ELSEVIER

Journal of Banking & Finance 19 (1995) 803–821

Journal of BANKING & FINANCE

Conditional volatility and the informational efficiency of the PHLX currency options market

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

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

Received January 1993; final version received November 1993

Abstract

The relative performance of implied and historical volatility predictors is compared for four exchange rates from 1985 to 1991. For three currencies, ARCH models estimated up to 1989 show that PHLX implied volatilities provide specifications for daily conditional variances which can not be significantly improved by using past returns. This result is consistent with the informational efficiency of the Philadelphia currency options market. Out-of-sample forecasts of the average volatility over four-week periods are evaluated for 1990 and 1991. Once more the implied predictors are superior to historical predictors.

Keywords: ARCH models; Forecasting; Informational efficiency; Options; Volatility

JEL classification: G13; G14

1. Introduction

Options markets are often viewed as markets for volatility trading. Options prices provide forecasts of the future average variance of returns from the underlying asset over the life of the option. The ability of the volatility forecast implied by options prices to predict future volatility is considered a measure of the

* Corresponding author. Tel. 44 161 275 4030, Fax 44 161 275 4023.

SSDI 0378-4266(95)00086-7

information content of option prices (Day and Lewis, 1992a). Until recently most research into information content has focused on relationships between implied volatility and measures of the subsequent realised volatility (Latane and Rendleman, 1976; Chiras and Manaster, 1978; Gemmill 1986; Shastri and Tandon, 1986; Scott and Tucker, 1989). This prior research concludes that volatility predictors calculated from option prices are better predictors of future volatility than standard deviations calculated from historical asset price data. However, Canina and Figlewski (1993) claim that neither implied volatilities nor historical volatilities have much predictive power for the S&P 100 index.

Day and Lewis (1992a) develop a new methodology based upon ARCH models. They examine the information content of implied volatilities, obtained from call options on the S&P 100 index, relative to ARCH estimates of conditional volatility by adding the implied volatility to ARCH models as an exogenous variable. They find that both implied volatility and historical series of asset returns contain incremental information. The null hypothesis that returns contain no volatility information additional to that found in options prices is rejected. Likewise, the hypothesis that options prices have no additional information is rejected. These results imply that the US stock options market either does not use all publicly available information in setting options prices or that the dividend-adjusted version of the Black–Scholes formula used for computing implied volatility is misspecified. Similar results have been found by Day and Lewis (1992b) for the conditional volatility of oil futures and by Lamoureux and Lastrapes (1993) for the conditional volatility of individual stocks. However, it should be noted that these ARCH tests might be unreliable because the implied volatility variable used in the tests is not necessarily appropriate. The ARCH conditional volatility is always for the next period, typically the next day or week. As an implied volatility for the next period rarely exists, Day and Lewis (1992a) use implied volatilities obtained from options having short times to maturity, which range from 7 to 36 days. The implied volatilities used by Lamoureux and Lastrapes (1993) correspond to maturities of between 64 and 129 trading days. It is known that implied volatilities at any moment in time vary for different times to option expiry. This term structure of volatility expectations has been ignored in previous applications of ARCH methodology. This paper includes an empirical evaluation of whether or not this methodological issue is important.

The relative importance of implied and historical volatility predictors is examined once more in this paper using ARCH models, but we go one step further. The first implied predictor used here is an estimate of the volatility expectation for the next period calculated from the volatility term structure model developed in Xu and Taylor (1994). Results for this predictor are compared with results for a second predictor defined by short-maturity implied volatilities. We study four exchange rate series and find that for three series the implied volatility information alone gives optimal predictions of one-period-ahead conditional volatility. Consequently the volatility information provided by currency returns has no incremental predictive power. This is consistent with the informational efficiency of the Philadelphia currency options market. Out-of-sample forecast comparisons confirm this conclusion. The choice of implied volatility predictor (term structure or short maturity) does not affect the conclusions for our data.

The rest of the paper is organised as follows. Section 2 describes the data and a method for estimating volatility expectations for any set of future periods. Exchange rates and options prices from January 1985 to February 1992 are analyzed. Within-sample calculations are confined to the period until November 1989. The subsequent data are reserved for out-of-sample forecasting evaluations. Section 3 discusses briefly the definitions of ARCH models and the specifications of empirical tests of the informational efficiency of the options market. The empirical within-sample results are presented in Section 4 followed by out-of-sample forecasting results in Section 5. The paper concludes with a summary in Section 6.

2. Data and implied volatility methodology

2.1. Datasets

The primary source database for the options prices is the transaction report compiled daily by the Philadelphia Stock Exchange (PHLX). Daily closing option prices and the simultaneous spot exchange rate quotes have been used for the (British) Pound, (Deutsche) Mark, (Japanese) Yen and (Swiss) Franc quoted against the US Dollar from January 2, 1985 to January 8, 1992. However, the transaction report is not available for some trading days. For some other days the report is not complete or in a few cases it is clearly erroneous. Prices have been collected manually from the Wall Street Journal (WSJ) whenever necessary. Approximately 10% of our implied volatilities are calculated from WSJ prices [1].

Uninformative options records are removed from the database. Options which violate boundary conditions, or are either deep in- or out-of-the-money, or have time to expiry less than ten calendar days are not considered. In addition, all implied volatilities more than five standard deviations distant from their sample mean are excluded.

