53

Applied Financial Economics, 1994, **4**, 241–247

Estimating the volatility of stock prices: a comparison of methods that use high and low prices

L. C. G. ROGERS, S. E. SATCHELL[‡] and Y. YOON[§]

School of Mathematical Science, Queen Mary and Westfield College, University of London, [‡]Trinity College, University of Cambridge and Birkbeck College, University of London, [§]Department of Applied Economics, University of Cambridge, Sidgwick Avenue, Cambridge CB3 9DE, UK

The volatility of stock prices has played an important role in the financial literature. Different methods of estimating the volatility are suggested and applied to British financial assets. Since we cannot observe the real volatility, we investigate the efficiency of the methods through simulation. The question of which estimator to use rather depends on the distributional assumption of returns. If it is log-normal, methods based on high/low prices are preferred. Furthermore, if there is drift in the data, then one may wish to use a procedure devised by Rogers and Satchell. If the drift is varying with time, the Rogers and Satchell's method is clearly superior.

I. INTRODUCTION

Since the recognition that the Black and Scholes option price depends upon only one unobservable parameter, the volatility of logarithmic stock prices, considerable attention has been paid by financial economists to efficient volatility estimation. Efficiency is of key importance, because of the recognition that volatility may change over long periods of time; a highly efficient procedure will allow researchers to estimate volatility with a small number of observations. Methods based on using opening, closing, high and low prices, as published in the financial literature, have arisen in response to this need, and readers should consult Parkinson (1980), Garman and Klass (1980), and Rogers and Satchell (1991) for details.

It is the purpose of this paper to describe a procedure put forward by Rogers and Satchell (1991), explain how to use it with stock price data and assess its merits relative to existing methods in the presence of discrete observations. A case of some interest is when volatility is fitted over a long data period during which expected returns are changing; provided the volatility remains constant, the Rogers and Satchell method can estimate it even though the expected return is non-constant. This is important because there is some evidence of changing expected returns on the assets over time; time-varying expected return is often attributed to risk premia or trends in prices[1]. The relation between the expected returns and the volatility can be also found in French, Schwert and Stambaugh (1987). We investigate, via simulation, how the different methods respond to various parametric changes. We present a description of the methods and some results in Section II and our simulation results in Section III. Conclusions follow in Section IV.

II. ESTIMATION OF VOLATILITY

We shall make the standard assumption in finance, namely that $s(t)$, the price of the asset at time t, is generated by the process of

$$ds(t) = \alpha s(t)dt + \sigma s(t)dW(t) \tag{1}$$

where α and σ are assumed constants for the moment, and $W(t)$ is a standard Brownian motion (BM). The use of

[1]The changing expected return of stock prices is related to mean reversion in stock prices, which implies the presence of autocorrelation in returns. See Poterba and Summers (1988), Fama and French (1988), Cecchetti, Lam and Mark (1990), Schwert and Seguin (1990).

0960–3107

Equation 1 is based on the hypothesis that the continuous geometric Brownian motion is followed during periods between transactions and during periods of exchange closure, even though prices cannot be observed in such intervals. It is well known that Equation 1 has the solution

$$s(t)=s(0)\exp\left[\left(\alpha-\frac{1}{2}\sigma^2\right)t+\sigma(W(t)-W(0))\right] \tag{2}$$

and that in the logarithmic form,

$$\ln\left(\frac{s(t)}{s(t-1)}\right)=\left(\alpha-\frac{1}{2}\sigma^2\right)+\sigma(W(t)-W(t-1)). \tag{3}$$

We see from Equation 3 that given $\ln(s(t-1))$, $\ln(s(t))$ follows a random walk with drift and that errors are i.i.d. $N(0,\sigma^2)$, i.e.,

$$\ln s(t)=\ln s(t-1)+\left(\alpha-\frac{1}{2}\sigma^2\right)+\mathscr{E}(t)$$

$$\text{where}\quad \mathscr{E}(t)=\sigma(W(t)-W(t-1))\sim N(0,\sigma^2). \tag{4}$$

Previous studies have considered the problem of estimating σ^2 from daily data, specifically, the opening, closing, high and low prices. In some cases, the daily volume of trade or the number of transactions have also been used. However, in the latter case Equation 3 needs to be modified in some way to relate it to these variables, which is not our concern in this paper. The models of the joint dynamics for stock returns and volume are abundant in the recent finance literature.

