

Option Implied and Realised Measures of Variance

Damien Lynch and Nikolaos Panigirtzoglou*

First Draft: January 2003

This version: December 2003

* Monetary Instruments and Markets Division, Bank of England, Threadneedle Street, London EC2R 8AH, email damien.lynch@bankofengland.co.uk , nikolaos.panigirtzoglou@bankofengland.co.uk . The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of England.

Abstract

This paper analyses the differences between forward (risk-neutral expectation) and realised variance. Four different assets are examined: S&P 500, FTSE 100, eurodollar and short sterling futures. Model-free measures of risk-neutral forward variance are estimated using the full cross section of option prices of a given maturity. Intra-day futures data are used in the estimation of realised variance. A mean squared error criterion is devised to inform the choice of optimal intra-day frequency. Market microstructure issues are considered for both equity and interest rate futures markets. The bias and efficiency of forward variance as a predictor of realised variance is re-examined.

JEL Classification: G13; G19

Keywords: implied volatility; realised volatility; risk premia; market microstructure; high-frequency data

1. Introduction

This paper examines the difference between forward (risk-neutral) and realised variance. Over a long period, this difference can be considered as an ex-post measure of variance risk premia if market participants are rational, that is, they don't make any systematic errors in their variance forecasts. Most of the past literature has focused on the informativeness and unbiasedness of implied volatility from option prices as a guide or predictor for the future volatility of the underlying asset. A smaller part of the literature dealt with measurement issues of both forward (option implied) and realised volatility. We provide a brief description of this literature.

Early studies in the literature generally concluded that, when compared with past volatility, implied volatility was a better predictor of future volatility. However some doubt was raised as to whether implied volatility reflected all publicly available information at a given point in time. Latane and Rendleman (1976) using a cross section of Black-Scholes implied volatilities for twenty-four firms found that the standard deviations implied in call prices were strongly correlated with actual standard deviation. Becker (1981) found, also using equity options, that both past and implied volatility were found to provide the best forecast of future volatility. Such a finding raises questions about the efficiency of implied volatility as a forecast (does it fully reflect all information available to the market of an option contract) and/or the accuracy of the Black-Scholes pricing model.

This early empirical work was undermined by a lack of data and small sample size. As more data became available with the expansion of option markets, a number of papers sought to revisit the issue in the late eighties and throughout the nineties, focussing on time series as opposed to cross-sectional analysis. The papers by Day and Lewis (1992), Canina and Figlewski (1993) and Lameroux and Lastrapes (1993) are noted for their challenge to the conclusions of the earlier work with findings that refuted the usefulness of implied volatility as a guide to the future variability of an asset price. Canina and Figlewski (1993) regressed daily one-month volatility for the S&P 100 index on observations of one-month implied volatilities on S&P 100 index options and on past S&P 100 index volatility. The obvious overlapping observation problem was addressed by adjusting the standard errors. Implied volatilities from a

range of strike prices were examined rather than simply the at-the-money volatility. Their regressions generally returned very low R-squared and insignificant coefficient estimates, prompting them to conclude that implied volatility is not a good predictor of future volatility. Further, historical volatility was actually found to be a more accurate forecast than implied volatility over their sample period of 1983-1987. Rather than provide a forecast of conditional volatility, they thus conclude that implied volatility is merely another piece of information to be used in its estimation.

Several studies since the mid-1990's have tackled these results from both methodological and econometric perspectives. This work has generally taken the line that these negative conclusions are the result of mis-measurement of implied uncertainty; the risk neutral nature of implied volatility; inappropriate statistical inference (due mainly to overlapping observations); and/or a wrongly specified pricing model.

Jorion (1995) focussed on foreign exchange (FX) rather than equity index options. He claims that biases in the measurement of implied volatilities are less likely for FX options as these options are traded side by side with the underlying asset (hence no synchronicity problems) and also because prices are rigorously checked and thus less susceptible to clerical errors. This could induce a stale price effect and also mean less scope for active arbitrage between derivative and spot markets. The implications of a potential bid-ask distortion were also investigated. Having adjusted for all of these features, Jorion finds consistent evidence across three currency pairs over the period 1985 – 1992 that implied volatility contains information about future volatility and is a superior, though biased, forecast. The bias in implied volatility was found to be positive so that the future variability of the underlying asset price implied by options prices is, on average, greater than that subsequently realised. These findings for FX implied volatility have been supported by others such as Xu and Taylor (1995).

Christensen and Prabhala (1998) re-assess and greatly expand on the Canina and Figlewski (1993) paper. In particular, they use a much longer sample period allowing them to use non-overlapping observations. In addition to the finding that implied volatility cannot be outperformed by purely historical-based forecasts, they also find that, having adjusted for a potential errors-in-variables problem, the degree of bias in

implied volatility forecasts is much less than that suggested by previous empirical work in the literature.

Fleming (1998) and Blair et al. (2001) both re-visit the issue for S&P 100 implied volatility. In a very rigorous paper, Fleming (1998) attempts to minimise the errors in measuring implied volatility for equity indices and also employs a greater degree of sophistication in his econometric and statistical analysis (providing GMM standard error corrections for telescopic-fixed expiry date maturity patterns). Similar to Jorion (1995) he finds that, though biased, implied volatility subsumes all information from historical volatility as a forecast for future volatility and that the forecast errors from implied volatility are not correlated with many variables that are associated with conditional volatility. These results were consistent across a range of forecast horizons – from one day to several months – and prompt Fleming (1998) to conclude that using a linear model that adjusts for the bias in implied volatility can provide a useful guide for future volatility. Blair et al. (2001) support these findings in a paper that contributes to the debate by using alternative sources of data for both realised and implied volatility. Realised volatility is measured using intra-day data (sampled every five minutes) on the S&P 100 index while the CBOE Market Volatility Index (VIX) is used as a measure of implied volatility.¹ By construction the VIX avoids many of the potential sources of measurement error mentioned earlier. With the exception of the one-day-ahead horizon, Blair et al. (2001) find that S&P 100 implied volatility can also outperform very high frequency estimates of historical volatility in forecasting future volatility. Bollen and Inder (2002) also use high frequency intra-day data to estimate daily realised volatility. They account for heteroskedasticity and autocorrelation with the VARHAC estimator developed by den Haan and Levin (1996) by using a weak set of assumptions about the data generating process.

Focussing on the nature of the bias in implied volatility as a measure of expected future volatility, some preliminary work by Chernov (2001) theoretically justifies, and seeks to explicitly account for, the bias in implied volatility by the existence of a volatility risk premium. He assumes a particular stochastic process for actual

¹ The VIX is a weighted average of S&P 100 option implied volatilities that is computed and made available on a real time basis. By construction the index is a conditional volatility measure that is

volatility so that the estimated time-varying volatility risk premium is model dependent. The volatility risk premium is proxied by using historical volatility. He finds that implied volatility is both an unbiased and efficient forecast for future volatility. However these findings are undermined when broken down into sub-samples and this is attributed to econometric and measurement difficulties.

Overall a review of the literature seems to support a tentative conclusion that implied volatility does provide a better forecast or guide for future volatility than forecasts based solely on historical information. As a forecast it is, however, biased and any attempt to use option implied volatilities as a guide to future volatility must account for the volatility risk premium. The analysis of this bias is the focus of this paper.

We contribute to the existing literature in a number of ways. The first contribution relates to the issue of measurement of both forward (risk-neutral expectation) and realised variance. In particular, we use a measure of forward variance that embodies information from the full cross-section of option prices on a given date and is model independent. In this respect it is theoretically more appealing than measures such as the at-the-money volatility or VIX index used in previous studies. To measure realised variance we use high frequency data. In choosing the optimal frequency we apply a mean squared error (MSE) criterion by combining the asymptotic theory on realised variance with an econometrically robust estimate of the bias in realised variance due to market microstructural effects.

In addition we examine equity and interest rate markets in both the US and the UK, for a substantial time periods that include events such as the 1987 crash, the 1992 ERM crisis, the 1997 Asian crisis and the 1998 Russian default and LTCM crisis. We provide evidence on the relative importance of microstructural effects in the different markets and their implications in terms of volatility and volatility risk premia.

Finally, using these measures of variance, we re-visit the predictability of realised variance from forward variance implied from option prices. In testing for efficiency we use, apart from an intra-day measure of historical variance, measures of

market-based, forward-looking and has a constant forecast horizon of one month. For more information see Fleming (1998).

asymmetry and ‘fatness’ of tails implied from option prices. We provide estimates of average, ex-post, variance risk premia in both equity and interest rate markets, assuming that expectational errors offset each other over our relatively long sample periods.

The paper is organised as follows. Section 2 provides discusses the estimation and data issues of forward variance. Section 3 focuses on the measurement of forward and realised variance. Section 4 examines the bias and efficiency of forward variance as a predictor of future variance and analyses the properties of ex-post variance risk premia. Section 5 concludes.

2. *Estimation and data issues for forward variance*

Option premia for contracts with different strike prices can be used to provide a measure of uncertainty expected by market participants. Crucially, this measure reflects the future volatility that would be expected to occur in a world where investors are risk neutral. Instead real-world investors are risk-averse and so care about risk and a premium will be required by investors to bear non-diversifiable volatility risk.