The interest rates used are London euro-currency rates, collected from Datas-

[1] For the period from December 18, 1990 to March 15, 1991 there is no data for the Swiss franc in the PHLX database and the WSJ did not list the closing SF/$ spot exchange rate, so spot rates were collected from Datastream while option prices were collected from the WSJ. This affects the accuracy of implied volatilities calculated from option prices and in turn affects the performance of implied volatility forecasts, and thus the results for the Swiss franc in Section 5 should be interpreted with some caution.

tream. Daily closing prices for futures [2] contracts traded at the International Monetary Market (IMM) in Chicago are also collected from Datastream. Each futures contract is used for the three months prior to its expiration month. At the rollover date, the closing price of the new contract on the previous day before the rollover is used in the calculation of the relative price change. Results from estimating ARCH models show that the choice of method and timing for rolling over futures contracts has insignificant effects.

The IMM closes at 1.20 p.m. while the PHLX closes at 2.00 p.m., and this could lead to bias in favour of the informational efficiency of the options market. However, it has been checked that the conclusions presented in Section 4 do not change even if we use the options estimate of volatility expectations from the previous day.

2.2. *Computation of implied volatility*

Implied volatilities have been calculated from American model prices, approximated by the very accurate functions derived in Barone-Adesi and Whaley (1987). The calculations use an interval subdivision method, which always converges to an unique solution.

All the implied volatilities are calculated for nearest-the-money options; the selected exercise price on a specific day for a specific maturity minimises $|S-X|$. [3] Nearest-the-money options are chosen for two reasons. First, given the widely reported 'strike bias' or 'smile effect' (Shastri and Wethyavivorn, 1987; Sheikh, 1991; Taylor and Xu, 1994), including out-of-the-money and in-the-money options would introduce further noise into volatility expectation estimates. Second, the approximation that the implied volatility of a rationally priced option will equal the mean expected volatility over the time to expiry is generally considered more satisfactory for an at-the-money option than for all other options (Stein, 1989; Day and Lewis, 1992a; Heynen et al., 1994).

2.3. *Estimating the term structure of volatility expectations*

Implied volatilities have a term structure. One set of volatility expectations is obtained by estimating a time-varying term structure model for volatility expectations. Complete details are given in Xu and Taylor (1994).

---

[2] The volatility of currency futures prices is identical to the volatility of spot exchange rates if domestic and foreign interest rates are non-stochastic. Differences between spot and futures volatilities will be minimal as the futures maturities are always less than four months. Futures prices, rather than spot prices, were available to us from Datastream for the whole period under study.

[3] Equal forward and exercise prices define the at-the-money option in theoretical arguments but these arguments are usually developed for European options.

Market agents will have expectations at time t about price volatility during future time periods. Suppose they form expectations of the quantities

$$\mathrm{var}(R_{t+\tau}),\ \text{with}\ R_{t+\tau} = \ln P_{t+\tau} - \ln P_{t+\tau-1},\ \tau = 1, 2, 3, \ldots \tag{1}$$

where P refers to the price of the asset upon which options are traded.

The volatility term structure model involves two factors representing short-term ($\tau = 1$) and long-term ($\tau \to \infty$) annualised volatility expectations, denoted α_t and μ_t, respectively. As the horizon τ increases, the volatility expectations are assumed to revert towards the long-term expectation and the rate of reversion, ϕ, is assumed to be the same for all t. Then the expected volatility at time t for an interval of general length T, from time t to time $t+T$, is the quantity ν_T given by

$$\nu_T^2 = \mu_t^2 + \frac{1-\phi^T}{T(1-\phi)}\left(\alpha_t^2 - \mu_t^2\right) \tag{2}$$

providing it is assumed that subsequent asset prices, $\{P_{t+\tau},\ \tau > 0\}$, follow a random walk.

Kalman filtering methodology applied to implied volatilities provides estimates for the term structure parameters, particularly ϕ, and also time series of short-term and long-term volatility expectation estimates, $\{\hat{\alpha}_t\}$ and $\{\hat{\mu}_t\}$.

3. ARCH methodology

3.1. Specifications of ARCH models using returns information

The expected variance for the next time interval, $t+1$, can be obtained from returns up to time t by using the conditional variance h_{t+1} from an ARCH model. In general:

$$h_{t+1} = \mathrm{var}(R_{t+1} \mid I_t) \tag{3}$$

where I_t denotes the information set of all observed returns up to time t. The most successful and parsimonious models are the GARCH(1,1) model of Bollerslev (1986) and the Exponential ARCH(1,0) model of Nelson (1991). These models have provided satisfactory descriptions of numerous financial time series (Bollerslev et al., 1992).

The GARCH(1,1) model defines the conditional variance recursively using residual terms, which are returns minus their conditional means, thus:

$$\varepsilon_t = R_t - E[R_t \mid I_{t-1}] \tag{4}$$

and

$$h_{t+1} = c + a\varepsilon_t^2 + bh_t. \tag{5}$$

A few examples of applications are Akgiray (1989), Baillie and Bollerslev (1989), Baillie and DeGennaro (1990) and Hsieh (1989).

Nelson (1991) introduces models whose conditional variances are an asymmetric function of the residuals ε_t. The Exponential ARCH(1,0) model involves standardised residuals, z_t,

$$z_t = \varepsilon_t / h_t^{\frac{1}{2}}, \tag{6}$$

and an AR(1) specification for $\ln(h_t)$:

$$\ln(h_{t+1}) - \lambda = \rho[\ln(h_t) - \lambda] + \theta z_t + \gamma(|z_t| - E[|z_t|]). \tag{7}$$

Examples of equity studies are Nelson (1991) and Poon and Taylor (1992).

