The distribution of realised volatility: Evidence of normality and long memory in UK bond futures

David G. McMillan* and Alan E.H. Speight

*School of Management, University of St Andrews, The Gateway, North Haugh,
St Andrews, KY16 9SS, UK.

Tel: +44 0 1334 462800, Fax: +44 0 1334 462812, E-mail: dgm6@st-andrews.ac.uk

Received (in revised form): 16th October, 2006

David G. McMillan is Reader in Finance in the School of Management at the University of St Andrews, UK. His research interests are in the time-series modelling of intra-day financial data, volatility forecasting, interest rate dynamics and the relationship between financial and macroeconomic variables. He has published widely on these topics in national and international journals

Alan E.H. Speight is Professor of Finance in the School of Economics and Business at the University of Wales, Swansea, UK. His research interests are in empirical finance, volatility modelling and forecasting, high frequency financial data, non-linear time series analysis and business cycle dynamics. He has published widely on these topics in national and international journals

Abstract

Daily 'realised' volatility for two UK bond futures is obtained from high frequency, five-minute, returns data. In contrast to previous results using daily data, which have reported that daily returns volatility and daily standardised returns exhibit conditional non-normality, we find that the 'realised' logarithmic standard deviation and returns standardised by 'realised' volatility are approximately normal. These results support the mixture-of-distributions hypothesis that daily returns are characterised by a normal–lognormal mixture. Finally, realised volatility exhibits highly persistent temporal dependence that can be characterised by a long memory process.

Derivatives Use, Trading & Regulation (2006) **12**, 200–208. doi:10.1057/palgrave.dutr.1850040

Keywords: realised volatility; high-frequency; futures markets

INTRODUCTION

The accurate estimation and forecasting of volatility in financial markets is an essential issue, as variation in market returns and other economy-wide risk factors are central to the theory and practice of risk management and asset allocation and pricing. It is now widely recognised that returns volatility is highly predictable and exhibits 'clustering', such that large (small) returns follow large (small) returns of random sign.^{1,2} While volatility is inherently unobservable, most of our knowledge regarding volatility dynamics has been garnered from fitting parametric models such as the generalised autoregressive conditional heteroscedasticity (GARCH) model,^{3,4} or by studying volatility implied by option prices in conjunction with specific option pricing models (eg Black–Scholes).⁵ Such approaches, however, rely upon specific distributional assumptions. Further, the

**Derivatives Use,
Trading & Regulation,**
Vol. 12 No. 3, 2006,
pp. 200–208
© 2006 Palgrave
Macmillan Ltd
1357-0927/06 \$30.00

existence of multiple competing models implies some mis-specification, as only one model could be ‘correct’. Furthermore, an alternative, model-free, approach based upon ex post squared returns, while providing unbiased estimates, is contaminated by measurement error. That is, squared returns are very ‘noisy’ volatility indicators that do not allow reliable inference of the true underlying latent volatility process to be made.

Given the drawbacks of these popular approaches to daily volatility estimation, recent research has suggested an approach based upon ‘realised volatility’.^{6–8} This approach involves constructing estimates of ex post realised daily volatility by summing the squares of intra-day returns. Volatility estimates thus constructed are model-free, and provided sufficiently frequent sampling of the underlying returns is undertaken, are largely free of measurement error (indeed, in the theoretical limit, such estimates are wholly free of measurement error). Hence, we can treat ‘realised’ volatility as observed and implement techniques to examine its distribution directly.

Extant work using this recently developed technique has focussed largely upon examining exchange rate and equity returns volatility.⁹ The results of such work have suggested that daily realised volatility follows a log-normal distribution, while daily returns standardised by realised volatility are conditionally normal. Both of these empirical results are consistent with the ‘mixture-of-distributions’ hypothesis,¹⁰ which states that temporal dependence in volatility arises from the speed and number of intra-day information flows. Further, empirical results also support strong dynamic dependence or ‘long-memory’ within volatility, such that it can be characterised by a fractional integration process which exhibits slow mean-reversion. The aim of

this paper is to apply these new techniques to provide daily volatility estimates of two UK futures markets, namely the Long-Gilt and three-month (short) sterling markets. Such estimates should provide improved distributional and dynamic characterisation of these series over previous daily volatility estimates, and be of interest to researchers in both academia and industry; notably those concerned with risk management, asset pricing, portfolio allocation, forecasting and analysis of market microstructure effects where persistent volatility fluctuations are a central concern.

