Roll, R.: 1984, Stock return variances : The arrival of information and the reaction of traders, *Journal of Financial Economics* **17**(1), 5–26.

Glosten, L.: 1994, Is the electronic open limit-order book inevitable?, *Journal of Finance* **49**, 1127–1161.

Glottlieb, G. and Kalay, A.: 1985, Implication of the discreteness of observed stock prices, *Journal of Finance* **40**, 135–153.

Harris, L.: 1990, Estimation of stock price variances and serial covariances from discrete observations, *Journal of Financial and Quantitative Analysis* **25**(3), 291–306.

Hasbrouck, J.: 1993, Assessing the quality of a security market: a new approach to transaction-cost measurement, *Review of Financial Studies* **6**(1), 191–212.

Hasbrouck, J.: 1996, Modelling market microstructure time series, *in* G. S. Maddala and C. R. Rao (eds), *Statistical Methods in Fianance*, Vol. 14 of *Handbook of Statistics*, North Holland, Amsterdam, pp. 647–692.

Hasbrouck, J.: 1999, The dynamichs of discrete bid and ask quotes, *Journal of Finance* **54**, 2109–42.

Hasbrouck, J. and Ho, T.: 1987, Order arrival, quote behavior, and the return-generating process, *Journal of Finance* **42**, 1035–48.

Merton, R. C.: 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* **8**, 323–61.

Parkinson, M.: 1980, The extreme value method for estimating the variance of the rate of return, *Journal of Business* **53**, 61–65.

Roll, R.: 1984, A simple implicit measure of the effective bid-ask spread in an efficient market, *Journal of Finance* **39**, 1127–1139.

Zhou, B.: 1996, High-frequency data and volatility in foreign-exchange rates, *Journal of Business & Economic Statistics* **14**(1), 45–52.