3.2. Tests of the informational efficiency of the options market

As noted in Section 2.3, an estimate $\hat{\alpha}_t^2$ of expected short-term squared volatility can also be obtained from a term structure model for option prices. The estimate can be rescaled to give a variance estimate for one period rather than an annualised quantity. The estimate is then for the following unobservable conditional variance

$$\alpha_t^2 = \mathrm{var}(R_{t+1} \mid M_t). \tag{8}$$

Here α_t is not an annualised figure and M_t is the information used by options market agents when they set prices at time t. The set M_t is presumed to include observed returns I_t. Day and Lewis (1992a) use implied volatility from short maturity options to approximate α_t. To evaluate their methodology we use the implied volatility from an option with the least time to maturity but greater than nine calendar days.

Options prices will provide optimal predictions of volatility when options markets use information efficiently and the pricing model correctly specifies the relationship between prices and volatility expectations. Information other than options prices should not have incremental predictive power when this joint hypothesis is true.

To test the hypothesis that options prices give optimal one-period-ahead volatility predictions two ARCH models are estimated. The first model only uses options information, the second model also uses returns information. For a GARCH(1,1) specification the two models are

$$h_{t+1} = c + d\hat{\alpha}_t^2 \tag{9}$$

and

$$h_{t+1} = c + a\varepsilon_t^2 + bh_t + d\hat{\alpha}_t^2. \tag{10}$$

For the symmetric version of Exponential ARCH(1,0) the models are

$$\ln(h_{t+1}) = \lambda + \delta \ln(\hat{\alpha}_t^2) \tag{11}$$

and

$$\ln(h_{t+1}) = \lambda(1-\rho) + \rho \ln(h_t) + \gamma(|z_t| - E[|z_t|]) + \delta \ln(\hat{\alpha}_t^2). \tag{12}$$

Likelihood ratio tests of the null hypothesis $a = b = 0$ or $\rho = \gamma = 0$ can be evaluated by comparing $LR = 2(L_1 - L_0)$ with χ_2^2, with L_0 and L_1 the maximum log-likelihoods either for (9) and (10) or (11) and (12). These tests make strong and possibly optimistic assumptions about the asymptotic distribution of the likelihood ratio [4]. The relative information content of the two sets of estimates for α_t^2 can be assessed by comparing their maximum values of the log-likelihood.

Estimates $\hat{\mu}_t^2$ of the squared long-term expectation should have no incremental power to predict short-term volatility if market expectations are rational. This hypothesis is tested by including an additional term $e\hat{\mu}_t^2$ in (9) and (10) and a term $\eta \ln(\hat{\mu}_t^2)$ in (11) and (12).

3.3. Seasonal volatility effects

It is known that returns measured over more than 24 hours often have higher variances than 24-hour returns. The equations for h_t require revisions to allow for this seasonality. It is assumed that equations (5), (7) and (9)–(12) all apply to a non-seasonal conditional variance h_{t+1}^* defined to be the conditional variance h_{t+1} divided by a seasonal term. We replace h_t, h_{t+1}, and ε_t by h_t^*, h_{t+1}^* and ε_t^* (where $\varepsilon_t^* = h_t^{*\,1/2} z_t$) and assume

$$h_t/h_t^* = \begin{cases} 1 \text{ if close } t \text{ is 24 hours after close } t-1, \\ M \text{ if } t \text{ falls on a Monday and } t-1 \text{ on a Friday} \\ H \text{ if a holiday occurs between close } t \text{ and close } t-1. \end{cases} \tag{13}$$

ARCH models for h_t^* combined with (13) define appropriate models for h_t.

The options estimates $\hat{\alpha}_t^2$ and $\hat{\mu}_t^2$ need to reflect the seasonal effects measured by M and H. Annualised estimates $\hat{\alpha}_{A,t}^2$ have been converted into non-seasonal, daily estimates using appropriate calendar constants for currency markets, as follows:

$$\hat{\alpha}_t^2 = \hat{\alpha}_{A,t}^2/(196 + 48M + 8H). \tag{14}$$

3.4. The conditional distribution

Empirical evidence decisively rejects the hypothesis that the distribution of a return R_t conditional upon the information set I_{t-1} of past returns is Normal for high frequency data (Engle and Bollerslev, 1986; Baillie and Bollerslev, 1989; Taylor, 1994). Two empirically better conditional distributions are the scaled t and the generalised error distribution (GED) (Taylor, 1994). The shape of the condi-

[4] We recognise that asymptotic theory for ARCH models is difficult (Bollerslev et al. 1992, Section 2.6, and sometimes unreliable (Lumsdaine, 1995), consequently we use very small nominal significance levels.

tional distribution z_t then depends on the degrees-of-freedom ω for the scaled t and the tail-thickness parameter ν for the GED. Normal distributions are given by $\omega \to \infty$ and $\nu = 2$.

3.5. Estimation

Model parameters occur in the definitions of the non-seasonal conditional variance, the seasonal multipliers and the conditional distribution. All of these parameters can be simultaneously estimated by maximising the likelihood function for a set of observed returns and volatility expectations implied by options prices. The likelihood function for a given parameter vector is calculated from the conditional variances and the standardised return residuals (Bollerslev et al., 1992; Taylor, 1994).

4. Empirical results within-sample

All the results presented from estimating ARCH models are for the period from January 1985 to November 1989. Prices for 1990 and 1991 are used in Section 5 for ex ante forecast evaluations.