One of the major issues has been the extent to which the behaviour of different estimators has been influenced by the fact that the high/low prices are under/over-estimates of the true high/low daily values; these are under/over-estimates because they are calculated from discretely observed samples, not the idealized continuous sample path for the day, which is typically unobservable. Another issue is the possible behaviour of $s(t)$ when the market is closed[2]; for simplicity, in this paper we assume that the process generating $s(t)$ stops when the market is closed and starts to work again when it is open. However, Garman and Klass (1980) relax this assumption in their paper.

First we define the notation for our method of estimating volatility. Let $h(t)$ and $l(t)$ be the high and low prices from day t. We shall divide $h(t)$, $l(t)$ and $s(t)$, the closing price, by the opening price $o(t)$ and define

$$\begin{aligned} H(t)&=\ln(h(t)/o(t))\\ S(t)&=\ln(s(t)/o(t))\\ L(t)&=\ln(l(t)/o(t)). \end{aligned} \tag{5}$$

This makes $H(t)$, $S(t)$ and $L(t)$ independent across days, since they are functions of non-overlapping increments of BM.

We now assume that we have daily observations $s(t)$, $t=1,2,\dots,n$, and $o(t)=s(t-1)$, then the naive estimator of variance in a driftless world ($\alpha=0$) is just

$$\hat{\sigma}_0^2=\frac{1}{n}\sum_{t=1}^{n}(S(t))^2. \tag{6}$$

A mean-adjusted unbiased alternative for the case when α is not equal to zero is obviously

$$\hat{\sigma}_1^2=\frac{1}{(n-1)}\sum_{t=1}^{n}(S(t))^2-\frac{1}{n(n-1)}(\ln(s(n))-\ln(s(0)))^2. \tag{7}$$

Garman and Klass (1980) derive an estimator that has a minimum-variance among the class of unbiased estimators which are quadratic in $H(t)$, $S(t)$ and $L(t)$. This is done under the assumption that $\alpha=0$ in Equation 1. Their practical estimator is the following

$$\hat{\sigma}_{GK}^2=0.511(H-L)^2-0.019[S(H+L)-2HL]-0.383S^2. \tag{8}$$

Ball and Torous (1984) derive a maximum likelihood estimator (MLE) which will be asymptotically efficient. Since it has no closed form solution and is again only relevant for the case $\alpha=0$, we shall not examine it here. We note that an MLE when α is not equal to zero will be extremely complicated.

For the above and other reasons, Rogers and Satchell (1991) suggest an estimator that has the attractive property that it is unbiased whatever the value of α is; this is given by

$$\hat{\sigma}_{RS}^2=H(H-S)+L(L-S). \tag{9}$$

Note that $\hat{\sigma}_{RS}^2$ is a member of the unbiased quadratic class of Garman and Klass, so $\hat{\sigma}_{GK}^2$ will outperform it if $\alpha=0$. However, as α increases, $\hat{\sigma}_{GK}^2$ goes badly astray, while $\hat{\sigma}_{RS}^2$ remains close to the true value, see Rogers and Satchell (1991). Having defined our basic estimators, we now address the problem of how to adjust our estimators when the path can only be monitored at discrete intervals. Since $H(t)$ and $L(t)$ are not true values, it turns out that all our estimators, except Equation 6 for $\alpha=0$ and Equation 7 for $\alpha\neq 0$, are biased.

The reasons for the observed bias are that the observed highs and lows from the market or simulation are less in absolute magnitude than the highs and lows of the idealized continuous process and that the time period over which it is estimated is shortened to the span of exchange trading hours excluding closing hours. These technical limitations produce a downward bias. Adjustments have been proposed for the estimators which perform well in simulations. The amount by which the random walk simulation under/over-estimates the real price will depend on the fineness of the

[2]French and Roll (1986) proved that asset prices are much more volatile during exchange trading hours than during non-trading hours due to the extent of the arrival of available information. This suggests that volatility is closely related to information processing.

