The incremental volatility information in one million foreign exchange quotations

Stephen J. Taylor [a,*], Xinzhong Xu [b]

[a] *Department of Accounting and Finance, Lancaster University, Lancaster, LA1 4YX, UK*
[b] *Department of Accounting and Finance, University of Manchester, Manchester, M13 9PL, UK*

Abstract

The volatility information found in high-frequency exchange rate quotations and in implied volatilities is compared by estimating ARCH models for DM/$ returns. Reuters quotations are used to calculate five-minute returns and hence hourly and daily estimates of realised volatility that can be included in equations for the conditional variances of hourly and daily returns. The ARCH results show that there is a significant amount of information in five-minute returns that is incremental to options information when estimating hourly variances. The same conclusion is obtained by an out-of-sample comparison of forecasts of hourly realised volatility. © 1997 Elsevier Science B.V.

JEL classification: G13; G14; G15

Keywords: ARCH models; Exchange rates; High-frequency data; Options; Volatility

1. Introduction

The volatility of a spot exchange rate S can be defined for many price models by the annualised standard deviation of the change in the logarithm of S during some time interval. For a diffusion process defined by $d(\ln S) = \mu \, dt + \Psi(t)dW$,

---

* Corresponding author. Tel.: +44-1524-593624; fax: +44-1524-847321; e-mail: s.taylor@lancaster.ac.uk.

with $\Psi(t)$ a deterministic function of time and $W(t)$ a standard Wiener process, the deterministic volatility $\sigma(0, T)$ from time 0 until time T is defined by

$$\sigma^2 = \frac{1}{T} \mathrm{var}(\ln S(T) - \ln S(0)).$$

Options traders make predictions of volatility for several values of T. These forecast horizons typically vary between a fortnight and a year and are defined by the times until expiration of the options traded. Insights into these predictions can be obtained by inverting an option pricing formula to produce implied volatility numbers for various values of T. Xu and Taylor (1994) show that these volatility expectations vary significantly for exchange rates, both across expiry times T and through time.

Options markets are often considered to be markets for trading volatility. It then follows that implied volatilities are likely to be good predictors of subsequent observed volatility if the options market is efficient. As options traders have more information than the historic record of asset prices it may also be expected that implied volatilities are better predictors than forecasts calculated from recent prices using ARCH models.

Day and Lewis (1992) investigate the information content of implied volatili-ties, calculated from call options on the S&P 100 index, within an ARCH framework. They conclude that recent stock index levels contain incremental volatility information beyond that revealed by options prices. Lamoureux and Lastrapes (1993) report a similar conclusion for individual US stocks. Xu and Taylor (1995), however, use daily data to conclude that exchange rates do not contain incremental volatility information: implied volatility predictions cannot be improved by mixing them with conditional variances calculated from recent exchange rates alone. Jorion (1995) also finds that daily currency implieds are good predictors.

The superior efficiency of currency implieds relative to implieds calculated from spot equity indices has at least two credible explanations. First, there is the theoretical argument of Canina and Figlewski (1993) that efficiency will be enhanced when fast low-cost arbitrage trading is possible. S&P 100 index arbitrage, unlike forex arbitrage, is expensive because many stocks must be traded. Second, as Jorion (1995) observes, index option implieds can suffer from substan-tial measurement error because of the presence of some stale quotes in the index.

This paper extends the study of Xu and Taylor (1995), hereafter XT, by using high-frequency exchange rates to extract more volatility information from the historical record of exchange rates. From probability theory it is known that it may be possible to substantially improve volatility estimates by using very frequent observations. Nelson (1992) shows that it is theoretically possible for volatility estimates to be made as accurate as required for many diffusion models by using ARCH estimates and sufficiently frequent price measurements. As trading is not

continuous and bid/ask spreads exist, there are, of course, limits to the benefits obtainable from high-frequency data.

Sections 2 and 3 review the definition of implied volatility and the low-frequency results of XT. Section 4 describes our estimates of Deutschemark/dollar volatility obtained from the high-frequency dataset of Olsen & Associates. Section 5 presents the results from estimating ARCH models when the conditional variance is a function of implied volatilities and/or high-frequency volatility estimates. Section 6 provides further evidence about the incremental information content of options prices and the O & A quotations database by evaluating the accuracy of volatility forecasts. Finally, Section 7 summarises our conclusions.

2. Implied volatility

The implied volatilities used in this paper are calculated from the prices of nearest-the-money options on spot currency. These options are traded at the Philadelphia stock exchange (PHLX). Standard option pricing formulae assume the spot rate follows a geometric Brownian motion process. The appropriate European pricing formula for the price c of a call option is then a well-known function of the present spot rate S, the time until expiration T, the exercise price X, the domestic and foreign interest rates, respectively r and q and the volatility σ (see for example Hull, 1995). The Philadelphia options can be exercised early and consequently the accurate approximate formula of Barone-Adesi and Whaley (1987) is used to define the price C of an American call option. This price can be written as

$$C = \begin{cases} c + e, & S < S^*, \\ S - X, & S \geq S^*, \end{cases}$$

with e the early exercise premium and S^* the critical spot rate above which the option should be exercised immediately. The implied volatility is the number σ_I that equates an observed market price C_M with the theoretical price C:

$$C_M = C(S, T, X, r, q, \sigma_I).$$

There will be a unique solution to this equation when $C_M > S - X$. As $\partial C / \partial \sigma > 0$ when $S < S^*(\sigma)$, the solution can be found very quickly by an interval subdivision algorithm. Similar methods apply to put options. A typical matrix of currency implied volatilities calculated for various combinations of time-to-expiry T and exercise price X will display term structure effects as T varies for a fixed X near to the present spot price. These effects have been modelled by assuming mean reversion in implied volatilities (Xu and Taylor, 1994). Matrices of implieds also display smile effects as X varies for fixed T (Taylor and Xu, 1994).

Traders know that volatility is stochastic, nevertheless they make frequent use of implied estimates obtained from pricing models that assume a constant volatil-

ity. The implied volatility can be interpreted as a volatility forecast if we follow the analysis of Hull and White (1987) and make three assumptions: first that the price $S(t)$ and the stochastic volatility $\sigma(t)$ follow diffusion processes, second that volatility risk is not priced and third that spot price and volatility differentials are uncorrelated. The first and second assumptions are pragmatic and the third is consistent with the empirical estimates reported by XT. With these assumptions, let $\bar{V}_T$ be the average variance $(1/T)\int_0^T \sigma^2(t)dt$. Also, let $c(\sigma^2)$ represent the Black–Scholes, European valuation function for a constant level of volatility, σ. Then Hull and White (1987) show the fair European call price is the expectation $E[c(\bar{V}_T)]$, which is approximately $c(E[\bar{V}_T])$ when X is near S. Thus the theory can support a belief that the implied volatility for time-to-expiry T is approximately the square root of $E[\bar{V}_T]$. Traders might obtain efficient prices if they forecast the average variance and then insert its square root into a pricing formula that assumes constant volatility.

3. Low-frequency results

The evidence for incremental volatility information can be assessed by making comparisons between the maximum likelihoods attained by different volatility models. An ARCH model for returns R_t based upon information sets Ω_t will specify a set of conditional variances h_t and hence conditional distributions $R_t | \Omega_{t-1}$, from which the likelihood of observed returns can be calculated. We consider information sets I_t, J_t and K_t respectively defined by (a) all returns up to time t, (b) implied volatilities up to time t and (c) the union of these two sets. We say that an information source has incremental information if it increases the log-likelihood of observed returns by a statistically significant amount.

