ELSEVIER

Journal of Empirical Finance 9 (2002) 551–562

Journal of EMPIRICAL FINANCE

www.elsevier.com/locate/econbase

Estimating daily volatility in financial markets utilizing intraday data

Bernard Bollen [a,*], Brett Inder [b]

[a] *Department of Accounting and Finance, Monash University, Clayton 3168, Australia*
[b] *Department of Econometrics and Business Statistics, Monash University, Clayton 3168, Australia*

Abstract

This study proposes a new approach to the estimation of daily realised volatility in financial markets from intraday data. Initially, an examination of intraday returns on S&P 500 Index Futures reveals that returns can be characterised by heteroscedasticity and time-varying autocorrelation. After reviewing a number of daily realised volatility estimators cited in the literature, it is concluded that these estimators are based upon a number of restrictive assumptions in regard to the data generating process for intraday returns. We use a weak set of assumptions about the data generating process for intraday returns, including transaction returns, given in den Haan and Levin [den Haan, W.J., Levin, A., 1996. Inferences from parametric and non-parametric covariance matrix estimation procedures, Working paper, NBER, 195.], which allows for heteroscedasticity and time-varying autocorrelation in intraday returns. These assumptions allow the VARHAC estimator to be employed in the estimation of daily realised volatility. An empirical analysis of the VARHAC daily volatility estimator employing intraday transaction returns concludes that this estimator performs well in comparison to other estimators cited in the literature.

Keywords: Daily volatility; Financial markets; Intraday data

1. Introduction

The key focus of this study is the estimation of daily risk (or volatility) in a financial market. This area of study is important because an overwhelming majority of papers dealing with financial risk regard the concept of risk as an unobserved latent

* Corresponding author. Tel.: +61-3-9905-2439; fax: +61-3-9905-5474.
E-mail address: Bernie.bollen@buseco.monash.edu.au (B. Bollen).

0927-5398/02/$ - see front matter © 2002 Elsevier Science B.V. All rights reserved.
PII: S0927-5398(02)00010-5

variable or as a useful theoretical construct to explain the behavior of asset returns. In trying to estimate daily volatility, Andersen et al. (2001a) utilized the 5-min intraday return series to estimate daily volatility and state that "by summing intraday returns sufficiently frequently, the realised volatility can be made arbitrarily close to the underlying integrated volatility". They further state that "for practical purposes we can treat the daily volatility as observed, which enables us to examine its properties directly". By treating daily volatility as an observation, we have a powerful tool to analyse "a broad range of issues in financial economics, both within and beyond the realm of market microstructure" (Andersen and Bollerslev, 1998a, p. 261). In this paper, we extend the literature in regard to estimating daily realised volatility from intraday data, paying close attention to market microstructure effects. In particular, we develop an efficient approach to the estimation of daily volatility, which can utilize all available transactions data.

A number of authors have utilized the notion of realised volatility to estimate daily realised volatility from intraday data. Beckers (1983), Anderson (1995), Parkinson (1980) and Rogers and Satchell (1991) all propose estimators of daily volatility based upon daily high, low, opening and closing prices. Schwert (1990), Hsieh (1991), Andersen and Bollerslev (1998a) and Andersen et al. (2001a) propose efficient unconditional daily volatility estimators based upon the intraday return series. The practical value of the concept of daily realised volatility has also been extensively demonstrated in the literature. Andersen and Bollerslev (1997, 1998a) show how high frequency intraday returns contain valuable information for the measurement of volatility at the daily level. Andersen and Bollerslev (1998b) show how the concept of realised volatility can be utilized to demonstrate how standard ARCH models do in fact provide accurate volatility forecasts. Andersen et al. (2001a,b) utilize the concept of realised volatility to examine the cross-correlation in various asset returns. Moosa and Bollen (2001) utilize realised volatility to test for relationships between volatility and time to maturity in futures markets. Moosa and Bollen (2002) also employ the concept of realised volatility to test for bias in Value at Risk estimates. Andersen et al. (2000) show that daily returns conditioned on realised volatility are nearly Gaussian, a result consistent with the Mixture of Distributions Hypothesis (see Harris, 1986, 1987). Using daily data, Canina and Figlewski (1993), Christensen and Prabhala (1998) and Gwilym and Buckle (1999) examine the information content of an option's implied volatility using realised volatility as the benchmark.