4.1. ARCH models that only use returns information

Initial comparisons are made between the GARCH(1,1) model and the symmetric Exponential ARCH(1,0) model (i.e. $\theta = 0$ in (7)) and between the three most popular conditional distributions, the Normal, the scaled-t and the GED. Comparisons are also made between two ways to define the conditional mean return: either always zero, which is reasonable for futures data, or the appropriate figure defined by five dummy variables, one for each day of the week. To maximise the log-likelihood an initial value for the conditional variance is required. Empirical evidence suggests that the choice of initial value does not matter much, and we report the results as if the initial value is an additional parameter.

The following conclusions are obtained for all models estimated. There is no significant increase in the maximum log-likelihood when dummy variables define the conditional mean, consequently it is assumed to be zero. There is no uniform statistical result across the currencies about the significance of the variance dummy variables, but all the point estimates are well above one. Consequently these dummy variables are included in all the models discussed here.

The results presented in Table 1 reflect the above conclusions. The results for the scaled-t distribution are not reported to save space because the GED distribution is always slightly superior to the scaled-t distribution. The differences in maximum log-likelihoods between the GED and the Normal conditional distribution all exceed 20 for the GARCH(1,1) and symmetric Exponential ARCH(1,0)

Table 1
Parameter estimates for GARCH(1,1) and Exponential ARCH(1,0) models. Maximum likelihood estimates and the maxima of the log-likelihood function for ARCH models fitted to daily BP/$, DM/$, JY/$ and SF/$ exchange rates between January 1985 and November 1989. The ARCH models contain seasonal dummy variables M and H for Mondays and holidays, and the estimates of these parameters are not given here. The conditional distribution is GED with tail-thickness parameter ν. The special case $\nu = 2$ defines a Normal distribution. The conditional mean is equal to zero.

GARCH(1,1): $h_t = h_t^*$, Mh_t^*, Hh_t^*, and $h_{t+1}^* = c + a\varepsilon_t^2(h_t^*/h\,t) + bh_t^*$

	$10^5 c/(1-a-b)$	a	$a+b$	ν	$\ln(L)$
BP	6.0576	0.0466	0.9839	2	4254.01
	5.9589	0.0391	0.9877	1.1854	4297.04
DM	6.1391	0.0837	0.9569	2	4292.89
	6.0004	0.0738	0.9583	1.2528	4327.31
JY	4.9830	0.1012	0.8712	2	4391.42
	5.1111	0.0960	0.9189	1.0336	4478.84
SF	7.4797	0.0642	0.9617	2	4152.34
	7.3561	0.0599	0.9611	1.3778	4173.67

Exponential ARCH(1,0): $h_t = h_t^*$, Mh_t^*, Hh^*, and $\ln(h_{t+1}^*) = (1-\rho)\lambda + \rho \ln(h_t^*) + \gamma(|z_t| - E[|z_t|])$

	λ	γ	ρ	ν	$\ln(L)$
BP	−9.4316	0.1206	0.9795	2	4251.67
	−9.8054	0.1095	0.9835	1.1802	4296.03
DM	−9.5816	0.1708	0.9594	2	4295.11
	−9.8108	0.1574	0.9617	1.2571	4329.18
JY	−9.8369	0.2155	0.8604	2	4388.97
	−9.9819	0.2089	0.9117	1.0299	4477.72
SF	−9.3939	0.1347	0.9689	2	4155.43
	−9.5704	0.1260	0.9687	1.3855	4176.14

models. As the Normal distribution is the GED with $\nu = 2$, doubling log-likelihood differences and comparing test values with χ_1^2 shows they are all statistically significant at the 0.1% significance level. A fat-tailed, non-Normal conditional distribution enhances the descriptive accuracy of the model. Comparing the log-likelihoods for GARCH(1,1) and symmetric Exponential ARCH(1,0) reveals very small differences. The differences between GARCH and Exponential ARCH are 1.01 for the Pound, −1.87 for the Mark, 1.12 for the Yen and −2.47 for the Franc when the conditional distribution is GED. There is thus no clear-cut difference between the two models.

Now we consider more general ARCH models with the GED conditional distribution. First, consider the asymmetric specification for the volatility response in the Exponential ARCH(1,0) model. The increases in maximum log-likelihoods are less than 0.7 for all four currencies; the hypothesis $\theta = 0$ can not be rejected at the 10% significance level. As noted by Taylor (1994), there are plausible theories for a negative θ in stock models but none for a non-zero θ in a currency model. Price and volatility innovations can therefore be assumed independent when pricing currency options.

Table 2

Parameter estimates for GARCH(1,1) models including the term structure volatility expectations, with GED conditional distributions. The terms $\hat{\alpha}_t$ and $\hat{\mu}_t$ are respectively short- and long-term volatility expectations obtained from a term structure model. $h_t = h_t^*$, Mh_t^*, Hh_t^*, and $h_{t+1}^* = c + a\varepsilon_t^2(h_t^*/h_t) + bh_t^* + d\hat{\alpha}_t^2 + e\hat{\mu}_t^2$