The rest of this paper proceeds as follows. The next section briefly presents the theoretical underpinnings of realised volatility measures. The subsequent section describes the data used, while the penultimate section describes the distributional characteristics and temporal dependences of realised volatility. The last section summarises and concludes.

REALISED VOLATILITY

The concept of realised volatility was initially introduced as a method for providing an improved measure of ‘true’ volatility against which to compare the forecasts of ARCH and other daily volatility models.¹¹ Empirical research to that point had suggested that ARCH models provided poor out-of-sample volatility forecasts when compared against daily squared returns as the measure of ‘true’ volatility. The use of such daily squared returns as the forecast metric was, however, flawed due to the presence of a large noisy component. That is, returns (r_t) can be written as $r_t = \sigma_t z_t$ with z_t an independent zero mean, unit variance stochastic process, while σ_t is the volatility process which evolves according to the specific model

entertained.¹² Since volatility is not a directly observable process, the common approach is to evaluate forecasts predictions with squared returns over the forecast period. Specifically, if the model for σ_t^2 is correctly specified, then $E_{t-1}(r_t) = E_{t-1}(\sigma_t^2 z_t^2) = \sigma_t^2$, which appears to justify the use of the squared returns over the relevant horizon as a proxy for ex post volatility. While the squared return, however, provides an unbiased estimate of the volatility process, it may yield very noisy measurements due to the idiosyncratic term, z_t^2 . This component displays a large degree of variability compared to σ_t^2 , such that the proportion of variability in squared returns that can be attributed to volatility is low. Thus, it is suggested that the measure of ex post volatility should be based on aggregated intra-day squared returns, which reduces the effects of noise.

Moreover, it can be shown that as the sampling frequency of returns approaches infinity, the ‘integrated’ volatility estimates constructed in this way are, in theory, free from measurement error. To set out the basic idea and intuition assume that the logarithmic $N \times 1$ vector price process, p_t , follows a multivariate continuous time stochastic volatility diffusion:¹³

$$dp = \mu_t dt + \sigma_t dW_t \quad (1)$$

where W_t denotes a standard N -dimensional Brownian motion, and σ the $N \times N$ positive definite diffusion matrix. Further, normalising the unit time interval to represent one trading day, that is $h = 1$, and conditional on the past realisations of μ_t and σ_t , the continuously compounded h -period returns $r_{t+h,h} \equiv p_{t+h} - p_t$ is then:

$$r_{t+h,h} = \int_0^h \mu_{t+\tau} d\tau + \int_0^h \sigma_{t+\tau} dW(\tau) \quad (2)$$

which constitutes decomposition into a predictable or ‘drift’ component of finite

variation and a local martingale. Finally, using the theory of quadratic variation, increments to the quadratic return variation process are of the form:¹⁴

$$[r, r]_{t+h} - [r, r]_t = \int_0^h \sigma_{t+\tau}^2 d\tau \quad (3)$$

which defines integrated volatility and provides a natural measure of the true latent h -period volatility. Moreover, the notion of integrated volatility already plays a central role in the stochastic volatility option pricing literature,¹⁵ where the price of an option typically depends on the distribution of the integrated volatility process for the underlying asset over the life of the option. This measure contrasts sharply with the common use of the squared h -period return as the simple ex post volatility measure which, although it provides an unbiased estimate for realised integrated volatility, is an extremely noisy estimator. Furthermore, for longer horizons any conditional mean dependence will tend to contaminate this latter variance measure, whereas the mean component is irrelevant for the quadratic variation.