The paper proceeds as follows: In Section 2, we review a number of estimators of daily volatility cited in the literature. It will be seen that these estimators are based upon various restrictive assumptions in regard to the data generating process for intraday returns. In Section 3, we summarise literature that shows how, at the transaction level, returns can be characterised by both time-varying volatility and time-varying autocorrelation which is both seasonal and unpredictable. With this market microstructure analysis in mind, we then propose a new estimator of daily volatility (the VARHAC estimator) based upon quite weak assumptions about how transactions data are generated. In Section 4, we compare the proposed VARHAC estimator with the other estimators using S&P 500 Index Futures and find that it performs very well. Section 5 concludes the discussion.

2. Review of estimators

We begin by assuming the daily return $r = \ln(p_t) - \ln(p_{t-1})$, where p_t is the closing price on day t, and r_t is a drawing from a normal distribution with time-varying volatility:

$$r_t \sim N(0, \sigma_t^2). \tag{1}$$

Our purpose is to estimate the daily volatility parameter σ_t.

The first and most basic estimator is based upon the observation that the expected value of the folded normal distribution is given by $E[|r_t|] = \sigma_t\sqrt{2/\pi}$. This result motivates the Simple volatility estimator:

$$\hat{\sigma}_{(\mathrm{Simple}),t} = \frac{|r_t|}{\sqrt{2/\pi}}. \tag{2}$$

However, the fact that, effectively, only one observation is used to estimate each σ_t suggests that this estimator will be highly inefficient. However, if Eq. (1) holds, then the Simple volatility estimator will be unbiased.

Extreme Value estimators make some use of the information in the intraday data—publicly available information such as the daily high, low, opening and closing prices. Define H_t as the highest price and L_t as the lowest price on day t. Parkinson (1980) develops the PARK daily volatility estimator, based on the assumption that intraday prices follow a Brownian motion process. However, consistent with recent finance literature, it will be assumed that a security's price follows a geometric Brownian motion process. The corresponding PARK daily volatility estimator under this assumption is defined as:

$$\hat{\sigma}^2_{(\mathrm{PARK}),t} = \frac{(\ln(H_t) - \ln(L_t))^2}{4\ln(2)}. \tag{3}$$

Garman and Klass (1980) present an estimator which they claim is more efficient than the PARK. The corresponding estimator assuming a geometric Brownian motion price process shall be referred to as the GK volatility estimator and is defined as:

$$\hat{\sigma}^2_{(\mathrm{GK}),t} = 0.5(\ln(H_t/L_t))^2 - 0.39(\ln(p_t/p_{t-1}))^2. \tag{4}$$

Andersen et al. (2001a,b) propose the Sum of Squared Returns (SSR) daily realised volatility estimator which sums the squares of intraday returns:

$$\hat{\sigma}^2_{(\mathrm{SSR}),t} = \sum_{i=1}^{n_t} r_{i,t}^2, \tag{5}$$

where n_t is the number of intraday returns and is the ith return on trading day t. It should also be noted that the SSR estimator does not require the intraday return series be homoscedastic, but is based on an assumption that intraday returns are uncorrelated. In order to deal with market microstructure effects, which may induce autocorrelation into the intraday return series, Andersen et al. (2001b) estimate an MA(1) model on the 5-min intraday return series and then use the residuals (or filtered returns) in Eq. (5).

3. Intraday returns and the VARHAC estimator

As motivation for our proposed estimator of daily realised volatility, we initially consider some stylised facts in regard to the properties of intraday returns, which leads us to propose a set of assumptions for the data generating process for intraday returns.