Britten-Jones and Neuberger (2000) provide a characterisation of all continuous stochastic processes consistent with a given set of option prices, which enables risk-neutral forecasts of variance to be obtained that are independent of the stochastic process of the asset price or volatility. The only assumption that is needed is that the stochastic process is continuous. They derive the following risk-neutral variance forecast:

$$E_t \left[\int_t^T \left(\frac{dS_t}{S_t} \right)^2 \right] = 2 \cdot e^{r_t(T-t)} \cdot \left(\int_0^{S_t} \frac{P_t(T, K)}{K^2} dK + \int_{S_t}^{\infty} \frac{C_t(T, K)}{K^2} \right) \quad (1)$$

where S_t is the underlying asset, r_t is the risk-free rate, $C_t(T, K)$, $P_t(T, K)$ are the European call and the put prices respectively with expiry date at time T and strike price K .

This variance forecast is related to the fair value of variance swaps that were first introduced by Carr and Madan (1999) and Demeterfi, Derman and Kamal (1999). Variance swaps are forward contracts on annualised variance, that is, the payoff depends on the difference between the realised variance of the underlying asset over the life of the contract and the delivery (forward) price for variance. Carr and Madan (1999) and Demeterfi, Derman and Kamal (1999) show that the fair price of a variance swap (forward variance) is determined by the cost of a replicating portfolio given by the right-hand side of equation (1). This portfolio should consist of option premia for contracts with strikes ranging between zero and infinity. Therefore, although the derivation of equation (1) is independent of the stochastic process of the underlying asset, it assumes the existence of a continuum of strikes, which in practice doesn't exist. This means that the forward variance of equation (1) depends on the interpolation method used to generate a continuum of strikes from a limited range of available strikes. This limitation increases the risk of hedging variance swaps.

The variance forecast in equation (1) differs from Black-Scholes implied variance in that it contains information from all traded strikes and not just the at-the-money. More importantly, the Black-Scholes model assumes constant volatility and, as noted in Britten-Jones and Neuberger (2000), it is inconsistent to forecast volatility from constant volatility models. Additionally, in contrast to VIX or similar measures, it is more theoretically appealing in that it does not rely on ad hoc weights or restrictive model assumptions.

Another important issue mentioned in Britten-Jones and Neuberger (2000) is the convexity adjustment bias when forecasting volatility instead of variance as implied by equation (1). In particular, the risk-neutral forecast of volatility, that is, the risk-

neutral forecast of the square root of variance $E_t \left[\sqrt{\int_t^T \left(\frac{dS_t}{S_t} \right)^2} \right]$ will be smaller than the

square root of the forward variance $\sqrt{2 \cdot e^{r_i(T-t)} \cdot \left(\int_0^{S_t} \frac{P_t(T, K)}{K^2} dK + \int_{S_t}^{\infty} \frac{C_t(T, K)}{K^2} \right)}$. In

Appendix 1 we derive a formula for the convexity adjustment. This is an important issue for the analysis of the differences between forward and realised uncertainty in

section 4. Although Table 1.1 shows that the average convexity adjustment is not very large (between 0.15-0.32%), the adjustment, itself a function of volatility, is time varying and so regressions in volatility space are likely to be affected. For this reason we compare realised and forward uncertainties in variance space as opposed to volatility (i.e. standard deviation) space.

As mentioned above, to estimate the risk-neutral forecast of variance, that is, the right-hand side of equation (1), we need a continuum of strikes. To generate a continuum of strikes we employ the cubic smoothing spline interpolation of the implied volatility smile in delta space², as described in Bliss and Panigirtzoglou (2002). After fitting³ the smoothed implied volatility function, 20000 points along the curve are converted into European call and put prices (using the Black-Scholes formula). These option prices are then used in the estimation of forward variance:

$$2 \cdot e^{r_t(T-t)} \cdot \left(\int_0^{S_t} \frac{P_t(T, K)}{K^2} dK + \int_{S_t}^{\infty} \frac{C_t(T, K)}{K^2} dK \right).$$

The option contracts used in this paper have fixed expiry dates, that is, the time to maturity changes with time. Our input dataset includes daily call and put option and futures prices on all traded (quarterly) contracts for the FTSE 100 index and S&P 500 futures options and three-month sterling (short-sterling) and eurodollar LIBOR interest rate futures options. The FTSE 100 and short sterling options contracts are all traded on London International Futures and Options Exchange (LIFFE). Daily settlement prices from LIFFE were obtained from 1987 for short sterling futures options and from 1992 for FTSE 100 index options. The associated value of the underlying was the futures price reported by LIFFE. Options contracts on eurodollar futures and the S&P 500 futures are traded on the Chicago Mercantile Exchange and daily settlement prices for these option contracts were obtained from 1983 for S&P 500 and from 1985 for eurodollar. The associated value of the underlying was the settlement price of the futures contract maturing on or just after the option expiry

² The interpolation of the implied volatility smile in delta space has the advantage that far out-of-the-money options are grouped together in the tails allowing for more shape near the centre of the distribution where more trading occurs.

³ The smoothing parameter controls the tradeoff between smoothness and goodness-of-fit. After experimenting with different values we chose a value of 0.99 for all contracts used in this study.

date. The risk free rates used are the British Bankers Associations 11am fixings for short sterling and eurodollar LIBOR rates reported by Bloomberg.

Only at- and out-of-the-money call and put prices were used because there is usually more trading in these, rather than, in-the-money options (see Bliss and Panigirtzoglou (2000)). Option prices that violated the monotonicity⁴ or convexity⁵ properties were discarded. Option prices for which an implied volatility was impossible to compute or with deltas smaller than 0.01 or greater than 0.99 (far out-of-the money options with usually little or no trading) were also discarded. Following this screening process, we fit the available implied volatilities of a given maturity cross section in delta space with the smoothing spline method. However, it is necessary to extrapolate the spline beyond the range of available data. Since the spline method extrapolates linearly outside the available range (resulting sometimes in negative or implausible large implied volatilities) we force the spline to extrapolate horizontally. This is done by introducing three pseudo strikes above and below the available range with implied volatilities equal to that of the respective available extreme strikes.

The output from the estimation process consists of daily forward variances for all available expiry dates. To compare these measures of forward variance with those from the Black-Scholes model in volatility space, we adjust the square root of the forward variance to account for convexity and examine the difference between this measure and the Black-Scholes at-the-money volatility.⁶ In calculating the convexity adjustment we use the convexity adjustment formula set out in Appendix 1, where $E_t[y]$ is the forward variance (i.e. that of equation (1) above) and $Var_t[y]$ is proxied with the variance of the forward variance, using a three-month rolling window. These differences are plotted in Appendix 2 and the average absolute differences are shown in Table 2.1. On average the difference between the two measures is around 1.5 percentage points and thus can have a significant effect when examining the bias in risk-neutral or forward volatility as a predictor of realised volatility.

⁴ Monotonicity requires that the call (put) prices are strictly decreasing (increasing) with respect to the exercise price.

⁵ It requires that a butterfly spread at a particular strike (formed by selling two call options at this strike and buying the two adjacent call options) is positive. In a continuum of strikes, this is equivalent to require that the call and put price functions are convex.

3. *Estimation and data issues of realised variance*

Much of the recent literature on finance has focussed on the accurate estimation of the volatility of asset prices. Volatility is usually considered to be unobservable, a corollary to the view that data generating processes for asset prices are latent in nature. However computational and theoretical advances (diffusions etc.), as well as the availability of transaction by transaction asset prices, have brought us closer to an unbiased, consistent and efficient estimator for volatility.

The estimator for realised volatility that is utilised in most studies is very simple. It is based on evaluating the quadratic variation of the price series – that is, a sum of the squared asset price returns (SSR) at a particular frequency over a particular horizon. This provides an estimate of the ‘integrated volatility’ of the underlying asset price process; the ‘actual volatility’ for a period within this horizon is then taken as the difference between successive integrated volatilities. Recent advances in the understanding of the properties of quadratic variation as an estimator of integrated volatility have found it to be a consistent and efficient estimator (Barndorff-Nielsen and Shephard (2002)). In finite/smaller samples, Monte Carlo experiments suggest that a high number of intra-day observations are needed to rely on the asymptotic results (Barndorff-Nielsen and Shephard (2003)). These results have led some papers in the literature that use intra-day data to measure volatility in this way, to regard the resulting estimated volatility series to be as near to the true volatility of the asset price process as to regard true volatility as observable (see for example Andersen et al (2000)). In regarding volatility as close to observable, it may be examined without the need to account for measurement error.