The following maximum log-likelihoods are reported by Xu and Taylor (1995, Table 3) for a model defined below, for five years (1985–1989) of daily DM/$ returns from futures contracts:

I_t 4327.31
J_t 4349.64
$K_t = I_t + J_t$ 4349.65

Source J_t has incremental information because its addition to I_t adds 22 to the log-likelihood with only one extra parameter included in the ARCH model. This is significant at very low levels. Source I_t, however, does not contain incremental information because its addition to J_t only adds 0.01 to the log-likelihood. Thus, in this low-frequency example, there is only incremental information in options prices.

The models estimated in XT use daily conditional variances h_t that reflect higher levels of volatility for Monday and holiday returns. These seasonal effects

are modelled by multiplicative seasonal parameters, respectively denoted by M and H. The quantity h_t^* represents the conditional variance with seasonality removed: it is defined by

$$
h_t^* = \begin{cases} h_t & \text{if period } t \text{ ends 24 hours after period } t-1, \\ h_t/M & \text{if } t \text{ falls on a Monday and } t-1 \text{ on a Friday,} \\ h_t/H & \text{if a holiday occurs between the two prices.} \end{cases} \tag{1}
$$

A general specification for h_t^* that incorporates information at time $t-1$ about daily returns R_{t-1}, implied volatilities i_{t-1} and their lagged values is given by

$$
h_t^* = c + aR_{t-1}^2(h_{t-1}^*/h_{t-1}) + bh_{t-1}^* + di_{t-1}^2/(196 + 48M + 8H). \tag{2}
$$

The quantity i_{t-1} here denotes the implied volatility for the nearest-the-money call option, for the shortest maturity with more than nine calendar days to expiration. Although i_{t-1} is an expectation for a period of at least ten days it is used as a proxy for the market's expectation for the single trading period t. The standard deviation measure i_{t-1} is an annualised quantity. It is converted to a variance for a 24-hour return in the above equation by assuming there are 48 Mondays, 8 holidays and 196 normal weekdays in a year.

An appropriate conditional distribution for daily returns from DM/\$ futures is the generalised error distribution (GED) that has a single shape parameter, called the thickness parameter ν. The parameter vector for the general specification is then $\theta = (a, b, c, d, M, H, \nu)$. All conditional means are supposed to be zero.

The maximum likelihood for information sets I_t is obtained by assuming $d = 0$ followed by maximisation of the log-likelihood over the remaining parameters. This gives:

$$
h_t = \mathrm{var}(R_t \mid I_{t-1}),
$$
$$
h_t^* = 2.5 \times 10^{-6} + 0.07R_{t-1}^2(h_{t-1}^*/h_{t-1}) + 0.88h_{t-1}^*,
$$
$$
\nu = 1.25, \qquad M = 1.16, \qquad H = 1.50. \tag{3}
$$

The estimate of ν has a standard error less than 0.1 and therefore fat-tailed conditional distributions describe returns more accurately than conditional normal distributions ($\nu = 2$), as has been shown in many other studies of daily exchange rates. The estimates of M and H are more than one but their standard errors, respectively 0.12 and 0.33, are substantial.

The maximum likelihood for options information J_t is obtained when a and b are constrained to be zero and all the other parameters are unconstrained. MLE then gives $c = 0$ and:

$$
h_t = \mathrm{var}(R_t \mid J_{t-1}), \qquad h_t^* = 0.97i_{t-1}^2/(196 + 48M + 8H), \qquad \nu = 1.33. \tag{4}
$$

The incremental importance of previous returns and options information is assessed by estimating the general specification without parameter constraints. The MLE estimates of a and c are zero, with b estimated as 0.04 (t-ratio 1.43) and d as 0.93 (t-ratio 3.11). Any conventional statistical tests accept the null hypotheses $a = 0$, $b = 0$, $c = 0$ and $d = 1$. They also reject $d = 0$ at very low significance levels.

XT conclude that all the relevant information for defining the next period's conditional variance is contained in the most recent implied volatility. This conclusion holds despite using a volatility expectation for at least a ten-day period as a proxy for the options market's expectation for the next trading period. XT also present results for volatility expectations for the next day calculated from a term structure model for implieds studied in Xu and Taylor (1994). These expectations are extrapolations ($T = 1$ day) from several implieds ($T \geq 10$ days). Such extrapolations provide both the same conclusions as short-maturity implieds and very similar maximum levels of the log-likelihood function. However, these extrapolations are biased.

Out-of-sample forecasts of realised volatility during four-week periods in 1990 and 1991 confirm the superiority of the options predictions compared with standard ARCH predictions based upon previous returns alone.

4. Volatility estimates and expectations

4.1. Intra-day data

Estimates of Deutschemark/dollar volatility have been obtained from the dataset of spot DM/\$ quotations collected and distributed by Olsen & Associates. The dataset contains more than 1,400,000 quotations on the interbank Reuters network between Thursday 1 October 1992 and Thursday 30 September 1993 inclusive. It is our understanding that the dataset is an almost complete record of spot DM/\$ quotations shown on Reuters FXFX page. The quotations are time stamped using GMT. We converted all times to US eastern time which required different clock adjustments for winter and summer.

Volatility estimates have been calculated for 24-hour weekday periods for comparison with daily observations of implied volatilities. The options market at the Philadelphia stock exchange closes at 14.30 US eastern time, which is 19.30 GMT in the winter and 18.30 GMT in the summer. A 24-hour estimate for a winter Tuesday is calculated from quotations made between 19.30 GMT on Monday until 19.30 GMT on Tuesday. We follow Andersen and Bollerslev (1997) and ignore the 48 hours from 21.00 GMT on Friday until 21.00 GMT on Sunday, because less than 0.1% of the quotations are made in this weekend period. Thus a 24-hour estimate for a winter Monday uses quotations from 19.30 to 21.00 GMT on the previous Friday and from 21.00 GMT on Sunday until 19.30 GMT on Monday.

4.2. Definition and motivation of the estimates

The realised volatility for day t is calculated from intra-day returns $R_{t,i}$ with i counting short periods during day t, in the following way:

$$v_t = \sqrt{m \sum_{i=1}^{n} R_{t,i}^2} \, . \tag{5}$$

Here m is a multiplicative constant that converts the variance for one trading day into an annual variance and v_t is an annualised measure of realised volatility. The number of short periods in one trading day is chosen to be $n = 288$ corresponding to five-minute returns.

We follow the methods of Andersen and Bollerslev (1997), hereafter AB, when five-minute returns are calculated. Their methods use averages of bid and ask quotations to define rates. They define the rate at any required time by a linear interpolation formula that uses two quotations that immediately precede and follow the required time. As in AB, suspect quotations are filtered out using the methods of Dacorogna et al. (1993). AB note that there is very little autocorrelation in the five-minute returns: the first-lag coefficient is -0.04. Negative dependence has previously been documented by Goodhart and Figliuoli (1991).

Some motivation for the above method of volatility estimation is provided by supposing that spot exchange rates $S(\tau)$ develop in calendar time τ according to a diffusion process described by

$$d(\ln S(\tau)) = \mu \, d\tau + s(\tau) \sigma(\tau) dW(\tau) \tag{6}$$

with $\sigma(\tau)$ an annualised stochastic quantity and $s(\tau)$ a deterministic quantity that reflects the strong intra-day seasonal pattern in volatility. This pattern has been investigated in detail by AB and has been described in earlier studies that include Bollerslev and Domowitz (1993) and Dacorogna et al. (1993). The square of the seasonal multiplier $s(\tau)$ averages one over a complete seasonal cycle, so if τ_1 and τ_2 denote the identical position in the cycle then $s(\tau_1) = s(\tau_2)$ and $\int_{\tau_1}^{\tau_2} s^2(\tau) d\tau = \tau_2 - \tau_1$.