Previous studies have already highlighted certain characteristics of intraday returns. Andersen and Bollerslev (1997) study intraday 5-min returns on S&P 500 Index Futures and propose a Fourier flexible functional form as the appropriate methodology required to model intraday seasonality in volatility. They found that the number of transactions per 5-min interval is highly seasonal and strongly related to the 5-min volatility pattern. Roll (1984) shows how prices fluctuating between the bid and ask spread will induce first-order autocorrelation into the intraday return series. Hasbrouck and Ho (1987) propose an ARMA model of the intraday return process. After testing for the autocorrelation structure of a number of securities, they find that "the pattern consists of a large negative autocorrelation at the first lag, followed by positive autocorrelations of decreasing magnitude that are statistically significant. . .through the fifth lag". They propose that positive autocorrelation at lag lengths greater than one may be the result of traders "'working an order', a trader may distribute purchases or sales over time" (Hasbrouck and Ho, 1987, pp. 1039–1040). Admati and Pfleiderer (1988) develop a theoretical model in which concentrated trading at certain times is the result of the strategic trading behavior of both informed and liquidity traders. Consistent with the theoretical findings of Admati and Pfleiderer (1988), beyond predictable periods of intraday volatility, there also exist periods of unpredictable volatility in response to the arrival of new information to the market place.

We can conclude from these studies that each trading day may be characterised by both predictable (seasonal) and unpredictable periods of volatility. Consequently, each trading day may be characterised by a different autocorrelation structure, since autocorrelation in intraday returns appears to be closely related to periods of concentrated trading (as is volatility and the number of transactions per unit of time). The existence of autocorrelation at lag lengths of greater than one would, in theory at least, render all of the estimators of daily volatility reviewed in Section 2 as unsuitable volatility estimators. We will thus propose an approach to estimation of daily volatility, which allows for these properties of intraday returns, and makes use of all available intraday data.

Assume that we observe n_t returns on trading day t (where n_t may be, for example, the number of 5-min returns or transaction returns). Now define $r^*_{i,t} = n_t^{1/2} r_{i,t}$ so that the daily return can be written as $r_t = \sum_{i=1}^{n_t} r_{i,t} = n_t^{-1/2} \sum_{i=1}^{n_t} r^*_{i,t}$. The parameter we wish to estimate is the variance of daily returns $\sigma_t^2 = E[r_t^2]$. This parameter can be expressed as a function of intraday data:

$$\sigma_t^2 = E[r_t^2] = E\left[\left(\sum_{i=1}^{n_t} r_{i,t}\right)^2\right],$$

which can be written as,

$$\sigma_t^2 = \frac{1}{n_t}\sum_{i=1}^{n_t}\sum_{j=1}^{n_t} E[r_{i,t}^* r_{j,t}^*] \tag{6}$$

Apparent from this representation is that σ_t^2 represents the spectral density of $r_{i,t}^*$ at frequency zero. We thus follow den Haan and Levin (1996), and propose the use of the VARHAC estimator for this parameter.

Let us work through the connection with spectral density estimation more closely, and also focus on the required assumptions for this estimator to be appropriate. Adopting their notation, one of the concerns of den Haan and Levin (1996) is to estimate parameters underlying the data generating mechanism for a process $\{V_t\}_{t=-\infty}^{\infty}$. In particular, they are interested in estimating

$$S_T = \frac{1}{T}\sum_{s=1}^{T}\sum_{t=1}^{T} E[V_s V_t], \tag{7}$$

(see den Haan and Levin, 1996, Eq. (2.3)). Note that for our purposes, we simplify their notation and let V_t be a scalar. By comparing Eqs. (6) and (7), it is easy to see that if we let $V_t = r_{i,t}^*$ and $T = n_t$ then S_T is equal to σ_t^2. Consequently, we can appeal to the literature on spectral density estimation to guide us in our estimation of σ_t^2.

In this regard, den Haan and Levin (1996) argue strongly for the use of the VARHAC estimator, which we will outline below. Before this estimator can be adopted, however, it is necessary to verify whether $r_{i,t}^*$ satisfies the assumptions required for valid use of the estimator. We refer to den Haan and Levin (1996) who define conditions under which the VARHAC estimator can be implemented.

Conditions A*: Let $\{V_t\}_{t=-\infty}^{\infty}$ be a mean-zero sequence of random N-vectors, which satisfy the following conditions:

(a) $\sup_{t\geq 0} E[V_t V_t'] < +\infty$;
(b) $\Gamma(0) = \lim_{T\to\infty} 1/T \sum_{t=1}^{T} E[V_t V_t']$ is positive definite, and;
(c) $\sum_{j=1}^{\infty} \sup_{t\geq 1} |E[V_t V_{t+j}']|_{\infty} < +\infty$

(den Haan and Levin, 1996, Section 3.2). These conditions essentially allow V_t to possess a degree of heterogeneity and autocorrelation, and they provide sufficient conditions for valid estimation of S_T using the VARHAC estimator. Since we wish to make our assumptions in regard to intraday returns, we will rewrite conditions A* with reference to $r_{i,t}$, and then explore its implications.