$10^5 \times c$	a	b	d	e	ν	$\ln(L)$
Panel A: British Pound						
0.0735 (3.81)	0.0391 (2.56)	0.9485 (108.35)			1.1854 (17.73)	4297.04
			1.0000		1.2297 (20.57)	4300.23
1.2065 (3.53)			0.8459 (9.19)		1.2307 (17.74)	4312.17
1.1858 (3.22)	0.0297 (0.64)	0.0000	0.8276 (8.07)		1.2313 (17.75)	4312.41
0.5960 (0.81)			0.7538 (6.21)	0.1771 (1.22)	1.2344 (17.73)	4313.02
0.6152 (1.01)	0.0138 (0.43)	0.0000	0.7456 (6.11)	0.1682 (1.15)	1.2347 (17.73)	4313.11
Panel B: Deutsche Mark						
0.2500 (6.02)	0.0738 (3.81)	0.8846 (52.78)			1.2528 (17.62)	4327.31
			1.0000		1.2494 (20.77)	4328.74
1.1266 (3.84)			0.7941 (10.23)		1.3289 (17.51)	4349.69
1.1260 (3.84)	0.0000	0.0000	0.7943 (10.24)		1.3282 (17.51)	4349.69
1.1262 (3.84)			0.7945 (10.24)	0.0000	1.3285 (17.51)	4349.69
1.1265 (3.84)	0.0000	0.0000	0.7944 (10.23)	0.0000	1.3285 (17.51)	4349.69
Panel C: Japanese Yen						
0.4147 (6.39)	0.0960 (3.13)	0.8229 (26.41)			1.0366 (18.87)	4478.84
			1.000		0.9982 (24.57)	4442.75
1.6907 (4.83)			0.7139 (6.82)		1.0278 (19.15)	4480.48
1.5515 (4.40)	0.1252 (2.28)	0.0000	0.6185 (5.77)		1.0419 (19.03)	4485.21
1.6854 (4.51)			0.7151 (6.37)	0.0000	1.0278 (19.18)	4480.48

Table 2 (continued)

$10^5 \times c$	a	b	d	e	ν	$\ln(L)$
1.5476 (4.40)	0.1255 (2.29)	0.0000	0.6185 (5.78)	0.0000	1.0462 (19.03)	4485.21
Panel D: Swiss Franc						
0.2860 (7.42)	0.0570 (3.52)	0.9041 (56.60)			1.3777 (17.42)	4173.67
			1.0000		1.4512 (19.40)	4181.43
1.3539 (2.83)			0.8864 (8.98)		1.4519 (17.12)	4192.13
1.3537 (2.83)	0.0000	0.0000	0.8864 (8.98)		1.4519 (17.12)	4192.13
1.3533 (2.83)			0.8864 (8.98)	0.0000	1.4519 (17.12)	4192.13
1.3536 (2.83)	0.0000	0.0000	0.8864 (8.98)	0.0000	1.4519 (17.12)	4192.13

Notes: The numbers in parentheses are t-statistics estimated using the hessian and numerical second derivatives. All parameters are constrained to be non-negative. When a parameter estimate is zero or smaller than 10^{-6} then no estimated standard error is reported.

Second, consider higher order ARCH models. Results not reported here show such models have nothing extra to offer. Results for the tests of the informational efficiency of the currency options market are reported later for GARCH(1,1) and symmetric Exponential ARCH(1,0) models with the GED conditional distribution. However, not unexpectedly, all the conclusions also stand for other distributions and higher order ARCH models.

4.2. Informational efficiency of the currency options market

Tables 2 and 3 present the model estimates used for tests of informational efficiency. Table 2 uses market volatility expectations for the next period given by the term structure model outlined in Section 2.3. Table 3 uses the short-maturity implied volatility.

Eq. (10) includes options market volatility information as an exogenous variable in the GARCH(1,1) model for the conditional volatility. The increases in maximum log-likelihoods compared with those from the standard GARCH(1,1) model, i.e. Eq. (5), are as follows:

	Pound	Mark	Yen	Franc
Table 2	15.37	22.38	6.37	18.37
Table 3	18.47	22.34	4.94	17.36

Table 3

Parameter estimates for GARCH(1,1) models including short-maturity implied volatilities, with GED conditional distributions. The term $\hat{\alpha}_t$ is the implied volatility for the shortest maturity option with more than nine calendar days to expiry. $h_t = h_t^*$, Mh_t^*, Hh_t^*, and $h_{t+1}^* = c + a\varepsilon_t^2(h_t^*/h_t) + bh_t^* + d\hat{\alpha}_t^2$

$10^5 \times c$	a	b	d	ν	$\ln(L)$
Panel A: British Pound					
0.0735	0.0391	0.9485		1.1854	4297.04
(3.81)	(2.56)	(108.3)		(17.73)	
			1.0000	1.2632	4315.33
				(18.75)	
0.0000			1.0313	1.2524	4315.49
			(18.75)	(17.76)	
0.0000	0.0061	0.0000	1.0250	1.2524	4315.51
	(0.19)		(15.74)	(17.75)	
Panel B: Deutsche Mark					
0.2500	0.0738	0.8846		1.2528	4327.31
(6.02)	(3.81)	(52.78)		(17.62)	
			1.0000	1.3196	4349.49
				(17.85)	
0.0000			0.9719	1.3320	4349.64
			(19.42)	(17.42)	
0.0000	0.0000	0.0429	0.9295	1.3321	4349.65
		(1.43)	(3.11)	(17.40)	
Panel C: Japanese Yen					
0.4147	0.0960	0.8299		1.0336	4478.84
(6.39)	(3.13)	(26.41)		(18.87)	
			1.0000	1.0470	4473.03
				(22.11)	
0.9941			0.8698	1.0245	4478.53
(2.08)			(6.02)	(19.13)	
0.6033	0.1202	0.4017	0.3922	1.0420	4483.78
(1.72)	(2.58)	(1.56)	(1.64)	(18.98)	
Panel D: Swiss Franc					
0.2860	0.0570	0.9041		1.3777	4173.67
(7.42)	(3.52)	(56.00)		(17.42)	
			1.0000	1.4669	4187.16
				(18.56)	
0.5761			0.9993	1.4480	4190.66
(1.22)			(10.23)	(17.09)	
0.2339	0.0000	0.2144	0.8211	1.4515	4191.03
(0.47)		(0.80)	(3.08)	(17.05)	

Notes: The numbers in parentheses are t-statistics estimated using the hessian and numerical second derivatives. All parameters are constrained to be non-negative. When a parameter estimate is zero or smaller than 10^{-6} then no estimated standard error is reported.