Having outlined the techniques underlying the construction of realised volatility, we proceed to obtain the realised volatility measure for our two futures series through summing their intra-day squared returns, before subsequently examining their distributional characteristics and long-memory properties.

DATA

The data analysed consists of the prices of UK government bond (long gilt) and three-month (short) sterling futures contracts traded on the London International Financial Futures and Options Exchange (LIFFE). The sample covers

the period from 24th January, 1992 to 30th June, 1995. The price series on a given day consists of the first traded price of the day in the relevant contract, followed by the last traded price in each five-minute interval for the same contract.¹⁶ The contract price data, p , is sampled at five-minute intervals and transformed to yield the returns series, $r_t = \log(p_t/p_{t-1})$, with the overnight return excluded so as to ensure consistent time-series. With 846 trading days in the sample period, this yields 80,163 and 81,216 intra-day observations for the long gilt and short sterling, respectively. One advantage of using LIFFE futures prices is that they provide a full history of all transactions, as compared to exchange rate data, where only indicative quotes are available, and the 'time and sales' futures data of the Chicago Mercantile Exchange which only records transactions where price changes take place. Hence, the returns data analysed here are potentially far more informative.

Since multiple maturities are traded at any given time for a futures contract, a criterion must be set in order to create a continuous price series. The procedure undertaken here is to use price data on the nearest maturity (front month) contract until the day that traded volume in that contract is exceeded by traded volume in the second nearest contract. This is frequently termed the rollover date, from which point price data is taken from this next maturity contract. This procedure ensures that the most liquid contract is used. Additionally, by deleting overnight returns there is no possibility of calculating returns based on two different maturity contracts. This is the standard method of constructing an intra-day price series from high frequency futures data and given that the contract examined is heavily traded, such a sampling procedure should induce no systematic bias.^{17,18}

Following the construction of five-minute returns from the logarithmic differences between recorded prices at (or before) each five-minute interval, we proceed to construct the daily realised volatility for each series. As stated in Equation (3) and the subsequent paragraph, the limiting result is independent of the drift parameter; however, the use of discrete time intervals could induce dependency into the conditional mean and systematically bias volatility measures. Therefore, in order to remove any serial correlation induced by either the uneven spacing of the observed prices or market microstructure effects (eg traders who distribute their transactions over time¹⁹), we estimate appropriate autoregressive models for each five-minute returns series over the full sample.²⁰ Finally, daily realised volatility is given by:

$$v_t^2 = \sum_{j=1, \dots, [h/\Delta]} r_{t+j\Delta, \Delta} r'_{t+j\Delta, \Delta} \quad (4)$$

where $t = 1, \dots, T$, with T the total number of observations and $\Delta = 1/N$, with N the number of intra-day intervals.²¹ In addition, we consider the realised standard deviation and log standard deviation given by $v_t = \sqrt{v_t^2}$ and $h_t = \log(v_t)$, respectively. Given that volatility now becomes an observable variable, we can implement standard statistical procedures to examine its distributional and dynamic characteristics.

RETURN AND VOLATILITY DISTRIBUTIONS AND VOLATILITY DYNAMICS

The first row in Tables 1 and 2, respectively, present the distributional characteristics of the daily return for short sterling and long gilt series over our sample period. In common with the

Table 1: Distributional characteristics – short sterling

	<i>Mean</i>	<i>Std. dev.</i>	<i>Skewness</i>	<i>Kurtosis</i>
Returns	−0.0008	0.0996	5.3167	90.8066
Standardised returns	−0.5200	0.1019	0.1524	3.4297
Volatility	0.0001	0.0006	19.9531	424.7110
Standard deviation	0.0008	0.0007	10.3532	157.8041
Logarithmic Standard deviation	−0.7216	0.4187	1.3715	7.7419

Notes: All numbers are multiplied by 1,000.