Conditions A*: We assume $r_{i,t}$ is a mean-zero sequence of scalars satisfying:

(a) $\sup_{i\geq 0} n_t E[r_{i,t}^2] < +\infty$;
(b) $\sigma_t^2 = \lim_{n_t\to\infty} \sum_{t=1}^{n_t} E[r_{i,t}^2] > 0$, and;
(c) $\sum_{j=1}^{\infty} \sup_{i\geq 1} |E[r_{i,t} r_{i+j,t}]|_{\infty} < +\infty$.

Condition (a) requires that intraday returns have finite variance, and condition (b) rules out the possibility of too many values having zero variance. Condition (c) controls the degree of dependence across observations; it implies that the $r_{i,t}$'s are asymptotically independent, and that the sums of autocovariances are bounded. Note that no further restrictions are placed on $r_{i,t}$, other than those stated in Eq. (8). These conditions allow for a wide range of possible heteroscedasticity and autocorrelation in intraday returns. The autocovariances are not restricted to be constant across time—the only constraint is the absolute summable restriction in condition (c).

The key advantage of this weak set of conditions on $r_{i,t}$ is that they are consistent with the accepted wisdom about the time series behavior of intraday returns, as discussed earlier in this section. Such conditions also allow $r_{i,t}$ to be unequally spaced transaction returns. Unequally spaced data produces another potential source of heteroscedasticity, and also causes the autocovariances to vary across observations. Neither of these implications leads to a violation of assumptions stated in Eq. (8).

Following den Haan and Levin (1998), we can apply the VARHAC estimator to obtain a consistent estimator for σ_t^2, which uses all of the intraday information and is thus efficient. The steps in this estimator are as follows:

Step 1 (Lag order selection for each day). Consider a series of observations on transaction returns $\{r_{1,t}, r_{2,t}, \ldots, r_{n_t,t}\}$ for a specific trading day t, and furthermore assume, we have data for S such trading days (i.e. $1 \leq t \leq S$). For each of the S days of data, the following model is estimated by least squares:

$$r_{i,t} = \sum_{k=1}^{K} \hat{\alpha}_{k,t}(K) r_{i-k,t} + \hat{e}_{i,t}(K) \qquad \text{for } i = K+1, \ldots, n_t, \tag{8}$$

for each possible lag order $K = 1, \ldots, \bar{K}$. For $K = 0$, we set $\hat{e}_{i,t}(0) = r_{i,t}$. Next, we calculate the Schwarz BIC:

$$\mathrm{BIC}(K,t) = \ln\left(\frac{\sum_{i=k+1}^{n_t} \hat{e}_{i,t}^2(K)}{n_t}\right) + \frac{\ln(n_t)K}{n_t}. \tag{9}$$

For each day, the optimal lag order K_t is chosen as the value of K which minimises BIC(K,t).

Step 2 (Calculate the estimate of σ_t^2). For each day, the selected lag order K_t, and estimates of $\hat{\alpha}_{k,t}(K_t)$ and $\hat{e}_{i,t}(K_t)$ from Eq. (8) can be used to compute the estimate:

$$\hat{\sigma}_t^2 = \frac{n_t \hat{\sigma}^2}{\left[1 - \sum_{k=1}^{K_t} \hat{\alpha}_{k,t}(K_t)\right]^2}, \qquad \text{where } \hat{\sigma}^2 = \sum_{i=K_t+1}^{n_t} \hat{e}_{i,t}(K_t)^2 / n_t. \tag{10}$$

The estimator defined in Eq. (10) will be referred to as the VARHAC volatility estimator $\hat{\sigma}^2_{(\mathrm{VARHAC}),t}$.