Recent additions to the literature have sought to use intra-day data to provide more reliable estimates of realised volatility (for example see Taylor and Xu (1997), Andersen et al (2000, 2001), Bollen and Inder (2002)). Some of these papers have used the sums-of-squares estimator mentioned above. However the use of this estimator with very high frequency data can be problematic. The literature on market microstructure has identified several features that are observed empirically in high

⁶ The Black-Scholes implied volatility can be regarded as an expectation of future volatility for near-

frequency data. These can induce autocorrelation at very high frequencies (e.g. 1-minute or 5-minute). Chief among such features is the impact of the bid-ask spread, infrequent trading, order-spreading etc. Use of the SSR estimator requires serially uncorrelated data. Andersen et al. (2000) specifically guard against the indiscriminate use of the SSR in these circumstances. They explicitly recognise a tradeoff in choosing a high enough frequency to produce an estimate with minimal sampling variation on the one hand, but avoiding the biases that may be induced by microstructural features in very high frequency data on the other. In choosing the ‘optimal’ frequency they propose a tool called the ‘volatility signature plot’. Choice of frequency involves visual inspection of a plot of realised volatilities over various intra-day frequencies. The frequency at which the plot flattens out – i.e. produces successively similar volatilities – is taken to be the optimal frequency. Bollen and Inder (2002) revisit the theoretical basis for using the SSR estimator in measuring realised volatility and seek to take account of the nature of these empirical features. They adapt a procedure first developed by den Haan and Levin (1996) that allows consistent estimation of covariance matrices in vector autoregressions in the presence of both heteroscedastic and autocorrelated errors. Bollen and Inder (2002) show this estimator, known as the VARHAC estimator (described below), to be well suited to the high frequency return data that we are interested in. Theoretically, they provide an accurate estimate of realised volatility that is not confounded by microstructural effects, at the highest frequency available.

This paper seeks to contribute to the existing literature by also utilising intra-day data in the estimation of realised volatility. The methodologies set out above to estimate realised volatility, though intuitively appealing, can involve either some subjective choice for the practitioner (Andersen et al (2000) case), or are computationally intensive, and thus impractical, over large samples (Bollen and Inder (2002)). Our approach is to combine these two methodologies with the recently developed asymptotic theory on realised variation by Barndorff-Nielsen and Shephard (2002, 2003). In doing so we devise a framework that will provide for an objective choice of intra-day frequency at which to apply the SSR estimator to produce a measure of realised volatility that is robust (econometrically) to unwanted microstructural effects.

expiration, at-the-money options. See Feinstein (1989) for further details.

The core of our approach to choosing an optimal intra-day frequency lies in a mean square error (MSE) criterion. For any estimator, the MSE may be decomposed as the sum of the square of the bias of the estimator and the squared standard error of the estimator. Employing this decomposition, we seek to combine, for each frequency, the bias in the SSR estimator, due to microstructural effects, with an estimate of the standard error of the SSR estimator. In doing so, we explicitly recognise both elements of the tradeoff - the autocorrelation bias and the reduction in standard error due to larger samples at higher frequencies - involved in choosing an optimal intra-day frequency. We then choose that frequency that is found to have the minimum MSE, thus minimising the tradeoff.

3.1 Estimating the serial correlation bias

In arriving at an estimate of the bias in the SSR estimator of realised variance, we compare SSR estimates of realised variance with the consistent variance estimates obtained by implementing the estimation techniques set out in Bollen and Inder (2002). The methodology developed in that paper is based on applying the vector autoregressive heteroskedastic and autocorrelation consistent (VARHAC) estimate of the variance/covariance matrix of a data process as set out in den Haan and Levin (1996). Bollen and Inder (2002) find that the relatively weak conditions for the use of the VARHAC estimator are well satisfied by intra-day data. In the case of a univariate analysis, the core of the VARHAC estimator involves estimation of an autoregression of the squared intra-day returns. The residuals from this autoregression are then used to calculate the spectral density of the returns. Spectral densities are commonly used in time series analysis to study cycles in time series (at various trigonometric, as opposed to time, frequencies) and are closely related to the autocovariance function of a time series. Evaluation of the spectral density at a trigonometric frequency of zero provides an estimate of the variance of the time series. Thus, Bollen and Inder (2002), following den Haan and Levin (1996), evaluate the spectral density of the returns, obtained using the residuals from the autoregression, at a trigonometric frequency of zero to estimate the realised variance using intra-day frequency. The use of the residuals from the autoregression means

that the data used in estimating the realised variance of the series has been ‘cleansed’ of the effects of any serial correlation that may be present.

Implementation of the VARHAC estimator is extremely intensive. Autoregressions are repeatedly run for a given time frequency, to estimate the daily realised variation. In choosing the order of the autoregression, a testing down procedure where, beginning with a maximum lag, autoregressions at successively lower orders are estimated. The choice of lag length for each frequency, each day, is based on maximising an information criterion, such as the Akaike (1973) Information Criterion (AIC) or Schwartz (1978) Bayes Information Criterion (BIC). This procedure must be repeated each day for each frequency and involves the estimation of a very substantial number of regressions. For our purposes comparison of the realised with implied variance involves estimating realised variances over, very often, much longer periods of time than just one day. The time to expiry of an option, and so the horizon for the forward variance that we calculate, declines each day. This horizon may be as great as six months and as low as five days for example. As a result implementation of this strategy with samples of intra-day data over periods ranging from eleven to twenty years is not feasible. To address this, we obtain estimates of realised variance using the VARHAC procedure at a set of fixed numbers of days to maturity for each contract, rather than at every trading day. The process can be summarised as follows:

1. We focus on ten intra-day frequencies – 1, 5, 10, 15, 30, 45, 60, 120 and 400 minutes. The 400 minute frequency provides one observation per trading day (i.e. the last trade).
2. We estimate both SSR and VARHAC realised variance, at each of frequencies, over the remaining lives of the option contracts on trading days that are approximately 30, 60, 90, 120 and 150 days away from expiry.
3. For each contract, the average difference between the VARHAC and SSR variances is then taken to provide an estimate of the bias in the SSR estimate. In this way, the VARHAC estimated variance is taken to be a proxy for the true realised variance of the asset concerned.

Implementation of the VARHAC estimator follows the procedure set out in Bollen and Inder (2002) and den-Haan and Levin (1996) and so is just summarised in the following steps:

- a) For a specified frequency and a given contract with time to maturity T , we estimate the following regression for each day $t = T-30x$; $x=1,2,\dots,5$, using all of the intra day observations, $j=1, \dots, J(T-30x)$, where J is the total number of intra-day observations at the specified frequency between t and T :

$$r_{j,t}^2 = \sum_{i=1}^{\tilde{K}} \beta_i r_{j,t-i}^2 + e_{j,t}, \text{ for each lag length } k = 1, 2, \dots, \tilde{K}.$$

In line with the advice in den Haan and Levin (1996) we take \tilde{K} to be $(J(T-30x))^{1/3}$.

- b) Again following den Haan and Levin (1996) we compute BIC values for each lag length k and date t . For each day t , we then select the residuals from the regression that minimises the BIC across the different values of k . We denote the optimal value of k on date t , \hat{k}_t , the corresponding residuals, \hat{e}_t , where \hat{e}_t is a $(J(T-30x)) \times 1$ vector and the $\hat{k}_t \times 1$ vector of coefficients $\hat{\beta}_t$.
- c) The VARHAC realised variance is then estimated as:

$$\sigma_{t,T}^{VH} = \frac{\hat{e}_t' \hat{e}_t}{\left(1 - \sum_{k=1}^{\hat{k}_t} \hat{\beta}_{k,t} \right)}$$

3.2 Estimating the standard error of the SSR realised variances

The second element of the MSE relates to the accuracy with which the SSR estimator can estimate realised variance. Theoretically, the more data that is used the greater the accuracy, and the one minute frequency should provide as close an estimate as possible of actual realised variance. The gain in terms of accuracy must be measured relative to the loss in terms of bias at the very high frequencies. To capture the

standard error of the realised variances estimated using the SSR estimator, we use the relatively recent literature on the asymptotic theory of realised variation. In the context of a fairly general stochastic volatility model, Barndorff-Nielsen and Shephard (2002) develop a limit theory for the difference between realised and integrated volatility, where integrated volatility refers to the continuous time analogue of sums of squared returns and thus refers to actual volatility in a theoretical sense.

This limit theory provides us with information on the asymptotic properties of the SSR estimator. One of the properties of the SSR estimator, under this theory and assuming a very general stochastic volatility model, is that it is a consistent estimator for actual realised variance. That is, Barndorff-Nielsen and Shephard (2002) show that the SSR estimator converges asymptotically to a normal distribution, with a mean given by the actual volatility. However, the usefulness of this result in finite samples is questioned in Barndorff-Nielsen and Shephard (2003), where they use Monte Carlo methods to show that very high intra-day frequencies are needed for a good approximation in finite samples. They also find that the approximation in finite samples improves when a transformation to logarithmic space is used (that is, the SSR estimator is lognormally distributed), even for lower intra-day frequencies. That volatility is better approximated in finite samples by a lognormal distribution is a result that is acknowledged in the literature. Andersen et al (2001) find that in the case of equity volatilities, logarithmic standard deviations for 30 stocks from the Dow Jones Industrial Average are well approximated by a normal distribution. In the case of interest rate implied volatilities, Lynch and Panigirtzoglou (2003) also find logarithmic implied volatilities to have normal distributions over samples of at least ten years of data. Following Barndorff-Nielsen and Shephard (2003), we have that

$$\frac{(\log SSR_n - \log \sigma_n^2)}{\sqrt{\frac{2}{3} \frac{\sum_{j=1}^M y_{j,n}^4}{\left(\sum_{j=1}^M y_{j,n}^2\right)^2}}} \xrightarrow{L} N(0,1), \quad (2)$$

denoting actual volatility over horizon n by σ_n^2 , the SSR estimator by SSR_n and $y_{n,j}$ as the j th period return of M intra-day periods. The denominator of (2) provides an estimate of the standard error of the logarithm of the SSR estimator. To estimate the standard error for the SSR estimator itself using this result we make the following transformation from the normal variable $\log SSR_n$ to a lognormal variable SSR_n :

$$Var(SSR_n) = \{E(SSR_n)\}^2 \cdot (e^{Var(\log SSR_n)} - 1). \quad (3)$$

Proxying $E(SSR_n)$ with $\sum_{j=1}^M y_{j,n}^2$ and using the square of denominator from (2) as the variance of $\log SSR_n$, we estimate the standard error of the SSR estimator as:

$$Std.error(SSR_n) = \sqrt{\left\{ \sum_{j=1}^M y_{j,n}^2 \right\}^2 \cdot \left(\exp \left[\frac{2 \cdot \sum_{j=1}^M y_{j,n}^4}{3 \cdot \left(\sum_{j=1}^M y_{j,n}^2 \right)^2} \right] - 1 \right)}. \quad (4)$$