Table 2: Distributional characteristics — long gilt

	<i>Mean</i>	<i>Std. dev.</i>	<i>Skewness</i>	<i>Kurtosis</i>
Returns	0.0002	0.0453	−0.1739	4.8583
Standardised returns	0.1054	0.2077	0.2431	3.8689
Volatility	0.0003	0.0003	2.8946	16.1318
Standard deviation	0.0051	0.0020	1.3571	5.5627
Logarithmic Standard Deviation	−0.0536	0.3729	0.2586	3.2004

Notes: All numbers are multiplied by 1,000.

established literature on the unconditional returns distribution of asset prices, these series exhibit substantial excess kurtosis, that is, their returns distribution have fatter tails than the normal distribution. In contrast with this result, however, the returns standardised by our measure of daily realised volatility, which are presented in the second row in Tables 1 and 2, exhibit an unconditional distribution which approximates the normal distribution, with significant reductions in both skewness and kurtosis, which become closer to the normal distribution values of zero and three, respectively. This result is to be further compared with the conditional distribution of the standardised returns from the appropriate daily GARCH model. More specifically, the skewness and kurtosis values for this conditional distribution

are 0.60 and 8.65 for short sterling and −0.23 and 4.72 for the long gilt, which still indicate the presence of ‘fat’-tailed distributions.

The next three rows in each table present our measures of volatility, namely realised volatility, realised standard deviation and logarithmic standard deviation. Of particular note is that both the realised volatility and standard deviation measures, for both series, exhibit a high degree of right skewness and excess kurtosis (although these values are substantially less for the standard deviation measure). This contrasts with our final measure of volatility, the logarithmic standard deviation, which appears to more closely approximate the normal distribution.²² These results compare well with those reported previously in the literature. In particular, with the recent work on realised volatility,²³ as well as

earlier research, which had used daily data to produce monthly logarithmic standard deviations.^{24,25} Moreover, these two key results for each series, namely, the normally distributed standardised returns and the normally distributed logarithmic standard deviations, suggests that the overall unconditional distribution for daily returns is well approximated by the lognormal-normal mixture.

Having examined the distributional characteristics of unconditional volatility, it is instructive to consider the conditional volatility dynamics. Moreover, the need to provide improved volatility forecasts is of great importance to agents engaged in asset pricing. In order to examine these dynamics, evidence is presented regarding the temporal dependence of the logarithmic standard deviation for each series; therefore Figure 1 presents the time-series plots of the series and Figure 2 the autocorrelation functions. It is clearly evident from Figure 1 that both series exhibit substantial fluctuation, strong persistence and positive serial correlation, which are associated with the well-documented volatility clustering effect. This finding is confirmed by the autocorrelation functions presented in Figure 2 both of which exhibit slow decay. The dashed line in Figure 2 is the conventional upper 95 per cent confidence interval, and suggests rejection of the null hypothesis of no correlation up to a displacement lag of around 75 days for short sterling and 85 days for the long gilt series.

Finally, in an effort to further evaluate the long-memory properties of volatility for these series, using frequency-domain methods, we estimate the degree of fractional integration, d . Specifically, we implement the (log-) periodogram regression estimator d_{GPH} ,²⁶ and the alternative regression estimator d_R ,²⁷ based

Figure 1: Realised logarithmic standard deviation

on the smoothed periodogram using the Parzen window. Briefly stated, these estimators are obtained from the respective least squares regressions:

$$\log[I(w_j)] = \alpha_0 + \alpha_1 \log[2 \sin(w_j/2)]^2 + e_j \quad (5)$$

$$\log[f_s(w_j)] = \alpha_0 + \alpha_1 \log[2 \sin(w_j/2)]^2 + u_j \quad (6)$$

where $w_j = 2\pi j/n$, $j = 0, 1, 2, \dots, m$, defines the set of harmonic frequencies, $I(w_j)$ denotes the sample periodogram, $f(w_j)$ and $f_s(w_j)$ denote the spectral density and smoothed periodogram of the process respectively, $e_j = \log[I(w_j)/f(w_j)] - E \log [I(w_j)/f(w_j)]$, $u_j = \log[f_s(w_j)/f(w_j)]$, $\hat{d} = -(\hat{\alpha}_1/2)$, and the standard error of $\hat{d}$ is