When the volatility is constant during a one-day cycle, of length Δ years, and the multipliers are constants s_i during intra-day intervals, then

$$\sum_{i=1}^{n} \text{var}(R_{t,i} | \sigma(\tau)) = \Delta \frac{1}{n} \sum_{i=1}^{n} s_i^2 \sigma^2(\tau) = \Delta \sigma^2(\tau), \tag{7}$$

with τ the calendar time associated with trading period t. The quantity v_t^2 is the estimate of $\sigma^2(\tau)$ obtained by setting $m = 1/\Delta$ and using $R_{t,i}^2$ to estimate the above conditional variance of $R_{t,i}$. We set $m = 260$ which is appropriate when it can be assumed that there is no volatility during the weekend and a year contains exactly 52 weeks.
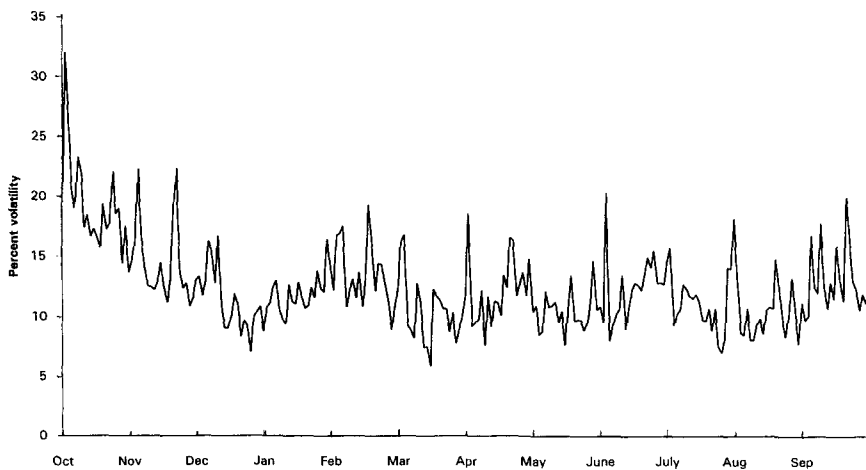
Fig. 1. Volatility estimates from intra-day quotations.

The estimate v_t will not be the optimal estimate of $\sigma(\tau)$ when volatility is constant within cycles. However, the estimate is consistent ($v_t \rightarrow \sigma(\tau)$ as $n \rightarrow \infty$) and it does not require estimation of intra-day seasonal volatility terms.

4.3. The estimates from intra-day quotations

Fig. 1 is a time-series plot of the volatility estimates v_t for the 253 days that the PHLX was open between 1 October 1992 and 30 September 1993 inclusive. The average of these estimates is 12.5% and their standard deviation is 3.6%. Further descriptive statistics are presented in Table 1. The estimates have also been calculated for US holidays and are smaller numbers as should be expected. The two extreme holiday estimates are 2.5% on Christmas Day and 1.9% on New Year's Day; the other six holiday estimates range from 6.7 to 10.5%.

The estimates are higher in October 1992 than in any other month, with the two highest estimates, 32 and 26%, respectively, calculated for Friday 2nd and Monday 5th October. The October average is 19.3% compared with 14.4% for November and 11.7% for the other ten months. The difference may be associated with events that followed the departure of Sterling from the EMS in September 1992.

The estimates display a clear day-of-the-week effect. The average estimate increases monotonically as the week progresses, from 11.4% on Monday to 14.1% on Friday. This pattern reflects the predominance of important scheduled macroeconomic announcements on Fridays and less important announcements on Thursdays. Parametric (ANOVA) and non-parametric (Kruskal–Wallis) tests have p-values below 0.2% for tests of the null hypothesis that the distribution of the

Table 1

Summary statistics for volatility estimates v calculated from intra-day price quotations and implied volatilities i calculated from options prices

	Intra-day estimates v		Implied volatilities i	
	Oct./Sept.	Dec./Sept.	Oct./Sept.	Dec./Sept.
Sample size	253	211	253	211
Mean	12.53	11.66	13.57	12.64
Standard deviation	3.57	2.61	2.68	1.27
Minimum	5.83	5.83	9.71	9.71
Lower quartile	10.30	9.76	11.92	11.81
Median	11.86	11.26	12.82	12.45
Upper quartile	13.90	12.81	13.90	13.36
Maximum	32.05	20.32	24.24	16.69
Monday mean	11.43	10.42	13.74	12.67
Tuesday mean	11.96	11.19	13.68	12.72
Wednesday mean	12.03	11.31	13.53	12.73
Thursday mean	13.14	12.26	13.53	12.66
Friday mean	14.12	13.09	13.36	12.42
p-value, ANOVA	0.001	0.000	0.966	0.816
Autocorrelation				
Lag 1	0.628	0.386	0.914	0.800
Lag 2	0.444	0.042	0.863	0.699
Lag 3	0.392	0.077	0.821	0.632
Lag 4	0.382	0.038	0.777	0.603
Lag 5	0.382	0.120	0.734	0.565
Partial autocorrelation				
Lag 2	0.083	0.184	0.169	0.165
Lag 3	0.140	0.150	0.067	0.094
Lag 4	0.123	0.119	0.001	0.127
Lag 5	0.109	0.172	−0.017	0.037

Summary statistics are calculated for the 12 months from October 1992 to September 1993 and for the 10 months commencing December 1992.

estimates is identical for the five days of the week. Removing the high volatility months of October and November reduces the mean estimate by about 1.0% for each day but the monotonic pattern and the low p-values remain.

The autocorrelations and partial autocorrelations of the volatility estimates are similar to those expected from an AR(1) process. The first-lag autocorrelation is 0.63 for all the estimates but it falls to 0.39 when October and November are excluded.

4.4. Implied volatilities

Fig. 2 is a time-series plot of implied volatility estimates i_t for the same days as are used to produce Fig. 1. Each estimate is the average of two implied
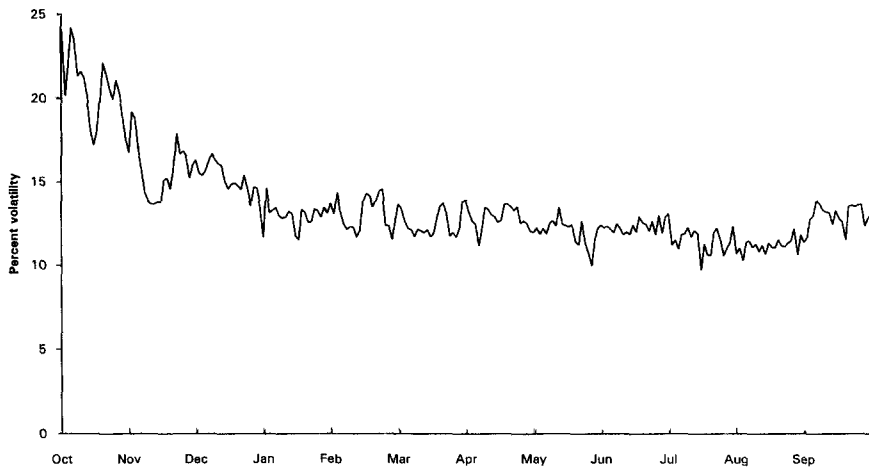
Fig. 2. Implied volatilities.

volatilities, one calculated from a nearest-the-money (NTM) call option price and the other from a NTM put option price. The last options prices before the PHLX close at 14.30 local time are used. These are the only useful options prices supplied to us by the PHLX: high and low options prices are supplied but they do not usually define high and low implied volatilities. The spot prices used for the calculations of the implieds are contemporaneous quotations supplied by the PHLX.