3.3 Implementing the mean square error criterion and choice of frequency

Our dataset includes transaction-by-transaction futures prices for the S&P 500 (1983-2001) and FTSE 100 (1992-2001) equity indices and for short sterling (1988-2001) and eurodollar (1985-2001) interest rate futures. Trading hours for the exchanges concerned varied during the samples. In the US, the Chicago Mercantile Exchange (CME) S&P 500 futures traded between 08:30 – 15:30, and eurodollar futures between 07:20 – 14:00. In the UK, trading hours at London International Futures and Options Exchange (LIFFE) for FTSE 100 futures were 08:30-17:30 for pre-September 1999 period and 08:00-17:30 in the post-1999 period; while short sterling futures traded from 08:20-16:02 up to March 1990, between 08:20-18:00 to March 1990, 08:00-18:00 to October 1991 and finally changing to 07:30-18:00 in July 2001. The raw data obtained from LIFFE (UK) and the Institute for Financial Markets (US) were run through filters to clean the data and remove implausibly large observations due to

errors in the recording of the transaction data. Time series of prices at each intra-day frequency, between the relevant exchange opening and closing times, for each asset were then constructed by sampling the filtered data at each frequency every trading day between a given trade date and the corresponding option expiry date. The return series were then obtained by taking the first difference of the logarithm of the demeaned price series at each frequency, for a given number of days to the expiry date.

Average realised variances using the SSR and VARHAC estimators, across all ten intra-day frequencies, are plotted in Charts A2.1-4 for the four assets. These averages are based on a 'trimmed' sample whereby 5% - that is, upper and lower 2.5% - of the total number of observations were disregarded.

The profiles of SSR realised variances over the intra-day frequencies clearly indicate the importance of microstructural effects in estimating realised volatility in very high frequency data. For almost all assets, the SSR estimator produces the highest estimates of realised variance at the highest –1-minute – frequency with subsequent estimates declining as the frequency decreases. It is clear from these charts that estimating realised variance by arbitrarily selecting an intra-day frequency is ill-advised. The profile of most SSR variance estimates appears to flatten at around 45-60 minutes. The rate of decline is particularly marked in the case of short sterling and eurodollar, relative to that for the two equity indices. The declining patterns of SSR realised variance for short sterling and eurodollar returns suggests greater degrees of negative serial correlation as frequency increases. This finding is consistent with what we would expect to see for a highly liquid asset where negative serial correlation could be induced by the bid-ask bounce, as posited by Roll (1984).

The shape of the equity index SSR variance profiles over various frequencies suggests that futures returns on these indices may be less affected, on average, by this kind of negative serial correlation. It is possible to posit a number of potential explanations for why this may be so. Equity index futures are less traded than interest rate futures so that the bid ask-bounce effect could be weaker.⁷ In addition, futures on equity

⁷ Evidence about the relative turnovers of equity and interest rate futures contracts can be found in BIS Quarterly Review, June 2003. For example, for 2002, equity futures turnover was approximately 5% of turnover in the interest rate futures.

indices may be subject to other types of microstructural features. For example, equity indices may suffer from a non-synchronous trading problem whereby a market-wide shock is not reflected in all individual constituent equity prices simultaneously. As a result, the shock may impact on the index value, which is the weighted-value of the constituent prices, on a more gradual basis as individual constituent equities are subsequently traded. The working of orders in financial markets could also serve to induce a positive serial correlation in equity index returns. In this case, large orders cannot be filled in one attempt as market makers can specify a volume limit with each bid/ask quote. This means that such orders must be worked through the market, starting with the best quote and continuing until the order is filled. It is not likely that this would be completed simultaneously so that the price information for the index in the large order is reflected gradually and not instantaneously. Thus the 'order working' and non-synchronous effects may well be inducing a degree of positive serial correlation in equity index returns that could potentially offset, partially at least, any induced negative serial correlation due to a bid-ask bounce effect.

Turning to the estimated MSEs to formally choose a frequency for each asset, the estimates of the average squared biases, squared standard errors and corresponding MSEs for the SSR estimators for all four assets are presented in Appendix 3, Charts A3.1-4. As before, trimmed means are shown. The charts highlight the relative contributions of the two components in the calculation of the MSE criteria.

For most assets, the squared serial correlation bias is of greater, or at least equal, magnitude to that of the squared standard error of the SSR estimator. As a result the gain, in terms of accuracy of estimation of realised variance, from using intra-day data at increasingly higher frequencies is usually matched or outweighed by the biases that are incurred. Interestingly, the bias profiles, and their relative weight in the MSE criteria, differ substantially between equity index and interest rate futures.

As alluded to in the comparison between the average VARHAC and SSR estimates of realised variance in sub-section 4.1 above, the serial correlation biases for the interest rate futures SSR estimates are much greater than those for the two equity indices. SSR estimates for eurodollar futures were found to have the highest squared biases among all of the assets, exceeding those of the FTSE 100 and S&P 500 by a factor of

up to 160. Generally the interest rate biases were found to be greatest at the 1 minute frequency, declining out to around 45 or 60 minutes. For all frequencies, the interest rate biases dominate the estimated MSEs, thus ruling out very high-frequency data in terms of estimating realised variance using the SSR estimator. The minimum MSE criteria instead suggest frequencies of 60 minutes for short sterling futures, and 45 minutes for eurodollar futures.

For equity index futures, the greatest biases occurred at the 400 minute or daily frequencies, though the FTSE 100 also exhibited a strong bias at one minute intervals. In contrast to the interest rate futures, the equity index squared biases and standard errors were similar in size, suggesting that market microstructure effects are much weaker in equity index futures intra-day data. The equity index futures squared biases are relatively flat between the daily and ultra-high frequencies and the estimated MSE's are minimised at 45 minutes for FTSE 100 and 30 minutes for S&P 500.

Using the chosen frequencies for the four assets, the SSR estimator was then applied to all trading days for all contracts to produce a series that matches that of the forward variance series in terms of time to maturity each trading day, over the entire sample, for each asset. It is the difference between these two – the matched forward and realised variances – that is the main subject of the analysis in the remainder of the paper.

4 Predicting realised with forward variance

In this section we consider the information content of forward for realised variance. Our analyses are conducted by examining coefficients, and their statistical significance, from OLS regressions. However, there are two features in our data that we must first address.

Our variables consist of variances, estimated each day, over a horizon that is determined by the number of days to a fixed option contract expiry date. As a result, the variance horizon declines over successive days, eventually reaching zero at the expiry date. Thus variances estimated for a given contract are overlapping and so highly serially correlated. Moreover, as pointed out by Fleming (1998), the data that

we use is ‘telescopic’ in nature. The telescopic feature arises as instead of looking ahead for a similar fixed period of time, the data are overlapping for the life of an option contract only. As a result, there is no overlap among variance observations across contracts. The presence of these features in the data means that a simple application of OLS would yield inefficient coefficient estimates, with inaccurate standard errors, thus undermining statistical inference. Given the telescopic nature of the data, use of the Newey-West heteroskedastic and autocorrelation consistent (HAC) method of covariance matrix estimation is not appropriate.