interval chosen; the more steps taken by a random walk in the time interval, the better the approximation to the real price we shall obtain. We shall denote by V the number of steps taken by the random walk during the time interval, the proposed estimator $\hat{\sigma}_{ARS}$ of this adjustment is the positive root of the equation

$$\hat{\sigma}^2_{ARS}=2b\hat{\sigma}^2_{ARS}(1/V)+2(H-L)a\hat{\sigma}_{ARS}\sqrt{1/V} + H(H-S)+L(L-S) \tag{10}$$

where

$$a\equiv\sqrt{2\pi}\left[\frac{1}{4}-\frac{\sqrt{2}-1}{6}\right]\approx 0.4536,$$

$$b\equiv\left(1+\frac{3\pi}{4}\right)\Big/12\approx 0.2797.$$

For a proof see Rogers and Satchell (1991), we note that Equation 10 always has a solution[3]. Adjustment in the same manner to Equation 10 for the Garman and Klass estimator produces the estimator $\hat{\sigma}^2_{AGK}$, where $\hat{\sigma}_{AGK}$ solves

$$\begin{aligned}\hat{\sigma}^2_{AGK}=0.511((H-L)^2&+4(H-L)a\hat{\sigma}_{AGK}\sqrt{1/V}\\ &+2\hat{\sigma}^2_{AGK}(1/V)(b+a^2))-0.019S(H+L)\\ &+0.038(HL-(H-L)a\hat{\sigma}_{AGK}\sqrt{1/V}\\ &-a^2\hat{\sigma}^2_{AGK}(1/V))-0.383S^2.\end{aligned} \tag{11}$$

The practical question arises as to how one would use these formulae given daily data and the trading volume and/or the number of transaction. Since volume is typically available whilst the number of transactions is not, we can use the former as a proxy for the latter. This is clearly unsatisfactory and Becker (1983) has attempted to estimate a variant of the Garman and Klass estimator using the average number of daily transactions for the New York Stock Exchange, but with no obvious gains as a consequence. Since we obtained daily transactions data from the London Stock Exchange for five shares chosen from the FT100 index for the period of 1 November 1989 to 11 October 1991, we will use the number of transactions as a proxy for V in our estimation.

The procedure we propose is as follows. Given our different estimation procedures, Equations 6–11, which are based on daily data, our procedure is to calculate our daily variance estimator $\hat{\sigma}^2_{lt}$, $t=1,\ldots,n$ where l refers to the method, i.e. l=naive(0), mean-adjusted(1), GK, AGK, RS, ARS. Our estimators of volatility will be based on the following equation

$$\hat{\sigma}^2_l=\sum_{t=1}^{n}\hat{\sigma}^2_{lt}/n. \tag{12}$$

To give the reader some idea of the magnitudes involved we estimate the variance of the daily logarithmic differences of five British companies for which we have transactions data. The companies are ASDA Group, British Telecommunication, Grand Metropolitan, ICI and Thorn-EMI. We present in Table 1 the values of the estimators $\hat{\sigma}^2_0$, $\hat{\sigma}^2_{GK}$, $\hat{\sigma}^2_{AGK}$, $\hat{\sigma}^2_{RS}$, $\hat{\sigma}^2_{ARS}$; we do not report $\hat{\sigma}^2_1$ since it is virtually identical to $\hat{\sigma}^2_0$. The data period is chosen from 1 November 1989 to 11 October 1991 due to the availability of transactions data.

Table 1 shows the difference among the estimates from the different methods even though the magnitude is quite small. The naive estimator is typically between 20% and 100% larger than the extreme-value estimators in columns 3 to 6, this is true in all cases, and is a phenomenon that requires explanation. In fact to strengthen this point, we note that the four high/low estimators are very similar in value, each lies in the other's 95% confidence interval in most cases. The naive estimator does not lie in the confidence intervals of any of the high/low estimators with the exception of ICI. Note that the adjusted estimators are always larger than the unadjusted; this follows from their definition. This indicates the substantial effect of transactions on the volatility. In general, most estimators are significantly correlated with the number of transactions even if not reported here. The only exception to this is British Telecommunication. It may be that the privatization of 1984, which led to a very broad ownership of BT, would have influenced the transactions costs for BT. These higher costs would mean less portfolio changes (transactions) to price changes (volatility), hence the non-significant correlation between the two.