Doubling these increases and comparing these test values with χ_1^2 shows clearly that the hypothesis that the options prices have no incremental information content can be rejected at the 0.5% significance level for each currency.

The more important question is whether options market volatility information is sufficient for predicting the next day's conditional volatility. As Eq. (9) only includes options market volatility information in the conditional variance equation, it is nested within Eq. (10). Thus a likelihood ratio test can again be used. The decreases in maximum log-likelihoods are as follows:

	Pound	Mark	Yen	Franc
Table 2	0.24	0.00	4.73	0.00
Table 3	0.02	0.01	5.25	0.37

Doubling these numbers and comparing these test values with χ_2^2 shows that the null hypothesis that returns contain no volatility information in addition to that already conveyed by options prices can not be rejected for any currency at the 0.5% significance level. The null can be rejected at the 1% level for the Yen. To conclude, the options market was informationally efficient for the three European currencies, the Pound, the Mark and the Franc, from 1985 to 1989; however, no such firm conclusion can apply to the Yen.

When the options market forms rational volatility expectations, long-term expectations have no incremental power to predict short-term conditional volatility. This null hypothesis is tested by evaluating the increases in maximum log-likelihood when the term structure estimate of long-term volatility is an additional variable in Eqs. (9) and (10). The results in Table 2 show the null hypothesis must be accepted.

The hypothesis that market expectations from option prices are unbiased estimates for one-period-ahead future volatility implies $c = 0$ and $d = 1$ in Eq. (9). Likelihood ratio tests are once more appropriate. The results in Table 2 reject the hypothesis at the 0.5% level for all four currencies. The results in Table 3 show much less evidence for bias. The term structure expectations for the next day are extrapolations and this can explain the bias identified by Table 2.

The maximum log-likelihoods in Tables 2 and 3 are very similar for our preferred model (Eq. 9, so $a = b = 0$). The higher values are in Table 2 for three currencies and in Table 3 for the Pound.

Tests of the informational efficiency of the options market using the symmetric Exponential ARCH(1,0) specification lead to the same conclusions as those reported above for the GARCH(1,1) specification.

5. Out-of-sample volatility forecasting

The tests in Section 4 characterise within-sample properties of volatility information because the likelihoods of both the ARCH and term structure models are maximised over the complete sample period. The direction of any within-sample biases is unknown (Day and Lewis, 1992a). In this section we compare the *ex ante* forecasting ability of historical volatility predictors, forecasts from standard ARCH models and options market forecasts over a longer time horizon than considered in Section 4.

The data cover seven years and we require a large sample to estimate the parameters in both the ARCH models and the volatility term structure model. Less than two and a half years of data remain after using five years to select and estimate the ARCH and term structure models. A four-week forecast horizon allows us to evaluate 30 non-overlapping, four-week ahead forecasts for the period from October 18, 1989 to February 4, 1992.

The non-seasonal, realised volatility is calculated ex post as follows:

$$V_{\mathrm{R},t} = \sqrt{\frac{1}{N}\sum_{i=1}^{N} \varepsilon_{t+1}^{*2}} \tag{15}$$

where N is the number of trading periods in some four-week interval. Note that the non-seasonal quantity ε_t^{*2} is one of R_t^2, R_t^2/M or R_t^2/H, with the choice determined by Eq. (13). The benchmark forecast is the simple historical volatility over the last four weeks, i.e.

$$V_{\mathrm{H},t} = \sqrt{\frac{1}{N}\sum_{i=0}^{N-1} \varepsilon_{t-i}^{*2}} = V_{\mathrm{R},t-N}. \tag{16}$$

The ARCH forecast for the realised volatility of the returns over N future periods can be obtained from N single-period forecasts all made at the same time. In the case of the GARCH(1,1) model, the volatility forecast can be calculated by:

$$V_{\mathrm{G},t} = \sqrt{\frac{1}{N}\sum_{i=1}^{N} \hat{h}_{t+i}} \tag{17}$$

where

$$h_{t+1}^{*} = c + a\varepsilon_t^{*2} + bh_t^{*}$$

and

$$\hat{h}_{t+i} = \frac{c}{1-a-b} + (a+b)^{i-1}\left(h_{t+1}^{*} - \frac{c}{1-a-b}\right), \quad i = 1, 2, 3, \ldots$$

All of the preceding volatility measures are annualised by multiplying by $(196 + 48M + 8H)^{1/2}$.

Forecasts can be derived from option prices, both from the term structure model and from a matched maturity option. The forecasts from the term structure model, $V_{\mathrm{TS},t}$, can be calculated from Eq. (2) with T equal to 28 calendar days. The matched forecast at time t is the implied volatility for an option whose time to maturity is nearest to 28 calendar days, denoted $V_{\mathrm{M},t}$.

The parameters in both the GARCH model and the term structure model of implied volatility are re-estimated as new observations come in on a rolling basis. We use a constant sample size of 250 weeks of daily data by adding the latest four weeks of observations and deleting the first four weeks of observations in the previous sample [5].