To deal with this, Fleming (1998) develops consistent generalised method of moments (GMM) estimators that can produce variance-covariance matrix estimates that are robust to residual serial correlation and conditional heteroskedasticity in the presence of telescoping observations. The Fleming (1998) covariance estimator is an extension of existing HAC methods. The Newey West (1987) HAC covariance estimator generalises the regular least squares covariance matrix estimator to allow for serially correlated errors. Extra cross product terms are included in the calculation of the covariance matrix by weighting observations with the autocovariances of the residuals.⁸ The number of extra cross product terms included is, in theory, determined by the order of autocorrelation of the residuals. The HAC covariance matrix estimator developed by Fleming (1998) limits the autocovariance terms to include only those observations that are common to an individual contract. Using this idea we modify the Newey West HAC covariance estimator to estimate the following HAC covariance matrix with telescopic observations:

$$\mathbf{\Omega} = \frac{1}{N} \sum_{t=1}^N \hat{\varepsilon}_t^2 .x_t' x_t + \sum_{k=1}^L \left(\sum_{i=k+1}^N \phi_k \hat{\varepsilon}_t .\hat{\varepsilon}_{t-i} (x_t' x_{t-i} + \hat{\varepsilon}_t^2 .x_{t-i}' x_t) \right),$$

where, $\hat{\varepsilon}_t$ are the OLS residuals, x_t is the t -th row of a regressor matrix, N is the total number of observations, L is the order of residual autocorrelation (taken to be 90 days) within contracts and ϕ_k takes a value of one if the t -th and $(t-i)$ -th observations relate to the same contract, and a value of zero otherwise. So ϕ_k acts to limit the

⁸ In addition to the autocovariances, the weights on the observations also include another factor to ensure a positive semi-definite covariance matrix.

overlap in observations to those within contracts only. As a result we can estimate a covariance matrix that is robust to the high serial correlation within contracts and so adjust for the telescopic nature of the data. In reporting the standard errors for our coefficient estimates in later regressions, we show both the Newey West standard errors and these ‘telescopic Newey West’ standard errors.

We begin our comparison of realised variance with forward option-implied variance by estimating traditional regressions of realised variance (annualised) on a constant and forward variance (annualised). If forward variance is an unbiased expectation of realised variance, then the constant should be close to zero and the slope coefficient on forward variance close to one. To assess this joint hypothesis, Wald tests are applied with the appropriate restrictions on the coefficients. The results of the regressions and statistical test are presented in Table 4.2. These show that for all assets and all sample periods the constant is close to zero but the slope coefficient is significantly different to one (at 1% significance levels) – generally of magnitude 0.5-0.6. Moreover, the Wald tests fail to accept the joint hypotheses in all cases, confirming the finding in the previous literature that forward variance is a biased expectation of future realised variance. Further, with all of the slope coefficients less than one, this suggests that forward variance, on average, overestimates subsequently realised variance. Nonetheless, forward variance does possess some information about future realised variance with the adjusted R-squared suggesting that it can explain between 22-48% of the variation in realised variance.

To provide an ex-post estimate of the average bias in forward variance, Table 4.1 shows the average differences between the forward and realised variances, for each of the four assets with varying sample periods, together with their standard errors (Newey West and telescopic Newey West adjusted). For the FTSE 100 and S&P 500 equity indices, most of the differences are of similar magnitude; are negative and highly significant with probability values of one percent or less. For the full sample for the S&P 500 we were, however, unable to reject the hypothesis that the difference between realised and forward variances is equal to zero. That this sample (1983-2001) contrasts with later S&P 500 samples (1988-2001 for example) is likely to be due to the large influence of the 1987 Crash. As the Crash was not anticipated by financial markets, realised variance exceeded that which would have been expected to

occur. It is this substantial underprediction of variance, in contrast to the usual overprediction that results in the failure to reject the null hypothesis of zero mean difference for the full S&P 500 sample.

Turning to interest rates, the average differences are again mostly negative. As for the equity indices, forward variance was found, on average, to exceed subsequently realised variance. Only in the period between 1988 and 2001, was the difference found to be statistically insignificant for short sterling. This period includes the ERM crisis and, as with the S&P 500 and the 1987 Crash, there was a large underprediction of volatility at that time.

The averages reported in Table 4.1(a), for both equities and interest rates, could be considered as ex-post measures of annualised variance risk premia, if expectational errors average out over the sample periods. The size of the (annualised) variance risk premium for interest rates (when it is significant) is about 0.005, half of that for equities (≈ 0.01).

Mapping these numbers for the variance risk premium into volatility (i.e. standard deviation) space is not straightforward. In doing so one has to take into account convexity adjustments which themselves, as functions of variance, are unlikely to be constant. To provide an indication of the potential magnitudes of the risk premia in volatility space we obtain an estimate of forward volatility by applying an adjustment to the square root of forward variance for convexity, as explained on page 10. Table 4.1(b) shows the mean differences between these 'convexity-adjusted' forward volatilities and the corresponding maturity-matched realised volatility. Again, as for the differences in variance space, the differences are all negative and, excepting the full sample for short sterling which includes the ERM crisis period, statistically significant. The magnitude of the ex-post volatility risk premia range from 2-3% for equity indices to 1-2% for interest rates. The difference in the sizes of the volatility risk premia between equity indices and interest rates may reflect higher equity index volatility, on average. For example, the average forward volatility over the sample 1988-2001 for the S&P 500 is 18% compared to 15% for eurodollar.

Comparing these risk premia estimates with the differences between the convexity-adjusted forward volatility and a Black-Scholes measure of implied volatility in Table 2.1, suggests that the use of Black-Scholes volatility would imply much lower volatility risk premia. This highlights the importance of measuring forward variance accurately using information from the whole cross-section of option prices as opposed to just the at-the-money contract price.

Next, we consider the efficiency of forward variance in predicting future realised variance by including a measure of backward-looking variance in our regressions. Some previous studies, such as Chernov (2001), also used historical variance as a proxy for the potential bias due to volatility risk premia.

Table 4.3 shows results from an extension of the regressions of realised on forward variance to include backward variance. The measure of backward variance is based on sums of squared historical intra-day returns of the same frequency as that used for realised variance. As with realised variance, the backward variance is estimated over a period that matches the horizon of the forward variance, that is, the time to expiry of the corresponding option contract. So with T days to expiry, backward variance is calculated using squared intra-day returns from the most recent T trading days, i.e. those days that are between $2T$ and $T+1$ days to expiry.

The results show that backward variance is statistically insignificant for all assets and sample periods. This is consistent with the more recent literature (e.g. Fleming (1998), Christensen and Prabhala (1998), Blair et al. (2001)). A corollary of our results is that the backward variance is unlikely to act as a proxy for variance risk premium.

Finally, in a further test of efficiency, we use some additional information from option prices. In particular, we use a measure of asymmetry (proxied by standardised risk reversal) and fatness of tails (proxied by standardised strangle).^{9, 10} Table 4.4 shows

⁹ Standardised risk reversal is given by the difference between the 25-delta call and 75-delta call implied volatilities, divided by the atm (50-delta) volatility. It reflects the slope of the volatility smile. Since it is divided by the atm volatility it adjusts for changes in uncertainty. A lognormal pdf has a risk reversal equal to zero, that is, the risk reversal shows the asymmetry of the implied pdf in excess of the benchmark lognormal pdf.

that there is little value to be added to forward variance in predicting realised variance by using the risk reversal and strangle and backward variance as additional variables (Table 4.4). As a result, the bias in forward variance is unlikely to be related to historical variance, asymmetry or “fat” tails in market expectations. Overall, the results in Tables 4.3-4 show that we cannot reject the hypothesis that forward variance is an efficient forecast of realised variance.

5 *Conclusions*

The issue of the relation between realised and forward (i.e. risk-neutral expectation) of variance has been extensively studied in the literature. The results have been mixed depending on the measures used, assets examined, sample periods covered and econometric techniques employed. We add to this literature by using a theoretically more consistent measure of forward variance; a measure of realised variance based on high frequency data; and by re-examining the bias and efficiency of forward as a predictor of realised variance.

We exploit the result of Britten-Jones and Neuberger (2000) to obtain a measure of risk-neutral expected variance based a static trading position in European call and put options. This measure is model independent in that it does not require any assumptions to be made about the dynamics of the underlying asset or its volatility. As such it is a superior measure to the Black-Scholes implied volatility typically used in previous studies. In addition, it is also a more theoretically consistent measure than the VIX, which uses an ad hoc weighting scheme to combine implied volatilities of different strikes or maturities.

High frequency intra-day data can improve the measurement of realised volatility by increasing its efficiency. However, microstructural effects at high frequencies can induce autocorrelations that create biases in the sum of squared returns (SSR)

¹⁰ Standardised strangle is given by the difference between the average of the 25- and 75-delta call implied volatilities and the atm volatility, divided by the atm volatility. It provides a measure of the degree of curvature of the volatility smile. Since it is divided by the atm volatility it adjusts for changes in uncertainty. A lognormal pdf has a strangle equal to zero, that is, the strangle shows the degree of ‘fatness’ of tails of the implied pdf in excess of the benchmark lognormal pdf.

volatility estimator. To deal with this, we derive a mean square error (MSE) criterion – comprised of bias and standard error of the volatility estimator - that allows for an objective choice of frequency. The bias term in the MSE is estimated by implementing the VARHAC methodology of den Haan and Levin (1996) and Bollen and Inder (2002). Recent additions to the asymptotic theory on realised volatility by Barndorff-Nielsen and Shephard (2003) were employed in estimating the standard error term.

The MSE criterion allows for an explicit analysis of the tradeoff between greater efficiency and the bias due to serial correlation empirically observed at high frequencies. This analysis presented a number of interesting results on market microstructure features in both equity and interest rate markets. We found strong evidence of negative serial correlation in interest rate futures markets, possibly induced by bid-ask bounce. In contrast, in equity futures markets there is little evidence of significant microstructure effects. This could be due to much lower turnover in equity futures relative to interest rate futures so that the bid-ask bounce effect is much weaker. It may also be due to offsetting positive autocorrelations due to non-synchronous trading in equity indices or ‘order working’.