On each day, the shortest maturity options with more than nine calendar days to expiration are selected. The time to maturity of the options is always between 10 and 45 calendar days. We only use the estimates i_t to represent options information about volatility expectations. We do not seek shorter-term expectations from the term structure of implieds because this involves extrapolations that produced no statistical benefits in Xu and Taylor (1995).

The average of the estimates i_t is 13.6%, which is slightly more than the average of the intra-day estimates. Table 1 provides information for comparisons of the distributions of the implied and intra-day estimates. Fig. 2 shows that traders expected a higher level of volatility in October and November and thereafter had expectations that were within an unusually narrow band. There are no day-of-the-week effects because the implieds are expectations for long periods that average 25 calendar days. The implied estimates i_t are markedly less variable than the realised estimates v_t again because the implieds are a medium-term expectations measure. This also explains why the serial correlation in the implied volatilities is substantial: 0.91 at a lag of one-day, using all the data and 0.80 when the first two months are excluded.
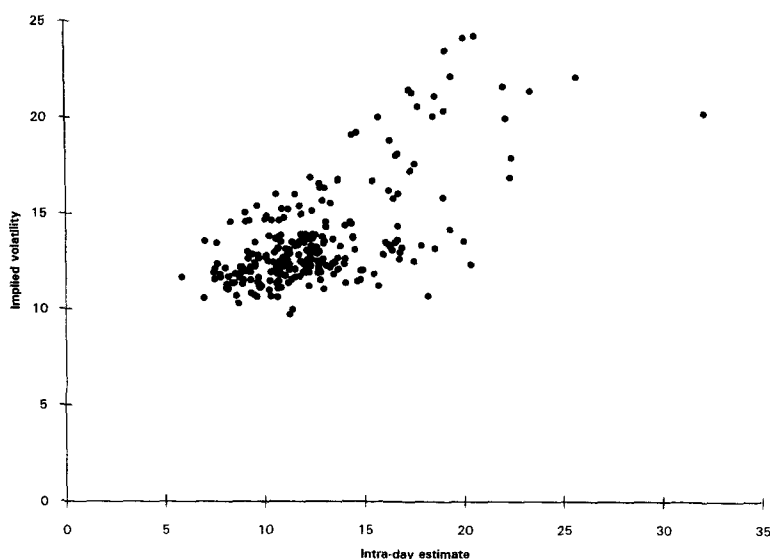
Fig. 3. Comparison of implied volatilities and intra-day estimates.

The correlation between the implied volatilities and the intra-day volatility estimates is 0.66. These two volatility measurements are plotted against each other on Fig. 3.

5. ARCH models with volatility estimated from intra-day quotations

Models and results are first discussed for daily returns and are subsequently discussed for hourly returns. Daily models are straightforward because they avoid estimation of intra-day, seasonal volatility patterns. Hourly models, however, are more incisive because of the much larger number of observed returns.

5.1. A general model for daily returns

ARCH models are estimated for daily spot returns, $R_t = \ln(S_t/S_{t-1})$, obtained from rates when the PHLX closes. All the ARCH models are estimated using data for the set of PHLX trading days. Our set of 253 daily returns is small.

The results are unusual and only need to be discussed when the conditional distribution of returns is normal with mean zero and a conditional variance h_t that depends on the information K_{t-1}, given by combining the information from options trades with the set of five-minute returns up to time $t - 1$. The options information is summarised by the implied volatility term i_{t-1}. The volatility information provided by the five-minute returns is summarised by the estimate v_{t-1}.

Table 2
Parameter estimates for daily ARCH models that include intra-day volatility estimates and short-maturity implied volatilities

$c \times 10^5$	a	b	d	e	max. $\ln(L)$
3.484 (3.55)			0.329 (2.01)		868.66
0.204 (1.95)		0.956 (45.87)	0.000		871.05
0.203 (1.94)	0.000	0.956 (45.78)			871.05
0.203 (1.89)	0.000	0.956 (44.38)	0.000		871.05
				1.000	870.95
0.897 (0.64)				0.683 (2.88)	873.04
0.897 (0.64)	0.000	0.000		0.683 (2.88)	873.04
0.897 (0.64)	0.000	0.000	0.000	0.683 (2.88)	873.04

The numbers in parentheses are t-statistics, estimated using the Hessian matrix and numerical second derivatives. t-statistics are not reported when an estimate is less than 10^{-6}.

The 24-hour conditional variance h_t is the product of the 24-hour deseasonalised conditional variance h_t^* and a multiplier that is either 1, M (for Mondays) or H (for holidays). The deseasonalised conditional variance is defined by $h_t^* = c + aR_{t-1}^2(h_{t-1}^* / h_{t-1}) + bh_{t-1}^* + dv_{t-1}^* + ei_{t-1}^*$. The terms R_{t-1}, v_{t-1}^* and i_{t-1}^* are, respectively, daily returns, intra-day volatility estimates and the squares of scaled implied volatilities. All parameters are constrained to be non-negative. In the fifth row, e is constrained to equal one. The estimates of M and H for the most general model are 1.44 and 1.74, standard errors 0.34 and 0.89, respectively.

The following model makes use of conditional variances h_t^* appropriate for 24-hour periods after removing multiplicative Monday and holiday effects, defined by Eq. (1):

$$R_t | K_{t-1} \sim N(0, h_t),\tag{8}$$

$$h_t = h_t^*, \; Mh_t^* \text{ or } Hh_t^*,\tag{9a}$$

$$h_t^* = c + aR_{t-1}^2(h_{t-1}^*/h_{t-1}) + bh_{t-1}^* + dv_{t-1}^* + ei_{t-1}^*,\tag{9b}$$

$$v_{t-1}^* = v_{t-1}^2/f,\tag{10a}$$

$$i_{t-1}^* = i_{t-1}^2/f,\tag{10b}$$

$$f = 196 + 48M + 8H.\tag{10c}$$

The parameter vector is $\theta = (a, b, c, d, e, M, H)$. The terms v_{t-1}^2 and i_{t-1}^2 are divided by f to convert these annual quantities into quantities appropriate for a 24-hour period.

5.2. Results for daily returns

Table 2 presents results for the general model and seven special cases. When the history of five-minute returns contains all relevant information about future volatility, the options parameter e is zero. An estimation with this constraint produces a surprise, when the initial value h_0^* is an additional parameter. As

$a = d = 0$, the conditional variances are deterministic, hence if the unconditional variance is $\mu_h = c/(1 - b)$ then:

$$h_t^* = \mu_h + b^t(h_0^* - \mu_h).\tag{11}$$

This result is less surprising when we recall the volatility estimates plotted on Fig. 1. The twelve months begin with high volatility followed by a long period during which volatility does not change much. The above edge solution is unlikely to be estimated if the period of exceptionally high variance is anywhere other than at the beginning of the sample. Ex post, the selection of dates for the sample period is rather unfortunate!

The edge solution is a consequence of an unusual volatility pattern found in a small sample. Small samples can give more ordinary results, for example $a = 0.035$ and $b = 0.917$ for GARCH(1, 1) estimated from the daily DM/$ rate from September 1994 to August 1995.

Next, consider models that make use of the information in implied volatilities. The specification

$$h_t^* = c + ei_{t-1}^*\tag{12}$$

has a maximum likelihood that is 1.99 above that of the edge solution. Estimation of the most general model simply produces the linear function of squared implied volatility above; the estimates of a, b and d are all zero.