Figure 2: Autocorrelations functions

dependent only on m , according to

$$m(\hat{d} - d) \sim N(0, \pi^2/24).$$

The estimates of the fractional integration parameter obtained using the d_{GPH} and d_R methods, respectively are (with t -statistics in parentheses) 0.3746 (2.6182) and 0.3377 (4.1222) for the long gilt and 0.2517 (1.7595) and 0.1633 (1.9931) for the short sterling, respectively. Thus, with the exception only of the d_{GPH} estimates for short sterling which is significant at the 10 per cent level, all d estimates are statistically significant and lie in the region $0 < d < 0.5$. These results thus support the belief that volatility exhibits long memory with a slow hyperbolic rate of decay to shocks. Moreover,

the long gilt series exhibits a stronger degree of persistence, confirming the results of the autocorrelation function, and suggest that empirical models of volatility must be able to capture these long-lag autoregressive effects.

SUMMARY AND CONCLUSION

This paper has implemented a recently proposed method for obtaining model-free measures of daily volatility using high-frequency intra-day UK bond futures data. More specifically, we follow the approach based upon ‘realised volatility’ (see, for example, Andersen *et al.*^{6–8}). This method involves constructing estimates of ex post realised daily volatility by summing the squares of intra-day returns. This volatility estimate is thus model-free, largely free of measurement error and can be treated as observed allowing the implementation of standard statistical techniques to examine its distribution directly.

The results presented here broadly support those found using exchange rate and equity price data, namely, that daily returns and daily realised variances and standard deviations exhibit non-normal unconditional distributions and, in particular, are characterised by excess kurtosis; but that the distributions of returns standardised by realised volatility and logarithmic standard deviations approximate the normal distribution. These results support the mixture-of-distributions hypothesis originally advocated by Clark¹⁰ that daily returns are characterised by a normal–lognormal mixture. Finally, our results indicate strongly persistent dynamics within realised volatility.

Moreover, the results presented in this paper should be of interest to market participants and practitioners. Perhaps most importantly volatility

estimation and forecasts appear prominently in many practical financial decisions, such as asset allocation, risk management and derivative pricing. The method of volatility construction examined here, which had been used only as a device for evaluating the forecasts of various parametric volatility models, could instead be the primary volatility measure of interest. Further, the improved understanding of the distributional characteristics of returns and volatility could allow for improvements in the accuracy of value-at-risk type calculations. Finally, this methodology can be extended to a multivariate context and therefore allow for the construction of realised covariances and hence generate realised betas.