These measures are then used in OLS regressions to address the information content of forward variance in predicting realised variance. Because of the unique (‘telescopic’) overlapping nature of options data, we modify the Newey-West heteroskedastic and autocorrelation consistent (HAC) method of covariance matrix estimation in line with Fleming (1998), to calculate robust standard errors.

The results of these regressions show that our measure of forward variance is a biased but nonetheless informative forecast of realised variance. It is found to be an efficient predictor in the sense that historical variance (also based on intra-day data) and measures of asymmetry and ‘fatness’ of tails implied from option prices were *not* found to contribute to the predictive power of forward variance. As a result, historical variance or measures of asymmetry or fatness of tails are unlikely to be proxies for variance risk premia.

Assuming that expectational errors cancel out over the sample periods examined, our results imply a significant negative variance risk premium (i.e. forward overpredicts realised variance). Comparing asset markets, the magnitude of the average variance risk premium in equity markets is around twice that in short-term interest rate markets. In volatility space, subject to our proxy for a convexity adjustment to forward volatility, volatility risk premia were found to average around 1-2% for interest rates and 2-3% for equity indices.

References

- Akaike, H. (1973), 'Information Theory and an Extension of the Maximum Likelihood Principle', *Second International Symposium on Information Theory*, B.N. Petrov and F. Csaki, eds., Akademia Kiado (Budapest), pp 267-281.
- Andersen, T. G. T., T. Bollerslev, F. X. Diebold and P. Labys (2000), 'Great Realisations', *Risk*, Vol. 13, pp 105-108.
- Andersen, T. G. T., T. Bollerslev, F. X. Diebold and P. Labys (2001), 'The Distribution of Realized Stock Return Volatility', *Journal of Financial Economics*, Vol. 61, pp 43-76.
- Barndorff-Nielsen, O. and N. Shephard (2002), 'Econometric Analysis of Realised Volatility and its Use in Estimating Stochastic Volatility Models', *Journal of the Royal Statistical Society, Series B*, volume 64, 2002, pp 253-280.
- Barndorff-Nielsen, O. and N. Shephard (2003), 'How accurate is the asymptotic approximation to the distribution of realised volatility?' forthcoming in *Identification and Inference for Econometric Models. A Festschrift for Tom Rothenberg*, (edited by Don Andrews, Jim Powell, Paul Ruud and Jim Stock), Econometric Society Monograph Series, Cambridge University Press, 2003.
- Beckers S. (1979), 'Standard Deviations Implied in Option Prices As Predictors of Future Stock Price Variability', *Journal of Banking and Finance*, Vol. 5, pp 363-81.
- BIS (Bank for International Settlements) Quarterly Review*, June 2003.
- Blair B.J., Poon S.H. and S.J. Taylor, 'Forecasting S&P 100 Volatility: The Incremental Information Content of Implied Volatilities and High-Frequency Index Returns', *Journal of Econometrics*, Vol. 105, pp 5-26.
- Bliss, R. R. and N. Panigirtzoglou (2002), 'Testing the stability of implied probability density functions', *Journal of Banking and Finance* 26(2-3), 381-422.
- Bollen B. and B. Inder (2002), 'Estimating daily volatility in financial markets utilising intraday data', *Journal of Empirical Finance*, Vol. 9, pp 551-62.
- Britten-Jones M. and A. Neuberger (2000), 'Option Prices, Implied Price Processes and Stochastic Volatility', *Journal of Finance*, Vol. 55, No 2, pp 839-66.
- Canina, L. and S. Figlewski (1993), 'The Informational Content of Implied Volatility', *The Review of Financial Studies*, Vol. 6, pp 659-81.
- Carr, P. and D. Madan (1999), 'Introducing the covariance swap', *Risk magazine*, February.
- Chernov M. (2001), 'Implied Volatilities as Forecasts of Future Volatility, the Market Risk Premia, and Returns Variability', *mimeo*, Columbia University.

- Christensen, B.J. and N.R. Prabhala (1998), 'The Relation Between Implied and Realised Volatility', *Journal of Financial Economics*, Vol. 50, pp 125-50.
- Day, T. and C. Lewis (1992), 'Stock Market Volatility and the Information content of Stock Index Options', *Journal of Econometrics*, Vol. 52, pp 267-87.
- den Haan, W. J. and A. T. Levin (1996), 'A Practitioner's Guide to Robust Covariance Matrix Estimation', *NBER Working Paper*, No. 197.
- Demeterfi K., Derman, E., Kamal, M. and J. Zou (1999), 'A guide to variance swaps', *Risk magazine*, June.
- Fleming J. (1998), 'The Quality of Market Volatility Forecasts Implied by S&P 100 Index Option Prices', *Journal of Empirical Finance*, Vol. 5, pp 317-54.
- Feinstein, S. (1989), 'The Black-Scholes formula is nearly linear in σ for at-the-money options; therefore implied volatilities from at-the-money options are virtually unbiased'. Federal Reserve Bank of Atlanta Working Paper.
- Hull J. and A. White (1987), 'The Pricing of Options on Assets with Stochastic Volatilities', *Journal of Finance*, Vol. 42, pp 281-301.
- Jorion P. (1995), 'Predicting Volatility in the Foreign Exchange Market', *The Journal of Finance*, Vol. 50, pp 507-28.
- Lamoureux, D. and W. Lastrapes (1993), 'Forecasting Stock-Return Variance: Toward an Understanding of Stochastic Implied Volatilities', *Review of Financial Studies*, Vol. 6, pp 293-326.
- Latane, H. and R. Rendleman (1976), 'Standard Deviations of Stock Price Ratios Implied in Option Prices', *Journal of Finance*, Vol. 31, pp 369-81.
- Lynch D. and N. Panigitzoglou (2003), 'Summary Statistics of Option-Implied Pdfs and Their Properties', *Bank of England Working Paper* forthcoming.
- Newey, W. and K. West (1978), 'A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix', *Econometrica*, Vol. 55, pp 703-708.
- Newey, W. and K. West (1994). 'Automatic Lag Selection in Covariance Matrix Estimation', *Review of Economic Studies*, Vol. 61, pp 631-653.
- Phillips, P.C.B. and P. Perron (1988), 'Testing for a Unit Root in Time Series Regression', *Biometrika*, Vol. 75, pp 335-346.
- Roll, R. (1984), 'A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market', *Journal of Finance*, Vol. 39, pp 1127-1139.

Schwartz, G. (1978), 'Estimating the Dimension of a Model', *Annals of Statistics*, Vol. 6, pp 461-464.

Xu, X and S.J. Taylor (1995), 'Conditional Volatility and the Informational Efficiency of the PHLX Currency Options Market', *Journal of Banking and Finance*, Vol. 19, pp 803-21.

Taylor S.J. and X. Xu (1997), 'The incremental information in one million foreign exchange quotations', *Journal of Empirical Finance*, Vol. 4, pp 317-340.

APPENDIX 1: convexity adjustment

The convexity adjustment of equation (1) when working with volatility instead of variance can be derived as follows. Lets call x the volatility and $y = x^2$ the variance of the asset:

$$x = \sqrt{\int_t^T \left(\frac{dS_t}{S_t} \right)^2} \quad \text{and} \quad y = x^2 = \int_t^T \left(\frac{dS_t}{S_t} \right)^2$$

Equation (1) can then be expressed as

$$E_t[y] = 2 \cdot e^{r(T-t)} \cdot \left(\int_0^{S_t} \frac{P_t(T, K)}{K^2} dK + \int_{S_t}^{\infty} \frac{C_t(T, K)}{K^2} \right)$$

Consider a Taylor expansion of the variable $x = \sqrt{y}$ around the point $y_0 = E[y]$:

$$x = \sqrt{y} = \sqrt{y_0} + (y - y_0) \frac{1}{2\sqrt{y_0}} + \frac{1}{2} (y - y_0)^2 \left(-\frac{1}{4} y_0^{-\frac{3}{2}} \right) + \dots$$

We take the conditional (at time t) expectation of both sides of the above equation:

$$E_t[\sqrt{y}] = \sqrt{y_0} + (E_t[y] - y_0) \frac{1}{2\sqrt{y_0}} + \frac{1}{2} E_t[(y - y_0)^2] \left(-\frac{1}{4} y_0^{-\frac{3}{2}} \right) + \dots$$

Given that $y_0 = E_t[y]$ we can get

$$E_t[\sqrt{y}] = \sqrt{E_t[y]} - \frac{1}{8} \text{Var}_t[y] \cdot (E_t[y])^{-\frac{3}{2}} + \dots$$

The above equation shows that the expectation of volatility \sqrt{y} is less than the square root of the expectation of variance by the convexity adjustment $-\frac{1}{8} \frac{\text{Var}_t[y]}{E_t[y]^2} \cdot \sqrt{E_t[y]}$.

The quantity $\frac{\text{Var}_t[y]}{E_t[y]^2}$ is the ratio of variance with respect to the mean squared, that is,

a scaled measure of the variance of the asset. If that scaled measure remains relative constant over time, we would expect the convexity adjustment to be proportional to uncertainty as measured by $\sqrt{E_t[y]}$. Table 1.1 provides average estimates of the volatility convexity adjustment.