Table 1. *Estimation of Volatility (1 November 1989–11 October 1991)*

	$\hat{\sigma}^2_0$	$\hat{\sigma}^2_{GK}$	$\hat{\sigma}^2_{AGK}$	$\hat{\sigma}^2_{RS}$	$\hat{\sigma}^2_{ARS}$
ASDA	7.146	3.681	4.063	3.552	3.887
	(2.154)	(0.628)	(0.660)	(0.604)	(0.629)
BT	1.447	0.871	0.933	0.831	0.888
	(0.105)	(0.047)	(0.050)	(0.051)	(0.054)
GMET	1.714	1.070	1.185	1.070	1.178
	(0.124)	(0.125)	(0.136)	(0.177)	(0.189)
ICI	1.268	1.059	1.141	1.132	1.212
	(0.125)	(0.099)	(0.104)	(0.125)	(0.132)
THN	1.421	0.704	0.861	0.654	0.795
	(0.120)	(0.047)	(0.055)	(0.049)	(0.056)

*The first number is calculated by Equation 12 and the number in parenthesis is its standard deviation. The number of observations used is 492. All numbers are scaled up by 10E+4.

[3]It is easy to show that the Equation 10 has real solutions, and hence one positive solution. Define

$$Q(\sigma)=(1-2b/V)\sigma^2-2(H-L)a\sigma\sqrt{1/V}-H(H-S)-L(L-S)$$

Then, for $V>2b$, $Q(\sigma)\to\infty$ as $|\sigma|\to\infty$ and $Q(0)<0$. This guarantees one positive root.

If we knew the explicit relationship between our estimators of σ^2 and V, we could introduce appropriate weights in Equation 12. But we do not assume here that daily volatility of stock price conditional on the number of transactions is a linear function of the number of transactions. Rather we assume that daily volatility is constant given the sample period by the definition of Equation 1, and that we only observe a finite number of price changes a day, whenever there is a transaction. If our assumption is true, we might expect a positive correlation between the extreme-value estimators and the number of transactions but not between $\hat{\sigma}_0^2$ and the number of transactions. However, the significant correlations suggest that $(\ln s(t) - \ln s(t-1))^2$ is positively correlated with the number of transactions. Thus the distributional assumptions in Equations 1 to 4 are not correct. The distribution of $s(t)$ conditional on V may be normal and there is a huge literature on this topic; we have already noted its importance earlier and we shall not pursue such a possibility in this paper.

Another potential candidate for the measurement of volatility is the implicit volatility. Even though we have no idea about the real volatility, in practice implicit volatility deduced from the Black and Scholes option pricing model is widely used. The remarkable feature of the Black and Scholes option formula is that it is unaffected by changes in the α function in Equation 1 of $s(t)$ for a fixed rate of interest, see Merton (1980), hence an implicit volatility estimator deduced from the options price may be useful for this model. However, the fact that the implicit volatility takes different values for different exercise prices and maturities makes it unreliable as a benchmark. What is clear from the empirical data is the large variation in volatility estimators across methods, and the significant correlation between high/low estimators and the number of transactions. These findings motivate a simulated experiment to try to understand these problems more deeply.

III. SIMULATION

In this section we present a generalization mentioned earlier in that we allow the drift (α) to vary from one day to the next, independently of the random walk. The drift may have a certain time-varying structure. The estimator $\hat{\sigma}_{RS}^2$ given by Equation 12 remains unbiased in the presence of changing drift (α) of the form described above. The adjusted estimator $\hat{\sigma}_{ARS}^2$ cannot be expected to be exactly unbiased so that changing drift may weaken its performance, a question we investigate below. We now give details of our simulation experiment. We assume our data are generated by Equation 1. We choose $\alpha = 0.00068$[4] and $\sigma = 0.013$ based on the average value of our empirical data during the estimation period, here we normalize $\sigma = 1.0$ and $\alpha = 0.052$ for the sake of simplicity. For each day given the number of transactions V_i, we generate V_i random normals with distribution, $N((\alpha - \sigma^2/2)/V_i, \sigma^2/V_i)$. From these we obtain prices. We set $n = 30$, for these 30 days we calculate our various estimators. Perhaps the different values of n can be used for longer periods to investigate the efficiency over the estimation period. We shall denote a typical one by $\hat{\sigma}_{lj}^2$[5], where l is the estimator type and j refers to the simulation numbers. Then we repeat 1000 simulations so that we have $\{\hat{\sigma}_{lj}^2\}$, $j = 1, \ldots, 1000$. As a preliminary test, we carried out the simple Monte Carlo simulation with a constant drift and a fixed number of transactions. Even though we do not report the results here, we find that with a small number of transactions the naive estimators perform better than the others. However, as we increase the number of transactions, the other four estimators are becoming more accurate. We also observe that the naive estimators are more volatile, and the high/low estimators lead to substantial gains in efficiency as is well known. As V increases, the bias in the high/low estimators is reduced, hence they are getting larger, explaining the positive correlation with V. One interesting finding is that the GK and AGK estimators are marginally better than the RS and ARS estimators, respectively for high V.