The results are compatible with the hypothesis that there is *no* incremental volatility information in the dataset of five-minute returns, when calculating daily conditional variances. However, the hypothesis that there is no incremental volatility information in the implied volatilities is dubious.

5.3. Intra-day seasonal multipliers

We now multiply the number of returns used to estimate models by 24. The much larger sample size provides a reasonable prospect of avoiding the unsatisfactory edge solutions found for daily returns. Before estimating conditional variances for hourly returns we must, however, produce estimates of the intra-day seasonal volatility pattern. We present simple estimates here. Our estimates ignore the effects of scheduled macroeconomic news announcements; we discuss the sensitivity of our conclusions to this omission in Section 5.6. Andersen and Bollerslev (1997) provide different estimates based upon smooth harmonic and polynomial functions.

It may be helpful to review some notation before producing the seasonal estimates. The time t is an integer that counts weekdays, n is the number of 5-minute returns in one day ($ = 288$) and $R_{t,i}$ is a 5-minute return; $i = 1$ identifies the return from 14.30 to 14.35 US eastern time (ET) on the previous day (i.e.

$t - 1) \ldots i = 288$ is the return from 14.25 to 14.30 ET on day t. Returns over 24 hours and over 1 hour periods indexed by j are respectively given by

$$R_t = \sum_{i=1}^{n} R_{t,i} \quad \text{and} \quad r_{t,j} = \sum_{i=12(j-1)+1}^{12j} R_{t,i}. \tag{13}$$

Sums of squared returns provide simple estimates of price variability and averages across similar time periods can be used to estimate the seasonal volatility pattern. Let N be the number of days in the sample. It would be convenient if the seasonal pattern could be described by 24 one-hour, multiplicative, seasonal variance factors s_j^2, with $\sum_{j=1}^{24} s_j^2 = 24$. A natural estimate of the variance multiplier for hour j is given by

$$\hat{s}_j^2 = \frac{24 \sum_{t=1}^{N} \sum_{i=12(j-1)+1}^{12j} R_{t,i}^2}{\sum_{t=1}^{N} \sum_{i=1}^{n} R_{t,i}^2}. \tag{14}$$

However, the seasonal pattern varies by day of the week, as might be expected from Table 1 and thus it appears preferable to estimate 120 multiplicative factors that average one over a complete week.

A second way to estimate variance multipliers takes account of the day of the week. Let S_t be the set of all daily time indices that share the same day-of-the-week
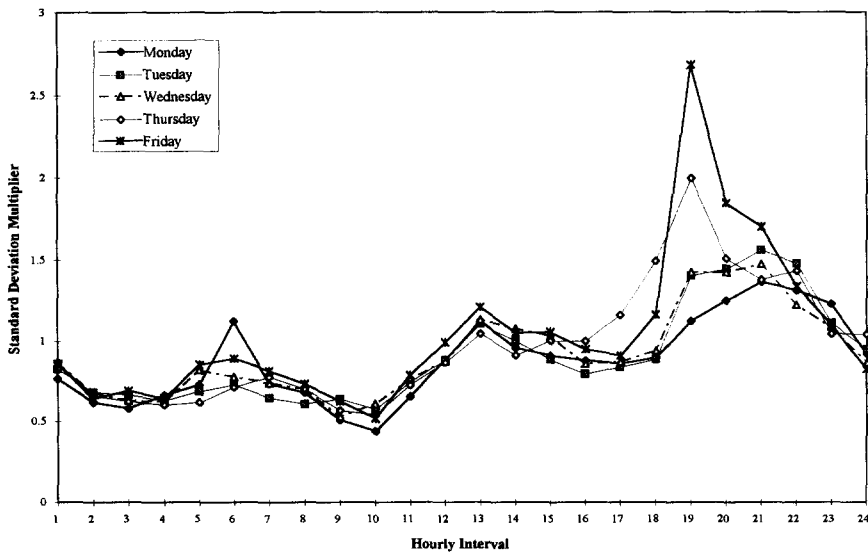

Fig. 4. DM/$ intra-day standard deviation multipliers.

as time index t. Let N_t be the number of time indices to be found in S_t. Then a set of 120 factors are given by:

$$\hat{s}_{t,j}^2 = \frac{24N}{N_t} \frac{\sum_{s \in S_t} \sum_{i=12(j-1)+1}^{12j} R_{s,i}^2}{\sum_{s=1}^{N} \sum_{i=1}^{n} R_{s,i}^2}. \tag{15}$$

Fig. 4 is a plot of standard deviation multipliers, $\hat{s}_{t,j}$. The final hourly interval, $j = 24$, is the hour ending at 14.30 ET (19.30 GMT, winter) when the options market closes. The first interval, $j = 1$, is the hour beginning after the previous day's options close.

The multipliers are generally higher for intervals 13 to 24, corresponding to 07.30 until 19.30 local time in London, with the highest levels in intervals 18 to 23 when both US and European dealers are active. The Thursday and Friday spikes, at interval $j = 19$, reflect the additional volatility when many US macroeconomic news reports are released in the hour commencing at 08.30 ET. Ederington and Lee (1993, 1995) provide detailed documentation of this link with macroeconomic news. The lower local maximum, at $j = 13$, occurs when trade accelerates in Europe in the hour commencing at 07.30 local time in London. The Monday spike earlier in the day, at $j = 6$, is the start of a new week in the Far East markets.

5.4. A model for hourly returns

An ARCH specification for hourly returns that is similar to that considered for the daily returns involves hourly returns $r_{t,j}$, information sets $K_{t,j-1}$, recent five-minute returns $R_{t,i}$, one-hour realised variances $V_{t,j}$, one-hour conditional variances $h_{t,j}$, one-hour deseasonalised conditional variances $h_{t,j}^*$ and the multipliers $\hat{s}_{t,j}$. The specification also incorporates the annualised implied volatility i_{t-1} calculated at the previous close; hourly implieds are not available to us, although we would not expect them to contribute much because the implieds change slowly. The information set $K_{t,j-1}$ is defined to be all relevant variables known at the end of hour j-1 on day t, namely the implieds i_{t-1}, i_{t-2},..., the latest five-minute return $R_{t,12(j-1)}$ and all previous five-minute returns.

The most general ARCH model that has been estimated for hourly returns is:

$$r_{t,j} | K_{t,j-1} \sim D_\nu \left(m_{t,j}, h_{t,j} \right), \tag{16a}$$

$$m_{t,j} = \phi r_{t,j-1}, \tag{16b}$$

$$h_{t,j} = \hat{s}_{t,j}^2 h_{t,j}^*, \tag{16c}$$

$$h_{t,j}^* = c + a r_{t,j-1}^2 / \hat{s}_{t,j-1}^2 + b h_{t,j-1}^* + d V_{t,j-1} / \hat{s}_{t,j-1}^2 + e i_{t-1}^*, \tag{16d}$$

$$V_{t,j-1} = \sum_{i=12(j-2)+1}^{12(j-1)} R_{t,i}^2, \tag{16e}$$

$$i_{t-1}^* = i_{t-1}^2 / f, \tag{16f}$$

for some number f that does not need to be estimated; we set f equal to the number of annual hourly returns (24×252). The subscript pair t,i refers to the time interval $t - 1$, $n - i$ whenever i is not positive. The distribution $D_\nu(m_{t,j}, h_{t,j})$ is GED with thickness parameter ν, mean $m_{t,j}$ and variance $h_{t,j}$. The parameter vector is $\theta = (a, b, c, d, e, \phi, \nu)$. As the autoregressive, mean parameter ϕ is always insignificant, we only discuss results when ϕ is constrained to be zero.