APPENDIX 1 continued

Table 1.1 Average volatility convexity adjustment

sample			sample		
S&P 500			Eurodollar		
1st contract	83-01	-0.17%	1st contract	85-01	-0.15%
FTSE 100			Short sterling		
1st contract	92-01	-0.16%	1st contract	88-01	-0.29%

APPENDIX 2: Differences between convexity-adjusted forward volatility and Black-Scholes at-the-money volatility

Chart A2.1 S&P 500 first contract




Chart A2.2 FTSE 100 first contract




Chart A2.3 Eurodollar first contract




Chart A2.4 Short sterling first contract




Table 2.1 Average absolute percentage point difference between convexity adjusted forward volatility and Black-Scholes at-the-money volatility

sample			sample		
S&P 500			Eurodollar		
1st contract	83-01	1.46	1st contract	85-01	1.54
FTSE 100			Short sterling		
1st contract	92-01	1.50	1st contract	88-01	0.80

APPENDIX 3

Chart A3.1 S&P 500 realised variance measures (averages[†])




Chart A3.2 FTSE 100 realised variance measures (averages[†])




Chart A3.3 Eurodollar realised variance measures (averages[†])

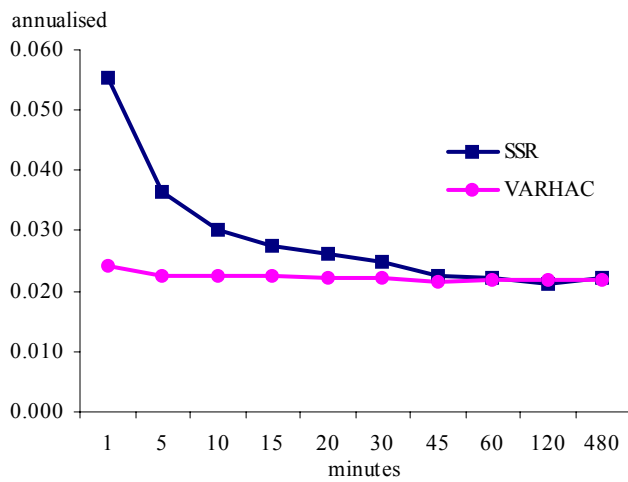



Chart A3.4 Short sterling realised variance measures (averages[†])



[†] averages are trimmed by excluding the upper and lower 2.5% of the observations

APPENDIX 3 continued

Chart A3.5 S&P 500 frequency choice for realised variance (averages[†])

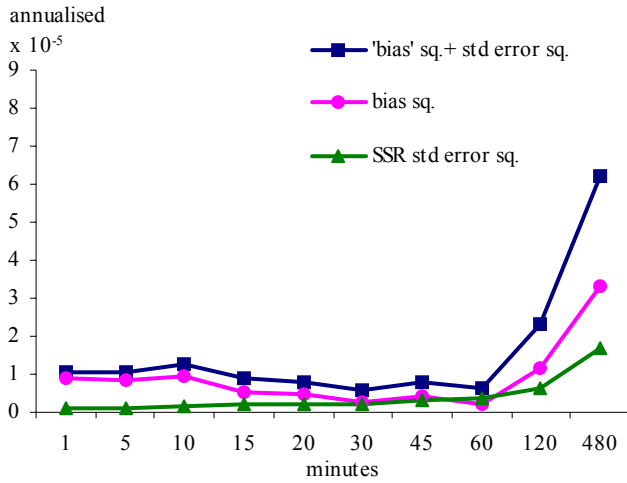


Chart A3.6 FTSE 100 frequency choice for realised variance (averages[†])

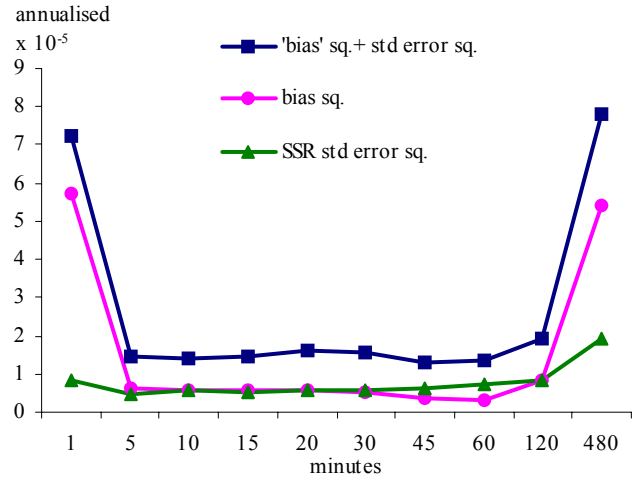


Chart A3.7 Eurodollar frequency choice for realised variance (averages[†])





Chart A3.8 Short sterling frequency choice for realised variance (averages[†])



[†] averages are trimmed by excluding the upper and lower 2.5% of the observations

APPENDIX 3 continued

Chart A3.9 S&P 500 realised variance at daily and optimal intra-day frequencies (first contract)

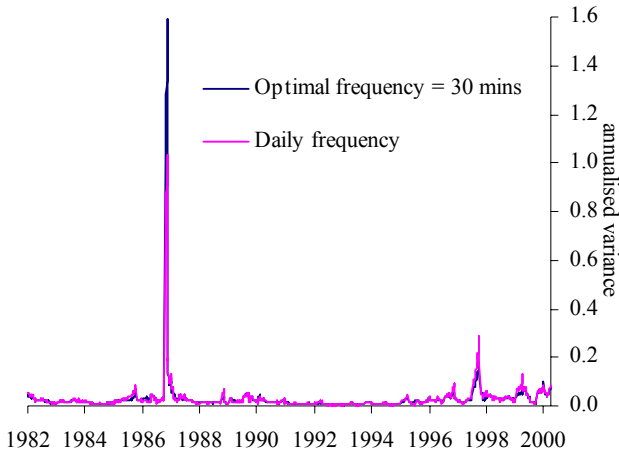


Chart A3.10 FTSE 100 realised variance at daily and optimal intra-day frequencies (first contract)

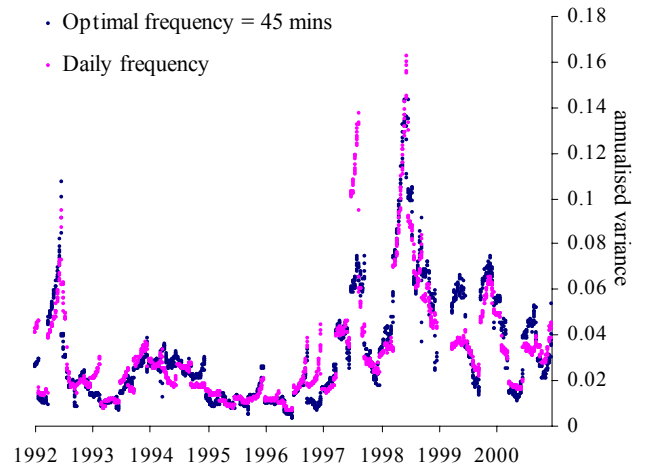


Chart A3.11 Eurodollar realised variance at daily and optimal intra-day frequencies (first contract)

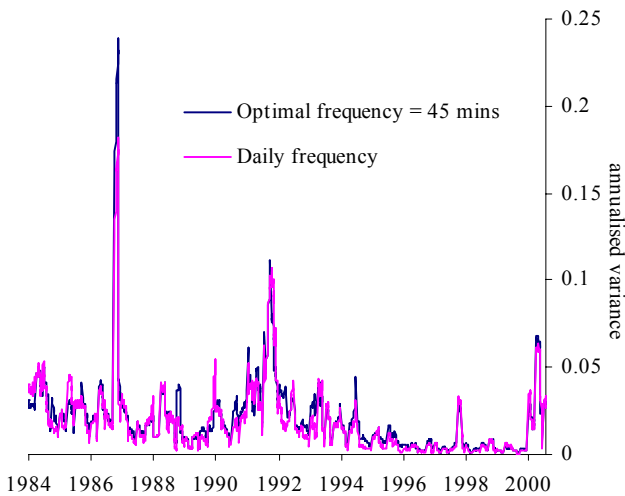
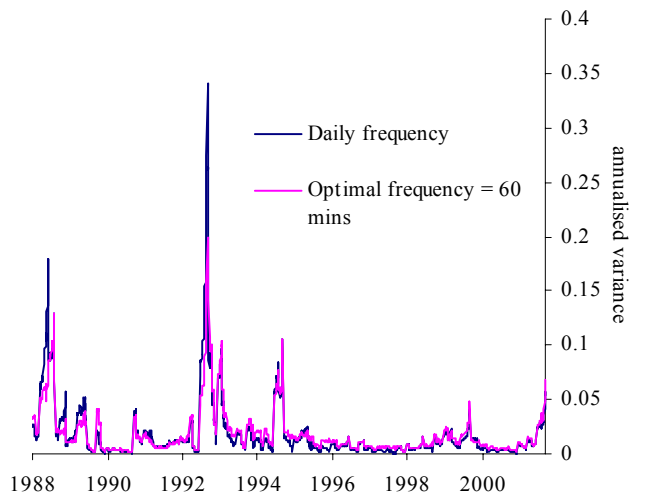