Now we carry out our experiment by letting α vary from day to day. The drift is generated by a sample from a normal distribution with mean equal to 0.052 and variance σ_α^2, and calculate (once for each value of σ_α^2) a sequence of n values of α_t. The choice of a normal distribution is arbitrary. One generalization would be to introduce some autoregressive process to reflect the mean reversion in stock prices. We let $\sigma_\alpha^2 = 0, 1^2, 2^2, 3^2$. Although our choices of σ_α^2 are too large to be considered 'realistic' we have purposely chosen them to highlight the different impacts of changing expected returns on our various estimators. This preserves the average value of drift whilst steadily increasing its volatility. Besides a time-changing drift, we simulated a time series with stochastic transactions, this is more realistic and applicable to economic data. Based on the empirical, the number of transactions is assumed to be uniformly distributed with mean 307 and standard deviation 228. Again the uniform distribution is chosen for the convenience. However, those numbers generated by the uniform distribution are quite arbitrary and may not reflect the actual number of daily transactions. All procedures should improve with increased time (n) but the assumption of fixed volatility becomes harder to maintain. There is ample evidence of heteroscedasticity of stock returns in the literature. Our choice of parameters is broadly consistent with the empirical data,

[4]We have the historical average equal to 0.0006 for the five companies, which is $(\alpha - \sigma^2/2)$ from Equation 4. Therefore, α is set equal to 0.00068. The rescaling approximately preserves the finite sample distribution in the sense that the simulated data is proportional to a normal distribution whose mean and variance are equal to those of the empirical data.

[5]There is a slight change in the meaning of the notation between here and the definition above in Equation 12.

Table 2. *MSE × 10^3 from the Simulations with $E(\alpha)=0.052$ and $\sigma=1.0$*

α	$\hat{\sigma}_0^2$	$\hat{\sigma}_1^2$	$\hat{\sigma}_{GK}^2$	$\hat{\sigma}_{AGK}^2$	$\hat{\sigma}_{RS}^2$	$\hat{\sigma}_{ARS}^2$
.052	128.118	65.188	16.344	10.548	23.360	12.657
	0.089	0.065	0.008	0.010	0.010	0.012
N(.052, 1)	1482.87	1297.30	10.959	29.416	28.010	14.635
	0.207	0.204	0.010	0.013	0.011	0.013
N(.052, 4)	17551.4	18288.9	124.23	274.70	38.662	17.612
	0.589	0.623	0.016	0.021	0.013	0.016
N(.052, 9)	85255.8	90927.8	757.573	1316.47	52.980	20.095
	1.234	1.319	0.024	0.032	0.014	0.018

*The second number indicates the variance of the estimator, $\text{var}(\hat{\theta})=E[\hat{\theta}^2]-(E[\theta])^2$, the first number is the MSE times 1000.

see 4. We note that $\alpha=0.00068$ and $\sigma=0.013$ corresponds to an annual expected return of 17% given our assumption that the same process governs the market when it is closed or open. This corresponds to the value over our data period (1 May 1990 to 11 October 1991) but may be thought to be too high.