4. An analysis of S&P 500 Index Futures realised daily volatility

In this section, we apply the VARHAC daily realised volatility estimator to S&P 500 Index Futures transaction data. Intraday data on S&P 500 Index Futures from the Chicago Mercantile Exchange (CME) for the period of January 1993 to December 1995 were obtained. Trading times for the CME are from 8.30 AM to 3.30 PM. Some 1,932,585 trades were recorded over this period, however, these recorded trades do not necessarily represent a record of each transaction. Transactions with zero price changes are frequently not recorded. A heteroscedasticity-consistent Variance Ratio Test (see Lo and MacKinlay, 1988) on both the 5-min and transaction returns series was conducted on each of the 755 trading days in the data set.[1] Table 1 shows the percentage of tests where the null hypothesis of no autocorrelation in the intraday return series is rejected for values of $Q = 2,\ldots,6$, where Q specifies the number of periods (number of transactions or 5-min intervals) over which the Variance Ratio Test is conducted.

It can be seen from Table 1 that the Variance Ratio Test provides strong evidence of autocorrelation at lag lengths greater than one in the transaction return series, but little evidence of autocorrelation in the 5-min return series. Further evidence of autocorrelation in the transaction return series is provided when AR($\cdot$) models are estimated on the transaction return series of each trading day to calculate VARHAC volatility estimates. AR(0), AR(1),. . .,AR(8) models were estimated on each trading day.[2] The BIC chose an AR(1), AR(2) or AR(3) model in 93% of trading days. On only 11 of the 755 trading days was an AR(0) model chosen by the BIC.

Next, the daily volatilities of the 755 trading days are estimated from intraday data using the Simple, PARK and GK daily volatility estimators reviewed in Section 2. The SSR daily volatility estimator was also calculated on each trading day using both transaction and 5-min data after intraday returns had been filtered using an estimated MA(1) process. Fig. 1 exhibits these volatility estimates over the period January 93 to December 95, as well as the VARHAC volatility estimates.

A simple visual inspection of Fig. 1 reveals that the Simple volatility estimates are far more variable than the other estimates, probably due to the inefficiency of this estimator. All estimates appear to move in a similar way through time, and appear to cluster in distinct high and low periods at the same time. The PARK and GK estimators appear, however, to display a greater degree of noise than the SSR and VARHAC volatility estimators. This is probably because the PARK and GK volatility estimators utilize very little intraday information, and consequently are also inefficient.

Daily volatility estimates appear to exhibit outliers on both 31/3/94 and 4/4/94, with estimates of 1.53% and 1.43%, respectively. The SSR (Transactions), SSR (5 min), and VARHAC volatility estimators all show these as unusually high volatility days, even though the daily return on these days was not unusually high. On 5/4/94,

[1] The Variance Ratio Test is designed to detect autocorrelation in returns.

[2] One variation on the procedure involved the treatment of a constant term. If mean return is zero, we would not expect a constant term in the AR models. Each AR model was fitted with and without a constant term, and the preferred model chosen by the BIC. It turned out that for none of the 755 trading days was a model selected which included a constant. The simplifying assumption that daily returns are mean zero thus seems reasonable.

Table 1
Variance Ratio Test results

Q	2	3	4	5	6
Percentage of times, H_0 is rejected in the transaction return series	65.03	75.89	80.00	80.26	81.06
Percentage of times, H_0 is rejected in the 5-min return series	5.70	4.64	3.31	3.05	2.91

the headline of the Australian Financial Review reads "Tokyo's wild day sparks new fears". The paper then goes on to report "intense volatility infected the Tokyo stock market yesterday as fears of a fresh post-Easter sell-off on Wall Street sent the main Nikkei share index on a roller coaster ride".

It is difficult to evaluate definitively which of the estimators is 'best' given that the true volatility is unobservable. However, benchmark criteria can be constructed to assess the performance of the different estimators.

Criterion 1 (Bias test (1)—comparison with Simple volatility estimates). Recall that the Simple volatility estimator is an unbiased estimator of daily volatility if daily returns are uncorrelated. It follows that the mean of the Simple volatility estimates over the 755 trading days will be an unbiased estimator of the mean of the true daily volatilities σ_t. Consequently, if another estimator yields a mean of the estimates across days which is significantly different to the mean using the Simple volatility estimator, then clearly this other estimator displays systematic bias. Such a hypothesis is tested by comparing $\overline{X}_{(\mathrm{Simple})} = 1/S\sum\hat{\sigma}_{(\mathrm{Simple}),t}$ and $\overline{X}_{(\cdot)} = 1/S\sum\hat{\sigma}_{(\cdot),t}$ using a standard test for the difference of two means, where $\sigma_{(\cdot),t}$ is the estimate from one of the daily realised volatility estimators.