Eq. (16d) contains terms, with coefficients a and d, that are both measures of hourly return variability. Both measures are included to permit comparisons of the information content of five-minute and hourly returns.

5.5. Results for hourly returns

Table 3 presents results for 6049 hourly returns when 120 seasonal volatility multipliers are included in the models. The maximum log-likelihood increases substantially when 120 day-of-the-week multipliers replace 24 hourly multipliers, typically by about 65 for conditional normal distributions and by about 22 for conditional GED distributions. Consequently, our discussion of the results is based upon models with 120 intra-day seasonal multipliers. All our observations and conclusions are also supported by numbers in a further table, available upon request, for models that have only 24 seasonal multipliers.

The lower panel of Table 3 shows that the conditional distribution of the hourly returns is certainly fat-tailed. The GED thickness parameter is estimated to be near 1.15 with a standard error less than 0.03. Conditional normal distributions are rejected for the most general specification and all the special cases. The log-likelihood ratio test statistic is 130.26 for the general specification with the null distribution being χ_1^2. A thickness parameter of 1 defines double negative-exponential distributions so the hourly returns have conditional distributions that are far more peaked and fat-tailed than the normal.

Our assumption of the GED for the conditional distributions does not ensure consistent parameter estimates and standard errors if the assumption is false. The quasi-ML estimates in the upper panel of Table 3 are consistent although they are not efficient.

The results in the lower panel of Table 3 fall into three major categories, and are discussed separately. The conclusions are the same if we focus on the upper panel for normal distributions.

First, consider models that only make use of returns information. The models that incorporate information more than one-hour old, through parameter b, have significant parameters for both recent information (the last hour, through a and d) and old information. This is the usual situation when ARCH models are estimated and so we no longer have the curious edge solutions discussed for the daily returns in Section 5.2. When five-minute returns are used, but hourly returns are not ($a = 0$; $b, d > 0$), the maximum of $\ln(L)$ is 23 more than the maximum when only

Table 3
Parameter estimates for ARCH models of hourly returns with 120 seasonal terms

$c \times 10^5$	a	b	d	e	v	max. $\ln(L)$
Panel A: normal distribution						
0.1321			0.2875		2	31848.67
(27.29)			(13.38)			
0.0012		0.9490	0.0352		2	31963.41
(2.35)		(115.84)	(6.49)			
0.0012	0.0196	0.9743			2	31933.21
(3.24)	(5.58)	(197.49)				
0.0012	0.0045	0.9480	0.0319		2	31964.09
(2.44)	(1.14)	(110.71)	(5.38)			
				0.6528	2	31945.15
				(54.99)		
0.0000				0.6528	2	31945.15
				(54.88)		
0.0000	0.1046	0.2926		0.3935	2	31999.55
	(6.98)	(3.78)		(8.20)		
0.0000		0.1766	0.1437	0.4141	2	31997.85
		(1.70)	(7.01)	(6.76)		
0.0000	0.0713	0.1812	0.0876	0.4127	2	32011.87
	(4.51)	(2.14)	(4.22)	(8.30)		
Panel B: GED distribution						
0.1232			0.3197		1.1025	32193.31
(19.12)			(10.68)		(42.42)	
0.0008		0.9434	0.0408		1.1460	32253.21
(1.25)		(89.55)	(5.53)		(41.90)	
0.0013	0.0227	0.9707			1.1376	32230.24
(2.38)	(4.45)	(136.35)			(41.88)	
0.0009	0.0028	0.9431	0.0387		1.1463	32253.36
(1.31)	(0.52)	(87.44)	(4.65)		(41.89)	
				0.6482	1.1418	32231.90
				(39.05)	(41.58)	
0.0000				0.6482	1.1458	32231.90
				(38.98)	(41.58)	
0.0000	0.1221	0.2590		0.4034	1.1598	32266.90
	(5.74)	(2.66)		(6.71)	(41.40)	
0.0000		0.2212	0.1678	0.3641	1.1594	32268.01
		(1.88)	(5.96)	(5.38)	(41.51)	
0.0000	0.0808	0.1801	0.1099	0.3887	1.1638	32277.00
	(3.64)	(1.83)	(3.87)	(6.71)	(41.45)	

The numbers in parentheses are t-statistics, estimated using the hessian matrix and numerical second derivatives. All parameters are constrained to be non-negative. t-statistics are not reported when an estimate is less than 10^{-6}.

The one-hour conditional variance is defined by $h_{t,j} = \hat{s}_{t,j}^2 h_{t,j}^*$, $h_{t,j}^* = c + a r_{t,j-1}^2 / \hat{s}_{t,j-1}^2 + b h_{t,j-1}^* + d(\sum_{i=12(j-2)+1}^{12(j-1)} R_{t,i}^2)/\hat{s}_{t,j-1}^2 + e i_{t-1}^*$.

The terms $r_{t,j-1}$, $R_{t,i}$ and i_{t-1}^* are, respectively, one-hour returns, five-minute returns and the squares of scaled implied volatilities. The conditional distributions are normal distributions in panel A and are generalised error distributions, with thickness parameter v, in panel B.

hourly returns are used ($d = 0$; $a, b > 0$). There is thus more relevant volatility information in five-minute returns than in hourly returns. This information comes from more than twelve five-minute returns, as expected, because the maximum of $\ln(L)$ decreases by 60 when older information is excluded ($a = b = 0$; $d > 0$). When all the returns variables are included in the model, a is insignificant and much smaller than d. The persistence estimates, given by the sum $a + b + d$, are between 0.984 and 0.993 when old information is included ($b > 0$).

Second, consider models that only make use of daily implied volatilities. The variable i_{t-1}^* is biased because estimates of the multiplier e are significantly smaller than 1. Some of this bias is presumably due to an unsuitable choice for the constant f that converts annual variances into hourly variances. When e and f are unconstrained the maximum of $\ln(L)$ is 21 less than the maximum when spot price quotations alone are used. This shows that five-minute returns are more informative than implied volatilities, at least when estimating hourly conditional variances.

Third, consider models that make use of five-minute returns, hourly returns and daily implied volatilities. The most general model in the final row of Table 3 is estimated to have a zero intercept c and the parameters a, d and e have t-ratios above 3.5 and thus are significant at very low levels. Deleting the implied volatility contribution from the most general model would reduce the maximum of $\ln(L)$ by 24. Alternatively, deleting the quotations terms would give a reduction of 45. It is concluded that both the quotations and the implied volatilities contain a significant amount of incremental information.

5.6. Results when scheduled news is incorporated

The hourly seasonal volatility multipliers are particularly high in the hour commencing at 08.30 ET when many US macroeconomic news reports are released. This effect is most prominent on Fridays. The volatility multipliers used in the preceding analyses are, for example, the same for all Friday hours commencing at 8:30 regardless of any news releases. This methodology might induce systematic mis-measurements of the volatility process. We have assessed the importance of this issue by comparing the results when there are 120 volatility multipliers with further results when either 121 or 144 multipliers are used.

Our first set of 121 multipliers contains two numbers for Friday 08.30 to 09.30 ET: one multiplier for those Fridays that have a relevant report and another multiplier for the remaining Fridays. Our first set of 144 multipliers contains two numbers for each of the 24 hours from Thursday 14.30 to Friday 14.30 ET, one used when there is a relevant report and the other when there is not. We have defined a relevant report as a news announcement about one or more of the six significant macroeconomic variables listed by Ederington and Lee (1993, p. 1189): employment, merchandise trade, PPI, durable goods orders, GNP and retail sales. These reports were issued on 25 of the Fridays in our sample.