Forecasting performance is initially evaluated using the mean forecast error (*ME*), the mean absolute error (*MAE*) and the root mean square error (*RMSE*), calculated from forecasts $V_{F,t}$ given by one of the five methods above and realised figures $V_{R,t}$ as follows:

$$ME = \frac{1}{n}\sum_{t\in S}(V_{F,t} - V_{R,t})$$

$$MAE = \frac{1}{n}\sum_{t\in S}|V_{F,t} - V_{R,t}|$$

$$RMSE = \left[\frac{1}{n}\sum_{t\in S}(V_{F,t} - V_{R,t})^2\right]^{1/2}$$

Here S indicates the set of times at which ex ante forecasts are produced and n denotes the number of forecasts made using each method [6]. For this study, $n = 30$.

The results, listed in Table 4, clearly demonstrate the superiority of the two volatility forecasts computed from options prices [7]. The smallest *RMSE* is obtained by the options forecasts for each currency. The differences in Table 4 between the two options forecasts are minimal. The options forecasts also have the smallest *MAE* for the Pound, the Mark and the Yen while the ARCH forecast with the GED conditional distribution achieves the smallest *MAE* for the Franc [8]. The *ME* values for both options forecasts are very small (less than 1% of the average volatility) for the Mark, the Yen and the Franc and are not statistically different from zero at the 5% level, although the *ME* nearest zero is obtained by the historical volatility forecast for the Pound and the Yen and by the ARCH forecast for the Mark and the Franc.

---

[5] We also estimated parameters and forecast the volatility on a continual updating basis, which only adds new observations. The results, not reported here in detail, are not very different from the results based on the rolling method; the only significant result, which is rather data related, is reported in footnote 8. Our results differ from the findings of Lamoureux and Lastrapes (1993), however the underlying assets, the number of observations in the rolling samples, and the forecasting horizon are all different.

[6] Note that an optimal forecast will not have $MAE = RMSE = 0$ because $V_{R,t}$ is only a point estimate of the asset's price volatility which is unobservable.

[7] Replacing $V_{R,t}$ and $V_{R,t}$ by $V_{R,t}$ in the definition of *ME*, *MAE* and *RMSE* provides identical conclusions to those presented in the text.

[8] Using the updating method to re-estimate parameters reduces the effect of the inferior data for the Swiss franc (footnote 1) and then the implied volatility forecast also has the smallest *MAE*.

Table 4
Comparisons of alternative out-of-sample volatility forecasts. The implied forecasts are estimates of the market's volatility expectation for the next 28 days obtained either from a term structure model or the option with maturity closest to 28 days. The numbers tabulated are mean forecast error (*ME*), mean absolute errors (*MAE*) and root mean square errors (*RMSE*)

	Historical voltatility	GARCH(1,1)		Implied volatility	
		Normal	GED	Term structure	Matched
Panel A: *British Pound* (Average Realised Volatility = 0.11433)					
ME	−0.001038	0.001737	0.002015	0.005234	0.004447
MAE	0.033370	0.029240	0.029384	0.028412	0.029083
RMSE	0.041493	0.036214	0.036531	0.032527	0.033399
Panel B: *Deutsche Mark* (Average Realised Volatility = 0.12121)					
ME	−0.001736	0.001095	−0.000083	0.000788	0.000279
MAE	0.032874	0.032247	0.031692	0.025364	0.025945
RMSE	0.040840	0.040525	0.039836	0.032931	0.034312
Panel C: *Japanese Yen* (Average Realised Volatility = 0.10585)					
ME	−0.000217	0.007983	0.008441	0.000469	−0.000369
MAE	0.037799	0.030106	0.030415	0.025612	0.025727
RMSE	0.045695	0.034854	0.035267	0.030865	0.030763
Panel D: *Swiss Franc* (Average Realised Volatility = 0.12789)					
ME	−0.001915	0.000674	0.000355	−0.001246	−0.002130
MAE	0.028903	0.023740	0.023654	0.024342	0.023883
RMSE	0.034371	0.029103	0.029066	0.028720	0.028315

The forecasts from the ARCH model offer a marked improvement over naive historical volatility forecasts for all four currencies. The ARCH forecasts have up to 20% smaller *MAE*s and *RMSE*s than naive historical volatility forecasts. Comparing the ARCH forecasts with different conditional distributions reveals that forecasts from non-Normal models do not convincingly outperform forecasts from Normal models. Some caution should be exercised in interpreting these results as we only predict 30 four-week realised volatilities.

Lamoureux and Lastrapes (1993) perform encompassing regressions of the realised volatility on their three alternative out-of-sample forecasts and argue that the regressions provide further insight into the nature of the different forecast models. Like them we are interested in the incremental predictive power of the forecasts. Care is required because all forecasts are highly correlated. We apply the stepwise regression technique to select statistically significant regressors from the forecasts. The only forecast selected for the Pound, the Mark and the Yen is the term structure implied predictor (5% significance level, *F*-test). This implies no bivariate predictor is significantly more accurate than the term structure predictor. A similar conclusion has been obtained by Day and Lewis (1992b) for oil futures. However, no forecast is significant for the Swiss franc. The poor performance in the case of the Swiss franc could well be due to the block of inferior data mentioned in footnote 1. The within-sample and out-of-sample methodologies give the same conclusion that the volatility forecast obtained from option prices is optimal: returns from the underlying asset do not contain significant incremental information for predicting future volatility.

Finally we perform the test of the null hypothesis that the volatility forecast from option prices is an unbiased estimate for the four-week ahead realised volatility. We run regressions of the realised volatility on the term structure forecast with and without a constant term. The results show that the unbiased hypothesis cannot be rejected. The slope coefficients are all very close to 1 when the constant term is suppressed.