Criterion 2 (Normality in standardised returns). The model for daily returns could be written as $r_t=\sigma_t z_t$, where $z_t \sim iid\ N(0,1)$. The normality of z_t is justified by appealing to a Central Limit Theorem for mixing processes to argue that if transaction returns are strong mixing, then returns over a reasonable aggregation (such as daily for heavily traded assets like S&P 500 Index Futures) should tend toward normality. Andersen et al. (2000) discuss the significance of the normality of scaled returns in far greater detail. We define the standardised return for a given volatility estimator as:

$$z_{(\cdot),t} = \frac{r_t}{\hat{\sigma}_{(\cdot),t}}. \tag{11}$$

If σ_t is adequately estimated, $z_{(\cdot),t}$ should be normally distributed with a variance of unity.[3] We can test for normality using the Jarque–Bera test on the S observations for $z_{(\cdot),t}$ for each chosen estimator.

Criterion 3 (Bias test (2)—variance of standardised returns). From Eq. (11), it is clear that any systematic bias inherent in a daily volatility estimator would yield standardised returns

[3] Again, we assume $E[r_t]=0$ on any given day.

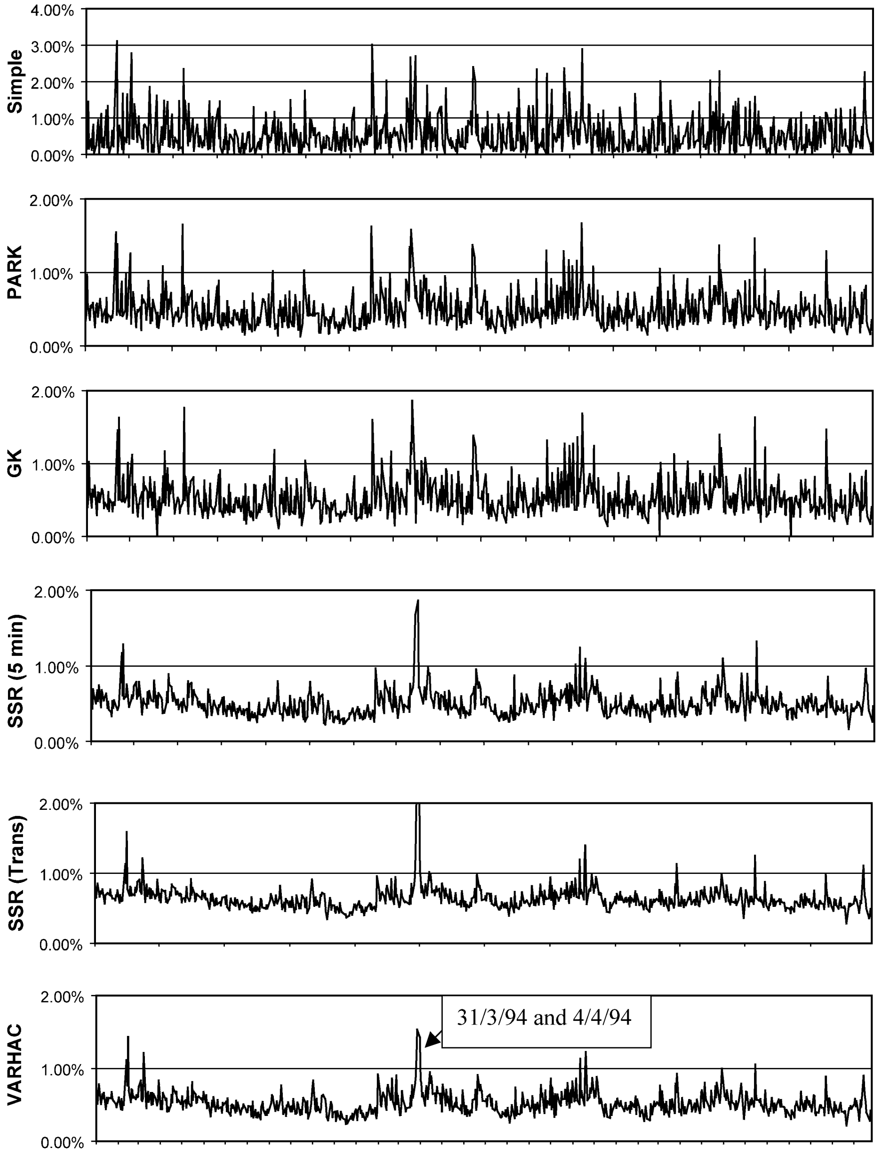

Fig. 1. Daily volatility estimates of S&P 500 Index Futures (1993–1995).