Chart A3.12 Short sterling realised variance at daily and optimal intra-day frequencies (first contract)



APPENDIX 4: differences between forward and realised variances

Table 4.1a: realised variance - forward variance = a

	sample	a	SE(a) [†]		sample	a	SE(a) [†]
S&P 500	83-01	-0.006	4.0E-03	Eurodollar	85-01	-6.4E-03**	1.1E-03
1st contract			5.8E-03	1st contract			1.7E-03
S&P 500	88-01	-0.014**	1.0E-03	Eurodollar	85-01	-5.5E-03**	1.0E-03
1st contract			1.9E-03	2nd contract			2.0E-03
S&P 500	92-01	-0.011**	1.1E-03	Short sterling	88-01	1.4E-04	1.2E-03
1st contract			2.3E-03	1st contract			4.2E-04
FTSE 100	92-01	-0.011**	1.6E-03	Short sterling	93-01	-2.6E-03**	1.0E-03
1st contract			3.3E-03	1st contract			3.5E-04
				Short sterling	88-01	-5.5E-05	1.3E-03
				2nd contract			4.4E-04
				Short sterling	93-01	-5.1E-03**	7.7E-04
				2nd contract			2.7E-04

[†] two standard errors are shown for each coefficient: the first is Newey-West heteroskedasticity and autocorrelation adjusted and the second is adjusted for the telescopic nature of the data.

* and ** denote significance at 5% and 1% levels respectively.

Table 4.1b: realised volatility - convexity adjusted forward volatility = a

	sample	a	SE(a) [†]		sample	a	SE(a) [†]
S&P 500	83-01	-0.029**	3.7E-03	Eurodollar	85-01	-0.021**	2.3E-03
1st contract			5.6 E-03	1st contract			3.8E-03
S&P 500	88-01	-0.038**	2.1E-03	Short sterling	88-01	-0.008	2.0E-03
1st contract			3.8E-03	1st contract			4.3E-03
S&P 500	92-01	-0.033**	2.6E-03	Short sterling	93-01	-0.010*	3.1E-03
1st contract			4.8E-03	1st contract			4.7E-03
FTSE 100	92-01	-0.025**	3.1E-03				
1st contract			5.6E-03				

[†] two standard errors are shown for each coefficient: the first is Newey-West heteroskedasticity and autocorrelation adjusted and the second is adjusted for the telescopic nature of the data.

* and ** denote significance at 5% and 1% level respectively.

As outlined on page 10, the convexity adjustment is estimated approximately and as a result the differences shown in this table between realised and forward volatility are subject to measurement error.

APPENDIX 4 continued: unbiasedness and orthogonality regressions

Table 4.2: *realised_variance = a + b.forward_variance*

	sample	<i>a</i>	SE(<i>a</i>) [†]	<i>b</i> [‡]	SE(<i>b</i>) [†]	R ² adj.	Wald [§] <i>H</i> ₀ : <i>a</i> =0, <i>b</i> =1	PP ^{§§} <i>H</i> ₀ : Unit root in residuals
S&P 500	83-01	0.011*	0.004	0.527**	0.099	0.03	12.13	-7.80**
1st contract			0.006		0.076		0.00	
S&P 500	88-01	0.005	0.002	0.500**	0.061	0.39	169.80	-4.19**
1st contract			0.003		0.103		0.00	
S&P 500	92-01	0.003	0.002	0.618**	0.076	0.48	67.88	-3.52**
1st contract			0.004		0.136		0.00	
FTSE 100	92-01	0.014**	0.002	0.434**	0.038	0.42	115.55	-5.78**
1st contract			0.003		0.070		0.00	
Eurodollar	85-01	0.007**	0.001	0.503**	0.050	0.24	51.60	-7.31**
1st contract			0.002		0.068		0.00	
Eurodollar	85-01	0.009**	0.001	0.546**	0.048	0.26	46.16	-5.23**
2nd contract			0.003		0.084		0.00	
Short sterling	88-01	0.006**	0.002	0.678**	0.118	0.22	7.67	-5.56**
1st contract			0.002		0.112		0.00	
Short sterling	93-01	0.006**	0.0014	0.520**	0.088	0.29	19.52	-8.30**
1st contract			0.0016		0.084		0.00	
Short sterling	88-01	0.009**	0.002	0.610**	0.080	0.17	14.59	-3.71**
2nd contract			0.003		0.105		0.00	
Short sterling	93-01	0.007**	0.001	0.445**	0.042	0.46	129.12	-4.38**
2nd contract			0.0013		0.049		0.00	

[†] two standard errors are shown for each coefficient: the first is Newey-West heteroskedasticity and autocorrelation adjusted and the second is adjusted for the telescopic nature of the data.

[‡] significance levels relate to test for null hypothesis that *b* = 1. * and ** denote significance at 5% and 1% level respectively.

[§] F-statistics with corresponding probability values underneath are shown for each Wald test.

^{§§} PP shows the Phillips-Perron (1998) test statistics for the null hypothesis of a unit root in the residuals of the corresponding regression. The test were performed using a spectral density estimation based using a Bartlett kernel function. The bandwidth for each test was chosen using the Newey-West (1994) data-based bandwidth parameter selection rule. * and ** denote significance at 5% and 1% level respectively.

APPENDIX 4 continued: unbiasedness and orthogonality regressions

Table 4.3: realised variance = a + b.forward variance + c.backward variance

	sample	a	SE(a) [†]	b [§]	SE(b) [†]	c	SE(c) [†]	R ² adj.
S&P 500	83-01	0.008*	0.003	0.636*	0.124	-0.037	0.015	0.03
1st contract			0.004		0.155		0.023	
S&P 500	88-01	0.005	0.002	0.479**	0.065	0.029	0.026	0.39
1st contract			0.003		0.105		0.027	
S&P 500	92-01	0.003	0.002	0.552**	0.109	0.094	0.098	0.48
1st contract			0.004		0.168		0.119	
FTSE 100	92-01	0.015**	0.002	0.474**	0.051	-0.08	0.059	0.42
1st contract			0.003		0.097		0.090	
Eurodollar	85-01	0.007**	0.001	0.486**	0.074	0.017	0.048	0.24
1st contract			0.002		0.119		0.072	
Short sterling	88-01	0.007**	0.002	0.589**	0.142	0.064	0.067	0.23
1st contract			0.002		0.129		0.045	
Short sterling	93-01	0.005**	0.0015	0.371**	0.093	0.180	0.090	0.32
1st contract			0.0016		0.085		0.110	

[†] two standard errors are shown for each coefficient: the first is Newey-West heteroskedasticity and autocorrelation adjusted and the second is adjusted for the telescopic nature of the data.

* and ** denote significance at 5% and 1% level respectively.

[§] significance levels relate to test for null hypotheses that $b = 1$.

Table 4.4: realised variance = a + b.forward variance + c.abs(risk reversal) + d.strangle + e.backward variance

	sample	a	SE(a) [†]	b [§]	SE(b) [†]	c	SE(c) [†]	d	SE(d) [†]	e	SE(e) [†]	R ² adj.
S&P 500	83-01	0.045	0.020	0.751	0.165	-1.1E-03	5.2E-04	-6.2E-03	3.2E-03	-0.041	0.017	0.07
1st contract			0.033		0.207		8.3E-04		5.2E-03		0.022	
S&P 500	88-01	0.019*	0.005	0.476**	0.070	-3.2E-04*	8.7E-05	-1.7E-03	7.2E-04	0.027	0.025	0.42
1st contract			0.009		0.114		1.5E-04		1.3E-03		0.025	
S&P 500	92-01	0.019	0.008	0.619*	0.113	-6.3E-04**	1.3E-04	-3.2E-05	1.0E-03	0.060	0.090	0.51
1st contract			0.014		0.180		2.4E-04		1.9E-03		0.106	
FTSE 100	92-01	0.018**	0.003	0.442**	0.058	3.7E-05	1.5E-04	-1.7E-03	8.1E-04	-0.08	0.056	0.42
1st contract			0.006		0.098		2.1E-04		1.5E-03		0.083	
Eurodollar	85-01	0.013**	0.002	0.455**	0.069	-1.2E-04	5.5E-05	-6.0E-04*	1.9E-04	0.008	0.043	0.26
1st contract			0.004		0.102		7.9E-05		3.0E-04		0.068	
Short sterling	88-01	0.011**	0.002	0.625**	0.148	-3.0E-04	1.5E-04	-1.0E-04	2.5E-04	0.042	0.067	0.24
1st contract			0.004		0.139		2.3E-04		3.8E-04		0.049	
Short sterling	93-01	0.004*	0.002	0.36**	0.098	8.6E-05	1.2E-04	9.2E-05	1.9E-04	0.184	0.091	0.32
1st contract			0.002		0.094		1.2E-04		2.5E-04		0.110	

[†] two standard errors are shown for each coefficient: the first is Newey-West heteroskedasticity and autocorrelation adjusted and the second is adjusted for the telescopic nature of the data.

* and ** denote significance at 5% and 1% level respectively.

[§] significance levels relate to null hypotheses that $b = 1$.