References

- 1 Mandelbrot B. (1963) 'The Variation of Certain Speculative Prices', *Journal of Business*, Vol. 36, pp. 394–419.
- 2 Fama E. (1965) 'The Behaviour of Stock Market Prices', *Journal of Business*, Vol. 38, pp. 34–105.
- 3 Engle R.F. (1982) 'Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom inflation', *Econometrica*, Vol. 50, pp. 987–1007.
- 4 Bollerslev T. (1986) 'Generalised Autoregressive Heteroscedasticity', *Journal of Econometrics*, Vol. 31, pp. 307–327.
- 5 For a review of the parametric ARCH literature, see Bollerslev T. Engle R.F. and Nelson D.B. (1994) 'ARCH Models, in Engle, R.F. and McFadden, D.L. (eds.), *Handbook of Econometrics*, Volume IV, Elsevier Science Publisher, Amsterdam.
- 6 Andersen T.G., Bollerslev T., Diebold F.X. and Ebens H. (2001) 'The Distribution of Realised Stock Return Volatility', *Journal of Financial Economics*, Vol. 61, pp. 43–76.
- 7 Andersen T.G., Bollerslev T., Diebold F.X. and Labys P. (2001a) 'Modelling and Forecasting Realised Volatility, NBER Working Paper #8160.
- 8 Andersen T.G., Bollerslev T., Diebold F.X. and Labys P. (2001b) 'The Distribution of Realised Exchange Rate Volatility', *Journal of the American Statistical Association*, Vol. 96, pp. 42–55.
- 9 See Andersen *et al.*^{6–8} and the references cited therein.
- 10 Clark P.K. (1973) 'A Subordinate Stochastic Process Model with Finite Variance for Speculative Prices', *Econometrica*, Vol. 41, pp. 135–155.
- 11 Andersen T.G. and Bollerslev T. (1998) 'Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts', *International Economic Review*, Vol. 39, pp. 885–905.
- 12 That is, following Andersen and Bollerslev¹¹ and Andersen T.G. Bollerslev T. and Lange S. (1999) 'Forecasting Financial Market Volatility: Sample Frequency *vis-a-vis* Forecast Horizon', *Journal of Empirical Finance*, Vol. 6, pp. 457–477.
- 13 For a more thorough treatment of the issues raised and discussion of alternative processes including jump processes, see Andersen *et al.*^{6–8}
- 14 The quadratic variation process measures the realised sample path variation of the squared returns process and suggests that we may approximate the quadratic variation by summing high-frequency squared returns (see Andersen *et al.*^{7,8} for further details).
- 15 Hull J. and White A. (1987) 'The Pricing of Options on Assets with Stochastic Volatilities', *Journal of Finance*, Vol. 42, pp. 381–400.
- 16 The choice of intra-day sampling frequency is a trade-off between the accuracy of the theoretical approximations underlying the construction of realised volatility and confounding influences from market microstructure frictions and clustering effects. The choice of five-minutes is seen as a reasonable balance between these latter effects and the necessity to ensure that daily realised volatilities are largely free of measurement error.
- 17 ap Gwilym O., Buckle M., Thomas S. and Woodhams M. (1998) 'Intraday Empirical Regularities in Interest Rate and Equity Index Futures Markets, and the Effect of Macroeconomic Announcements', *Journal of Business Finance and Accounting*, Vol. 2, pp. 921–944.
- 18 Abhyankar A., Copeland L.S. and Wong W. (1999) 'LIFFE Cycles: Intraday Evidence from the FTSE-100 Stock Index Futures Market', *European Journal of Finance*, Vol. 5, pp. 123–139.
- 19 Hasbrouck J. and Ho T. (1987) 'Order Arrival, Quote Behaviour, and the Return-generating Process', *Journal of Finance*, Vol. 42, pp. 1035–1048.
- 20 Autoregressive lag lengths were determined with reference to the BIC and residual diagnostic tests for remaining serial correlation.
- 21 The floor trading times for long gilt futures changed on 1st August, 1994 from 8:30am–4:15pm to 8:00am–4:15pm. So as not to induce any systematic bias into the daily realised volatility, conditional mean estimation and realised volatility construction was conducted over separate samples defined by the trading times change.

-
- The number of intra-day intervals in the first sample is 93 that increases to 99 in the second sub-sample. For the short sterling series there are 96 intra-day intervals.
- 22 This finding is even stronger for the short sterling series if account is taken of the sterling-ERM episode. Upon the exclusion of this period, the skewness and kurtosis values fall to 7.8 and 80.4, 3.2 and 20.5, 0.9 and 4.4 for volatility, standard deviation and logarithmic standard deviation, respectively.
- 23 See for example Andersen *et al.*⁶ on stock return data and Andersen *et al.*^{7,8} on exchange rate data.
- 24 Taylor S.J. (1986) 'Modelling Financial Time Series', John Wiley: Chichester.
- 25 French K.R., Schwert G.S. and Stambaugh R.F. (1987) 'Expected Stock Returns and Volatility', *Journal of Financial Economics*, Vol. 19, pp. 3–30.
- 26 Geweke J. and Porter-Hudak S. (1983) 'The Estimation and Application of Long Memory Time Series Model', *Journal of Time Series Analysis*, Vol. 4, pp. 221–238.
- 27 Reisen V.A. (1994) 'Estimation of the Fractional Difference Parameter in the ARIMA(p,d,q) Model using the Smoothed Periodogram', *Journal of Time Series Analysis*, Vol. 15, pp. 335–350.