with a variance different from one. A test is performed on the variance of standardised returns for each of the volatility estimators. The test statistic is given by:

$$\frac{\overline{Y}_{(\cdot)} - 1}{\hat{\sigma}_{Y_{(\cdot)}} / \sqrt{S}},$$

where $Y_{(\cdot)}=(z_{(\cdot),t}-\bar{z}_{(\cdot)})^2$, $\bar{Y}_{(\cdot)}$ and $\hat{\sigma}_{Y(\cdot)}$ are the mean and sample standard deviation of $Y_{(\cdot)}$. This statistic should follow a standard normal distribution if the true variance is one.

Table 2 contains results for S&P 500 Index Futures data relevant to each of the three criteria. The t-statistics in row three of the results reveal that H_0 cannot be rejected for both the VARHAC volatility estimator based upon transactions data, and the GK daily volatility estimator. Evidence points clearly to bias in the other estimators.

The Jarque–Bera normality test conducted on standardised returns suggests that the PARK and GK estimators clearly do not yield normally distributed standardised returns. Normality cannot be rejected for the SSR estimator based upon transaction and 5-min returns and for the proposed VARHAC estimator.

The final row of Table 2 indicates that we cannot reject the hypothesis that standardised returns for the PARK, GK and VARHAC volatility estimators have a variance of one. However, the SSR estimator based upon transaction data had a variance significantly different from one. This result indicates that the SSR estimator based upon transaction data is a systematically biased estimator of daily volatility. Such a conclusion can be explained by noting that the MA(1) filtering process does not adequately adjust for higher order autocorrelation in the transaction return series. When this test is conducted on the SSR estimator based upon 5-min data, the degree of bias is considerably reduced.

To summarise, the only volatility estimator which passed both the bias tests (Criteria 1 and 3) and the conditional normality test (Criterion 2) was the VARHAC volatility estimator. We believe that this estimator performed well in these tests because it is an efficient and robust estimator making full use of intraday returns data.

Table 2
Performance of different volatility estimators

	Volatility estimator					
	Simple	PARK	GK	SSR (Trans)	SSR (5 min)	VARHAC (Trans)
Volatility estimates						
Mean	0.005469	0.004920	0.005302	0.006395	0.005059	0.005234
Standard deviation	0.004978	0.002351	0.002517	0.001569	0.001689	0.001552
t-Test for difference between means		3.9965 *	1.0639	− 5.5649 *	2.4834 *	1.4242
Standardised returns						
Mean		0.1424	0.1575	0.1071	0.1609	0.1437
Standard deviation		1.0006	1.1339	0.8373	1.0575	1.0171
Coefficient of skewness		0.1027	1.7004	0.0057	0.1390	0.0682
Coefficient of excess kurtosis		− 0.4507	15.5762	0.1089	− 0.1502	− 0.0979
Jarque–Bera test for normality		7.6862 *	7964 *	0.3754	3.1298	0.8828
$z_{(\cdot)}$ Test for unit variance		0.0040	− 1.4515	8.1059 *	− 2.1128 *	− 0.6389

* Indicates significance at the 5% level.

5. Conclusion

The volatility estimator we develop arises directly from a study of stylised facts in regard to market microstructure, in particular by time-varying autocorrelation and heteroscedasticity in intraday returns. From this analysis, an estimator of daily volatility has been proposed, which we refer to as the VARHAC volatility estimator. The VARHAC volatility estimator was applied to S&P 500 Index Futures. Three criteria were proposed to evaluate a daily volatility estimator. The VARHAC volatility estimator performed well against all three criteria, while none of the competing estimators proposed in the literature were able to satisfy all criteria. The VARHAC estimator is attractive as it imposes few assumptions on the nature of the data generating process for intraday returns, and makes use of all information in all of the available data.