6. Summary

This paper examines the informational efficiency of the currency options market at the Philadelphia Stock Exchange using an ARCH methodology. By using likelihood ratio tests, we find that volatility forecasts estimated from call and put options prices contain incremental information relative to standard ARCH specifications for conditional volatility which only use the information in past returns. This is found for all four currencies. Furthermore, when predicting one-period-ahead volatility, the hypothesis that past returns have no incremental information content in addition to the information conveyed by the options market can not be rejected for the Pound, the Mark and the Franc, and is only marginally rejected for the Yen. We also find that market agents are rational in forming their expectations about future volatility, as the long-term expected volatility has no extra power when predicting short-term volatility.

The out-of-sample volatility forecasts confirm the above results. The two implied volatility forecasts markedly outperform the forecasts from past returns (the historical volatility forecast and forecasts from ARCH models) when predicting four-week ahead realised volatility. Further tests support the hypothesis that the options market's volatility expectations are unbiased predictions of future volatility.

These results suggest that the Philadelphia currency options market is informationally efficient in setting prices and the volatility expectations in option prices provide superior predictors of both one-period-ahead and longer horizon (i.e. four weeks) conditional volatilities. This conclusion contrasts with the lack of informational efficiency identified for US stock options markets by Day and Lewis (1992a), Lamoureux and Lastrapes (1993) and Canina and Figlewski (1993). The superior informational efficiency of the currency options market is consistent with the arguments developed in the final section of Canina and Figlewski (1993): they expect efficiency to be enhanced in an environment which permits low cost arbitrage trading.

The ultimate efficiency test is whether excessive profits can be made by some trading strategy. As Harvey and Whaley (1992) find, even if superior forecasts of future volatility are made then abnormal returns are not necessarily possible when transaction costs are taken into account. This issue will be explored in future research.

Acknowledgements

Xu was at the Financial Options Research Centre, University of Warwick when the paper was written and revised. We thank two anonymous referees, Stewart Hodges, Michael Selby, Martin Walker and participants at the 1993 meeting of the European Finance Association for their helpful comments and advice. We thank the Philadelphia Stock Exchange for providing their currency options data.

References

Akgiray, V., 1989, Conditional heteroscedasticity in time series of stock returns: evidence and forecasts, Journal of Business 62, 55–80.

Baillie, R.T. and T. Bollerslev, 1989, The message in daily exchange rates: a conditional-variance tale, Journal of Business and Economic Statistics 7, 297–305.

Baillie, R.T. and R.P. DeGennaro, 1990, Stock returns and volatility, Journal of Financial and Quantitative Analysis 25, 203–214.

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

Bollerslev, T., 1986, Generalised autoregressive conditional heteroscedasticity, Journal of Econometrics 31, 307–327.

Bollerslev, T., R.Y. Chou and K.F. Kroner, 1992, ARCH modeling in finance: a review of the theory and empirical evidence, Journal of Econometrics 52, 5–59.

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

Chiras, D.P., and S. Manaster, 1978, The information content of option prices and a test of market efficiency, Journal of Financial Economics 6, 213–234.

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

Day, T.E., and C.M. Lewis, 1992b, Initial margin policy and volatility in the crude oil future market, Proceedings of the Options Conference, FORC, Warwick University.

Engle, R. and T. Bollerslev, 1986, Modelling the persistence of conditional variances. Econometric Reviews 5, 1–50.

Gemmill, G.T., 1986, The forecasting performance of stock options on the London Traded Options Market, Journal of Business Finance and Accounting 13, 535–546.

Harvey, C.R., and R.E. Whaley, 1992, Market volatility prediction and the efficiency of the S&P 100 index option market, Journal of Financial Economics 31, 43–73.

Heynen, R., A.G.Z. Kemna and T. Vorst, 1994, Analysis of the term structure of implied volatilities, Journal of Financial and Quantitative Analysis 29, 31–56.

Hsieh, D.A., 1989, Modelling heteroscedasticity in daily foreign-exchange rates, Journal of Business and Economic Statistics 7, 307–317.

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

Latane, H. and R.J. Rendleman, 1976, Standard deviation of stock price ratios implied by option premia, Journal of Finance 31, 369–382.

Lumsdaine, R.L., 1995, Finite sample properties of the maximum likelihood estimator in GARCH(1,1) and IGARCH(1,1) models: a Monte Carlo investigation, Journal of Business & Economic Statistics 13, 1–10.

Nelson, D.B., 1991, Conditional heteroscedasticity in asset returns: a new approach, Econometrica 59, 347–370.

Poon, S. and S.J. Taylor, 1992, Stock returns and volatility: an empirical study of the UK stock market, Journal of Banking and Finance 16, 37–59.

Scott, E. and A.L. Tucker, 1989, Predicting currency return volatility, Journal of Banking and Finance 13, 839–851.

Shastri, K. and K. Tandon, 1986, An empirical test of a valuation model for American options on futures contracts, Journal of Financial and Quantitative Analysis 10, 377–392.

Shastri, K. and K. Wethyavivorn, 1987, The valuation of currency options for alternate stochastic processes, Journal of Financial Research 10 (No. 2), 283–293.

Sheikh, A.M., 1991, Transaction data tests of S&P 100 call option pricing, Journal of Financial and Quantitative Analysis 26, 459–475.

Stein, J.C., 1989, Overreactions in the options market, Journal of Finance 44, 1011–1023.

Taylor, S.J., 1994, Modelling stochastic volatility, Mathematical Finance 4, 183–204.

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

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