Our random numbers are generated by a random walk with a Gaussian step distribution. If V is the number of steps taken by the random walk in day t, then each step has mean $(\alpha_t-\sigma^2/2)/V$ and variance σ^2/V. In this way we preserve the mean and variance of the daily distribution whilst gradually 'smoothing' its daily sample path. For large V we would expect the data to be virtually 'Brownian'. These generate the values of $s(t)$, $h(t)$ and $l(t)$. This is done once for each simulation and the different methods are applied to these data. Thus we get 1000 values of each estimators. From this sample we calculate $(\Sigma\, \hat{\sigma}_{ij}^2)/1000$ and $(\Sigma\, \hat{\sigma}_{ij}^4)/1000$. These two statistics are used to calculate the sample bias in variance and sample mean squared error (MSE) of the estimators. The results are presented in Table 2. The accuracy of our simulations can be gauged by the use of analytic results for the case of $\hat{\sigma}_0^2$ and $\hat{\sigma}_1^2$ since we know their distributions. For our choice of parameters, the variance of $\hat{\sigma}_0^2$ for the case of constant drift will be 0.093 and for $\hat{\sigma}_1^2$, 0.066[6].

Turning to the MSE calculations in Table 2, we see that:

(i) The naive estimator, $\hat{\sigma}_0^2$ and mean-adjusted estimator, $\hat{\sigma}_1^2$, become less accurate as the volatility of drift increases; in particular, the bias becomes very large indeed. The values from the simulation are broadly in line with their true values, see 6.

(ii) We already noted in the early discussion that both the Garman and Klass and the adjusted Garman and Klass estimators perform very well for fixed and less volatile α. This suggests for the levels of less volatile expected return that one might have in practice the Garman and Klass estimator should be preferred.

(iii) As the volatility of α is increased, generally, for all estimators, bias becomes large, the only exception is the Garman and Klass estimator where the bias is reduced when $\sigma_\alpha=1$. This is because the Garman and Klass estimator is based on the assumption of $\alpha=0$ and the distribution $N(0.052, 1)$ can produce relatively small numbers. For highly volatile α, the adjusted Rogers and Satchell estimator outperforms the others significantly.

(iv) Adjustment procedures seem worthwhile. Unadjusted estimators seem to retain large bias. These are of the order of 100% for $\hat{\sigma}_{RS}^2$. In the same case, the adjusted

[6]In the case of a constant drift and $o(t)=s(t-1)$,

$$\hat{\sigma}_0^2=\frac{1}{n}\sum_{t=1}^{n}[S(t)]^2=\frac{1}{n}\sum_{t=1}^{n}[\ln(s(t))-\ln(s(t-1))]^2$$

$$=\frac{1}{n}\sum_{t=1}^{n}[(\alpha-\sigma^2/2)+\sigma\mathscr{E}(t)]^2$$

$$=(\alpha-\sigma^2/2)^2+\frac{2}{n}(\alpha-\sigma^2/2)\sigma\sum_{t=1}^{n}\mathscr{E}(t)+\frac{\sigma^2}{n}\sum_{t=1}^{n}\mathscr{E}(t)$$

$$E[\hat{\sigma}_0^2]=(\alpha-\sigma^2/2)^2+\sigma^2,$$

$$\text{Var}(\hat{\sigma}_0^2)=\frac{4(\alpha-\sigma^2/2)^2\sigma^2+2\sigma^4}{n}$$

In our simulation, $n=30$, $\alpha=0.052$, and $\sigma=1.0$. Then the true variance of the estimator is 0.093. Similarly, since $\hat{\sigma}_1^2$ is distributed as $\sigma^2\chi^2(n)/n$, $E(\hat{\sigma}_1^2)=\sigma^2$ and $\text{Var}(\hat{\sigma}_1^2)=2\sigma^4/n$, the variance of this estimator for the chosen values is equal to 0.066.

estimators have a bias to variance rate of about 20%. The adjusted Rogers and Satchell always outperforms the unadjusted but the same is true for the Garman and Klass estimators only in the constant drift case.

Some broad recommendations follow. If drift is constant, the adjusted Garman and Klass estimator seems to be preferred on the grounds of both bias and efficiency (small variance). As we allow the drift to vary, the picture changes. Naive estimators and Garman and Klass estimators perform badly although the adjusted Rogers and Satchell still has a reasonably small variance. In this simulation the adjusted Rogers and Satchell estimator outperforms all the others as from our earlier discussion.