Throughout this paper, we have referred to the VARHAC estimator as a volatility estimator. However, unlike all the other estimators considered, all transactions data are utilized; that is, we utilize the whole population of intraday returns rather than a sample of intraday returns. From a statistical perspective, one could argue that we have not estimated a population parameter, but rather have calculated a population parameter. On a philosophical note, if we accept that VARHAC 'estimates' are in fact the population parameter of interest, then one can regard VARHAC estimates based upon transaction returns as a 'true' volatility observation. Further research could endeavour to uncover how the VARHAC volatility estimator is able to address the "broad range of issues in financial economics, both within and beyond the realm of market microstructure" as proposed by Andersen and Bollerslev (1998a, p. 261).

Acknowledgements

The authors wish to thank Paul Kofman, Keith McLaren, Maxwell King, Neil Shephard and participants at the 1998 Econometric Society Australasian Meeting for their many invaluable comments. All errors and omissions are however the sole responsibility of the authors.

References

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

Andersen, T., Bollerslev, T., 1997. Intraday periodicity and volatility persistence in financial markets. Journal of Empirical Finance 4, 115–158.

Andersen, T., Bollerslev, T., 1998a. Deutsche mark-dollar volatility: intraday activity patterns, macroeconomic announcements, and longer run dependencies. Journal of Finance 53, 219–265.

Andersen, T., Bollerslev, T., 1998b. Answering the sceptics: yes, standard volatility models do provide accurate forecasts. International Economic Review 39, 885–905.

Andersen, T., Bollerslev, T., Diebold, F., Labys, P., 2000. Exchange rate returns standardised by realised volatility are (nearly) Gaussian, NBER Working Paper 7488.

Andersen, T., Bollerslev, T., Diebold, F., Labys, P., 2001a. The distribution of exchange rate volatility. Journal of the American Statistical Association 96 (453), 42–55.

Andersen, T., Bollerslev, T., Diebold, F., Ebens, H., 2001b. The distribution of realised stock return volatility. Journal of Financial Economics 61, 43–76.

Anderson, G., 1995. Variance estimation in the presence of drift and non-trading, Working paper, School of Economics and Commercial Law, University of Goteborg.

Beckers, S., 1983. Variances of security prices based on high, low and closing prices. Journal of Business 56, 97–112.

Canina, L., Figlewski, S., 1993. The information content of implied volatility. The Review of Financial Studies 6, 659–681.

Christensen, B., Prabhala, N., 1998. The relationship between implied and realized volatility. Journal of Financial Economics 50, 125–150.

den Haan, W.J., Levin, A., 1996. Inferences from parametric and non-parametric covariance matrix estimation procedures, Working paper, NBER, 195.

den Haan, W.J., Levin, A., 1998. A practitioner's guide to robust covariance matrix estimation, mimeo, Department of Economics, University of California at San Diego.

Garman, M., Klass, M., 1980. On the estimation of security price volatilities from historical data. Journal of Business 53, 67–78.

Gwilym, O., Buckle, M., 1999. Volatility forecasting in the framework of the option expiry cycle. The European Journal of Finance 5, 73–94.

Harris, L., 1986. Cross-security tests of the mixture of distributions hypothesis. Journal of Financial and Quantitative Analysis 21, 39–46.

Harris, L., 1987. Transaction data test of the mixture of distributions hypothesis. Journal of Financial and Quantitative Analysis 22, 127–141.

Hasbrouck, J., Ho, T., 1987. Order arrival, quote behavior, and the return-generating process. Journal of Finance 42, 1035–1048.

Hsieh, D., 1991. Chaos and non-linear dynamics: application to financial markets. Journal of Finance 5, 1839–1877.

Lo, A., MacKinlay, A., 1988. Stock market prices do not follow random walks: evidence from a simple specification test. Review of Financial Studies 1 (1), 41–66.

Moosa, I., Bollen, B., 2001. Is there a maturity effect in the S&P 500 Futures contract price. Applied Economic Letters 8, 693–695.

Moosa, I., Bollen, B., 2002. A benchmark for measuring bias in daily value at risk. International Review of Financial Analysis 11, 85–100.

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

Rogers, G., Satchell, S., 1991. Estimating variance from high, low and closing prices. The Annals of Applied Probability 1, 504–512.

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

Schwert, G., 1990. Stock market volatility. Financial Analysts Journal. May–June.