We find that the maximum of the log-likelihood function increases by similar amounts when there are more multipliers whichever model is estimated. Consider

the nine log-likelihood values reported in Table 3, panel B, for nine specifications with 120 multipliers. These values increase by between 6.6 and 8.6 when 121 multipliers are used and by between 16.7 and 19.6 when 144 are used. Consequently, as our conclusions depend on substantial log-likelihood differences across specifications these conclusions do not change when the additional multipliers are used.

It may be objected that the Fridays have been partitioned by announcements rather than by the impact of unexpected news. Second sets of 121 and 144 multipliers have been calculated by separating the 25 Fridays having the highest realised volatility from 08.30 to 09.30 from the remaining Fridays. The results are then similar, as 19 of the 25 high-volatility hours include a relevant announcement. The increases in the log-likelihoods from the values reported in Table 3, Panel B are now in the ranges 8.4 to 10.7 and 9.3 to 11.7, respectively, for 121 and 144 multipliers. We note that the first, second, fourth and fifth Fridays in the ranked list coincide with employment reports but the third ranked Friday has no relevant announcements.

The estimates of the parameters a, b, d and e change very little when the number of multipliers is increased above 120. The magnitudes of the changes are all less than 0.03 when the most general model is estimated. When the general model is constrained by ignoring the options information ($e = 0$), the persistence measure $a + b + d$ is always between 0.984 and 0.985.

5.7. Results for quarterly subperiods

It could be possible that some of the conclusions are only supported by the data during part of the year studied. The higher than average realised volatility during the first quarter, from October to December 1992, might be an unusual period whose exclusion would reverse some of the conclusions.

The models whose parameter estimates have been given in the lower panel of Table 3 for the whole year in the datasets have been re-estimated for the four quarters of the year commencing in October 1992 and in January, April and July 1993. The same 120 seasonal multipliers are used for the whole year and for each of the four quarters. All the conclusions for the whole year are supported by each quarter of the data: (1) when quotations alone are considered, five-minute returns have more volatility information than hourly returns and the relevant information is not all in the most recent hour (likelihood-ratio tests, 5% significance level), (2) five-minute returns are more informative than implied volatilities when estimating hourly conditional variances and (3) there is significant incremental information in both the quotations and the implied volatilities (likelihood-ratio tests, 5% significance level).

The reductions in the maximum of the log-likelihood when the quotations information is removed from the most general model are 14.7 for the first quarter, 8.0 for the second quarter, 8.4 for the third quarter and 17.3 for the fourth quarter.

Table 4
Parameter estimates for the most general ARCH model of hourly returns

Sample	$c \times 10^5$	a	b	d	e	v	max. $\ln(L)$
Full	0.0000	0.0808	0.1801	0.1099	0.3887	1.1638	32277.00
		(3.64)	(1.83)	(3.87)	(6.71)	(41.45)	
Q1	0.0000	0.0775	0.0000	0.1591	0.4841	1.1621	7726.38
		(1.78)		(2.54)	(11.01)	(20.94)	
Q2	0.0000	0.0435	0.2391	0.0995	0.3797	1.2155	8622.58
		(1.19)	(1.05)	(1.84)	(2.82)	(21.04)	
Q3	0.0164	0.1071	0.0000	0.0877	0.4334	1.1480	7676.79
	(0.32)	(2.06)		(1.53)	(2.05)	(20.05)	
Q4	0.0000	0.0750	0.4420	0.1185	0.2120	1.1421	8255.79
		(1.47)	(2.33)	(2.50)	(2.18)	(20.65)	

The general model has conditional variances defined in Table 3. There are 120 seasonal multipliers and the conditional distributions are generalised error distributions. The estimates are for the full year (October 1992 to September 1993) and for the four quarters that commence in October 1992, January 1993, April 1993 and July 1993. The numbers in parentheses are t-statistics.

The incremental information in the quotations information is thus of a similar order of magnitude in all the quarters and the first quarter is not clearly different to the other three quarters. The reductions in the maximum of the log-likelihood when the information in implied volatilities is removed from the most general model are 4.2 for the first quarter, 4.4 for the second quarter, 2.4 for the third quarter and 2.1 for the fourth quarter. These reductions are much smaller and are similar across quarters.

Table 4 presents the quarterly parameter estimates for the most general model. The estimates change little from quarter to quarter. The sum of the maximum log-likelihoods for the four quarters is only 4.54 more than the maximum when the same parameters are used for the whole year. Twice this increase in the log-likelihood is less than the number of extra parameters when four quarterly models are estimated compared with one annual model. There is no statistical evidence, therefore, that the parameters of the general model changed during the year. The variations in estimated parameters, by quarter, are minor relative to their estimated standard errors.

5.8. Residual diagnostic statistics and tests

A time series of standardised residuals from our most general model for hourly returns is defined by:

$$z_T = r_{t,j} / \sqrt{h_{t,j}}, \quad T = 24t + j. \tag{17}$$

The conditional variances are calculated using the maximum likelihood estimates of the model parameters for the whole year. In the unlikely event of our model

being perfect we would expect the standardised residuals to be approximately independent and identically distributed observations from a zero-mean and unit-variance distribution.

There are 6049 numbers in the time series $\{z_T\}$. Their mean is 0.005 and their standard deviation is 1.004. Their skewness is -0.01 and their kurtosis is 5.30, both of which are close to the values expected from a generalised error distribution with thickness parameter ν near one (skewness $= 0$ for all ν, and kurtosis $= 6$ when $\nu = 1$). A histogram of the z_T shows fat tails and a substantial peak around zero, which is a feature of the GED when ν is near one. Twenty of the standardised residuals are outside ± 4 although all of them are inside ± 5.5.

The autocorrelations of z_T, $|z_T|$ and z_T^2, from lags 1 to 10, are all within $\pm 0.025 = \pm 1.96 / \sqrt{6049}$ and therefore provide no evidence against the i.i.d. hypothesis, since all 30 tests accept this null hypothesis at the 5% level. The first-lag autocorrelations of the three series are 0.003, 0.007 and -0.007. Statistically significant dependence is found at lags that are multiples of 24: for z_T at lag 96 (correlation $= 0.056$), for $|z_T|$ at lags 24, 48, 72, 96 and 120 (the correlations are 0.045, 0.042, 0.049, 0.056 and 0.026) and for z_T^2 at lags 24, 72 and 96 (correlations 0.032, 0.022 and 0.062). These correlations show that the model is not perfect, presumably because of estimation errors in the hourly seasonal multipliers. Nevertheless, with all autocorrelation estimates within ± 0.07 the model is considered a satisfactory approximation to the process that generates hourly returns.

Estimates of spectral density functions, calculated from the autocorrelations at lags 1 to 240 of z_T, $|z_T|$ and z_T^2, confirm this conclusion. No statistical evidence against the i.i.d. hypothesis can be found in the estimates at frequencies corresponding to either 24-hour or 120-hour cycles. There is a significant spectral peak at zero frequency for the series $|z_T|$ (t-statistic $= 3.49$) that may simply reflect very small positive dependence at several lags.