IV. CONCLUSION

In this paper, we explained how to estimate volatility using different methods of incorporating high/low data. We considered two possible complications, namely changes in the expected return and the observations of a continuous process at discrete points in time. Both these problems may bias our estimators. Our simulation study suggests that the procedures advocated by the authors, given in Equations 9 and 10 of the paper, seem to be the best compromise and certainly outperform the other alternatives for the experiments chosen.

We address some limitations. Since each stock is isolated, we ignore the covariation thought to exist among stocks in various asset pricing models. Although this should not influence efficient estimation of volatility under our assumptions, one might envisage that there may be valuable information about σ^2 for some possible joint distributions. Dividends and other discrete capital payouts are neglected in our data, since these may violate the continuous nature of the assumed diffusion sample paths. As pointed out several times, there is a strong doubt whether the hypothesis of a diffusion process is the correct model of asset price fluctuations. Different assumptions from Equation 1 will dramatically change the properties of high/low, in particular, the estimators given by Equations 8 to 11 will be incorrect or need further adjustment. Also, in practice, bias may arise from the following sources: to the extent that transactions themselves may convey new information, daytime volatilities may be different from night-time volatilities; bid-ask spreads exist, within which the transactions process may be quite complicated; and volatilities may otherwise be non-stationary in a variety of fashions. Asset prices are much more volatile during exchange trading hours than during non-trading hours although our transactions based models may be able to take this into account. Also, if we assume the stock-specific relationships, using the cross-sectional differences in the relationship, we can obtain a better estimator since the informational content of the high/low data varies across stocks.

It is well known that the variance of aggregate stock returns changes over time. The presence of heteroscedasticity prevents us from estimating the variance over a long-time period. If the volatility of stock returns is not constant, recent data are preferred to predict the future volatility and by sampling the return process more frequently, we can increase the accuracy of the standard deviation estimate. Characterizations of the heteroscedasticity of stock returns are required for the study of distributional properties of stock returns and time-varying expected returns or mean reversion in stock returns.

The importance of high relative efficiency is obvious, in as much as prediction with improved confidence intervals may be constructed from our data bases. Consequently, investigators may adopt the tactic of purposely restricting data usage to combat unforeseen non-stationarities. If so, the procedures that we advocate give high efficiency without using too long a time series and seem able to cope with certain types of non-stationarity. Applied researchers should base their choice of estimation procedure by considering the validity of the various assumptions that we have made. The simplicity of the naive estimators and their robust distribution-free properties explain their persistence in the statistical literature. However, the simplified high/low estimators by Rogers and Satchell seem to perform well for changing drift and reasonable levels of market activity based on only 30 daily observations.

ACKNOWLEDGEMENTS

Financial support from the Newton Trust and Inquire (UK) is gratefully acknowledged. This paper has been produced as a discussion paper for the Autumn 1991 Inquire Conference. The authors thank Quentin Smith (London Stock Exchange) for providing the transaction data.

REFERENCES

Ball, C. A. and Torous, W. N. (1984) The maximum likelihood estimation of security price volatility: theory, evidence, and application to option pricing, *Journal of Business*, **57**, 97–112.

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

Cecchetti, S. G., Lam, P. and Mark, N. C. (1990) Mean reversion in equilibrium asset prices, *American Economic Review*, **80**, 398–418.

Fama, E. and French, K. (1988) Permanent and temporary components of stock prices, *Journal of Political Economy*, **96**, 246–73.

French, K. R. and Roll, R. (1986) Stock return variances; the arrival of information and the reaction of traders, *Journal of Financial Economics*, **17**, 5–26.

French, K. R., Schwert, G. W. and Stambaugh, R. F. (1987) Expected stock returns and volatility, *Journal of Financial Economics*, **19**, 3–29.

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

Merton, R. C. (1980) On estimating the expected return on the market, *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.

Poterba, J. and Summers, L. (1988) Mean reversion in stock prices: evidence and implications, *Journal of Financial Economics*, **22**, 27–59.

Rogers, L. C. G. and Satchell, S. E. (1991) Estimating variance from high, low, and closing prices, *Annals of Applied Probability*, **1**, 504–12.

Schwert, G. W. and Seguin, P. J. (1990) Heteroscedasticity in stock returns, *Journal of Finance*, **XLV**, 1129–55.