6. Forecasts of realised volatility

A comparison of volatility forecasts can provide further evidence about incremental information. We divide the whole year of data into an in-sample period from which ARCH parameters and intra-day seasonal volatility multipliers are estimated and an out-of-sample period for which the accuracy of forecasts of hourly realised volatility is evaluated. We split the year into a nine-month in-sample period followed by a three-month out-of-sample period. Relative accuracy measures for five forecasting methods are calculated using 120 seasonal volatility multipliers. The relative measures are not sensitive to the treatment of Friday macroeconomic announcements. Using our first set of 121 or 144 multipliers, defined in Section 5.6, has no effect on the rankings of the forecasts.

Two measures of hourly realised volatility are forecast, defined first by

$$a_{t,j,1} = \sum_{i=12(j-1)+1}^{12j} R_{t,i}^2 \qquad (18)$$

and second by the same quantity adjusted for intra-day seasonality using 120 volatility multipliers:

$$a_{t,j,2} = a_{t,j,1}/\hat{s}_{t,j}^2. \qquad (19)$$

Three forecasts of $a_{t,j,1}$ are defined by conditional variances $h_{t,j}$ obtained from variations on the most general ARCH model for hourly returns defined in Section 5.4. The first forecast excludes all quotations information by imposing the restriction $a = b = d = 0$ on the ARCH model. The second forecast excludes the options information by requiring $e = 0$ in the ARCH model. The third forecast is calculated from the general model without any parameter restrictions. These three forecasts are denoted $f_{t,j,1,l}$, $l = 1, 2, 3$. A fourth forecast, $f_{t,j,1,4}$, is defined by $\hat{s}_{t,j}^2$ multiplied by the in-sample average of the quantities $a_{t,j,2}$. Four forecasts $f_{t,j,2,l}$ of $a_{t,j,2}$ are defined in a similar way. The first three of these forecasts are now defined by deseasonalised conditional variances $h_{t,j}^*$ for the three ARCH specifications and the fourth forecast is the in-sample average of $a_{t,j,2}$.

The accuracy of a set of forecasts $f_{t,j,k,l}$ of the outcomes $a_{t,j,k}$ is reported here relative to the accuracy of a reference forecast given by the previous realised volatility

$$f_{t,j,1,5} = a_{t,j-1,2}\,\hat{s}_{t,j}^2, \qquad f_{t,j,2,5} = a_{t,j-1,2}. \qquad (20)$$

Table 5 presents values of the relative accuracy measures

$$F_{k,l,p} = \frac{\sum|a_{t,j,k} - f_{t,j,k,l}|^p}{\sum|a_{t,j,k} - f_{t,j,k,5}|^p} \qquad (21)$$

for powers $p = 1, 2$. The summations are over all hours in the out-of-sample period. The best of a set of five forecasts $f_{t,j,k,l}$, $l = 1, \ldots, 5$, has the least value

Table 5
Measures of relative forecast errors when forecasting hourly realised volatility out-of-sample

| Forecast | Error metric | Absolute | Absolute | Square | Square |
	Seasonal adjustment	No	Yes	No	Yes
1	options only	0.767	0.772	0.807	0.795
2	quotations only	0.781	0.769	0.812	0.783
3	options and quotations	0.731	0.731	0.797	0.776
4	in-sample average	1.055	1.080	0.821	0.803
5	lagged realised volatility	1	1	1	1

The accuracy of forecasts is measured by either the absolute forecast error or the squared forecast error. Hourly realised volatility is forecast, either without or with a seasonal adjustment.
Nine months are used for in-sample calculations and then three months for out-of-sample evaluations. The numbers tabulated are $\sum|a - f_l|^p/\sum|a - f_5|^p$ with a the realised volatility number, f_l forecast l and p either 1 or 2.

of F. The least value of F is considered for each of the four columns in Table 5. The columns are defined by all combinations of p (1 or 2) and k (1 or 2).

When accuracy is measured by absolute forecast errors, so $p = 1$, the best forecasts come from the general ARCH specification for both realised measures ($k = 1, 2$). This is further evidence that there is incremental volatility information in both the spot quotations and the options prices. The average absolute forecast error from the general specification is five per cent less than that from the next best specification. The second best set of forecasts are from quotations alone when the quantity forecast is adjusted for seasonality, but are from options prices alone when the quantity forecast is not adjusted, although the differences between the accuracies of the second and third best forecasts are small.

The results are similar but less decisive when accuracy is measured by squared forecast errors ($p = 2$). The most general ARCH specification again gives the best out-of-sample forecasts. However, the average of the squared forecast errors for the best forecasts are only slightly less than for the next best forecasts. This may be attributed to the marked skewness to the right of the distribution of the quantities to be forecast: this inevitably produces some outliers in the forecast errors whose impact is magnified when they are squared.

7. Concluding remarks

The evidence from estimating ARCH models using one year of exchange rate quotations for one exchange rate supports two conclusions. First, five-minute returns cannot be shown to contain any incremental volatility information when estimating daily conditional variances. This negative result may simply be a consequence of the small number of daily returns available for this study. Second, when estimating hourly conditional variances there is a significant amount of information in five-minute returns that is incremental to the options information. Furthermore, the quotations information then appears to be more informative than the options information. Thus there is significant incremental volatility information in one million foreign exchange quotations. This conclusion is confirmed by out-of-sample comparisons of volatility forecasts. Forecasts of hourly realised volatility are more accurate when the quotations information is used in addition to options information.

References

Andersen, T.G., Bollerslev, T., 1997. Intraday periodicity and volatility persistence in financial markets. Journal of Empirical Finance.

Barone-Adesi, G., Whaley, R.E., 1987. Efficient analytic approximation of American option values. Journal of Finance 42, 301–320.

Bollerslev, T., Domowitz, I., 1993. Trading patterns and prices in the interbank foreign exchange market. Journal of Finance 48, 1421–1443.

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

Dacorogna, M.M., Muller, U.A., Nagler, R.J., Olsen, R.B., Pictet, O.V., 1993. A geographical model for the daily and weekly seasonal volatility in the FX market. Journal of International Money and Finance 12, 413–438.

Day, T.E., Lewis, C.M., 1992. Stock market volatility and the information content of stock index options. Journal of Econometrics 52, 289–311.

Ederington, L.H., Lee, J.H., 1993. How markets process information: news releases and volatility. Journal of Finance 49, 1161–1191.

Ederington, L.H., Lee, J.H., 1995. The short-run dynamics of the price-adjustment to new information. Journal of Financial and Quantitative Analysis 30, 117–134.

Goodhart, C.A.E., Figliuoli, L., 1991. Every minute counts in financial markets. Journal of International Money and Finance 10, 23–52.

Hull, J., 1995. Introduction to Futures and Options Markets, 2nd ed. Prentice-Hall, Englewood Cliffs.

Hull, J., White, A., 1987. The pricing of options on assets with stochastic volatilities. Journal of Finance 42, 281–300.

Jorion, P., 1995. Predicting volatility in the foreign exchange market. Journal of Finance 50, 507–528.

Lamoureux, C.B., Lastrapes, W.D., 1993. Forecasting stock return variance: Toward an understanding of stochastic implied volatilities. Review of Financial Studies 6, 293–326.

Nelson, D.B., 1992. Filtering and forecasting with misspecified ARCH models I: getting the right variance with the wrong model. Journal of Econometrics 52, 61–90.

Taylor, S.J., Xu, X., 1994. The magnitude of implied volatility smiles: Theory and empirical evidence for exchange rates. Review of Futures Markets 13, 355–380.

Xu, X., Taylor, S.J., 1994. The term structure of volatility implied by foreign exchange options. Journal of Financial and Quantitative Analysis 29, 57–74.

Xu, X., Taylor, S.J., 1995. Conditional volatility and the informational efficiency of the PHLX currency options market. Journal of Banking and Finance 19, 803–821.