The Term Structure of Volatility Implied by Foreign Exchange Options

Xinzhong Xu*

and

Stephen Taylor**

*Financial Options Research Centre
University of Warwick

**Dept of Accounting and Finance
The Management School, Lancaster University

May 1992

Helpful comments and advice from Les Clewlow, Gordon Gemmill, Teng-Suan Ho, Stewart Hodges and Richard Stapleton are gratefully acknowledged. We are grateful to the Philadelphia Stock Exchange for providing their currency options data.

Financial Options Research Centre
Warwick Business School
University of Warwick
Coventry
CV4 7AL
Phone: 0203 523606

FORC Preprint: 92/32

THE TERM STRUCTURE OF VOLATILITY IMPLIED BY FOREIGN EXCHANGE OPTIONS

Abstract

This paper illustrates methods for estimating the time-varying term structure of volatility expectations, as revealed by options prices. Short and long-term expectations can be estimated using the Kalman filter. These expectations are estimated for four currencies from 1985 to 1989 using daily PHLX options prices. Throughout the five-year period there were important differences between short and long-term expectations. The slope of the term structure changed frequently and there were significant variations in long-term volatility expectations. The four currencies had very similar term structures, particularly in 1988 and 1989.

I. Introduction

Options provide information about the expected future volatility of the underlying asset. Implied volatilities at any moment in time vary, however, for different times to option expiry T and different exercise prices X. A matrix of implied volatilities is frequently available, say with columns ordered by T and rows ordered by X. Rational expectations of the average volatility during the next T years should vary with T because volatility is known to be stochastic, as reflected by the considerable literature on ARCH models reviewed by Bollerslev, Chou and Kroner (1992). Thus the rows of the implied volatility matrix will give information about the term structure of expected future volatility when option traders are rational. This paper describes and illustrates methods for estimating this term structure from one row of the implied volatility matrix, corresponding to nearest-the-money options. It is also rational for the implied volatility to vary with X when the asset volatility is stochastic

(Hull and White (1987), Stein and Stein (1991, Table 1)). Bid/ask spreads and nonsynchronous options and asset prices can also cause implied volatilities to vary (Day and Lewis (1988)).

The term structure of implied volatilities has been discussed by Poterba and Summers (1986), Stein (1989), Franks and Schwartz (1991) and Heynen, Kemna and Vorst (1991). Only two values of T are considered at any moment of time in these papers. Any number of T values can be studied using the estimation methods presented here and the number can vary from day to day. Time series studies involving one implied volatility figure per day are reported in several papers, for example Merville and Pieptea (1989) and Day and Lewis (1992), but such studies ignore term structure effects because T then varies from day to day. Our empirical analysis shows that it is possible to estimate interesting time series models for both short-term and long-term expected currency volatility using the Kalman filter.

Several methods of estimating volatility appear in the literature, including standard deviations calculated from recent asset returns. Although the historical volatility is a reasonable estimate when volatility varies slowly, Merville and Pieptea (1989) note that it fails to capture instantaneous changes. There is a general consensus in options research that the volatility implied by options prices and extensions of the Black-Scholes pricing model is a better predictor of future volatility than a standard deviation predictor based on historical data (Latane and Rendleman (1976), Chiras and Manaster (1978), Gemmill (1986), Shastri and Tandon (1986), Scott and Tucker (1989)). Insight into the relative importance of historical and implied predictors can be obtained by including both predictors in the conditional variance equation of an ARCH model (Day and Lewis (1992)). Further applications of this method for making comparisons would be interesting especially if the implied predictor is an estimate of short-term expectations.

A composite implied volatility can be obtained from the implied volatility matrix using some weighting scheme. We assign all the weight in each column to the nearest-the-money option, as in Stein (1989) and other recent studies, for reasons given at the end of Section IV. Other weighting methods are discussed by Latane and Rendleman (1976), Chiras and

Manaster (1978), Whaley (1982), Day and Lewis (1988), Turvey (1990) and Franks and Schwartz (1991).

Only a few studies have investigated implied volatility behaviour over varying horizons *T*, i.e. the term structure of expected volatility. Patell and Wolfson (1979) used two maturities after an earnings announcement to help show that there is more expected price variability in the days surrounding an announcement. Poterba and Summers (1986) obtained regression estimates from implied volatilities for three and six-month maturities, calculated from weekly Value Line indices. They concluded that when three-month volatility expectations change, the expected volatility three to six months hence also changes, but by much less than the change in expected short-term volatility. They also concluded that volatility shocks are short-lived after analyzing the volatility implied by CBOE indices.

Stein (1989) directly examined the term structure of implied volatilities, using two daily time series on implied volatilities for S&P 100 index options over the period December 1983 to September 1987. The values of *T* were less than one month for the first series and between one and two months for the second series. Based on the assumption that the volatility is mean reverting, as supported by his data, Stein found that the elasticity of volatility changes is larger than suggested by rational expectations theory : long-maturity options tend to "overreact" to changes in the implied volatility of short-maturity options. However, Heynen, Kemna and Vorst (1991) found that their conclusion about overreaction depended on the model used to represent changes in asset price volatility. They considered one year of European Options Exchange data and two values of *T*, one varying between zero and three months, the other between six and nine months.

Franks and Schwartz (1991) have tested possible explanations for stock volatility changes using weekly implied volatilities for the U.K. FTSE index from May 1984 to December 1989. Results for a capital structure hypothesis are compared in their Table 2 for shortest-maturity implieds and all-maturity implieds.

Engle and Mustafa (1992) infer volatility expectations for varying *T* without calculating implied volatilities. Their method seeks the GARCH(1,1) model which best explains observed options prices.

In this paper we model the term structure of expected volatility and the time series characteristics of the term structure. Section II describes a simple specification for the term structure at one moment in time. The specification involves two "factors" representing short-term expected volatility and long-term expected volatility and is more general than the approach of Stein (1989). The term structure specification is particularly appropriate when a satisfactory model for asset prices is GARCH(1,1), reviewed by Bollerslev, Chou and Kroner (1992) and Taylor (1992). Asset prices then follow a random walk with conditionally heteroscedastic steps. Section III describes estimation methods. A Kalman filter formulation has many advantages and allows estimation of time series models for the long-term expected volatility and the spread between short and long-term expected volatility; examples are given for AR(1) models.

Our empirical examples are for spot currency options on the British Pound, German Mark, Japanese Yen and Swiss Franc quoted against the U.S. dollar. Daily implied volatilities are modelled for the five-year period from January 1985 to November 1989. Section IV describes the Philadelphia Stock Exchange options data. Section V presents the empirical estimates of the term structure. Quasi-maximum likelihood methods are employed. We find that volatility expectations revert from their short-term level towards their long-term level with a half-life of approximately four weeks. We find considerable time-variation in the spread between short-term and long-term expectations and frequent changes in the slope of the term structure. We also document significant time-variation in the long-term volatility expectation, which can be modelled either by an AR(1) process or a random walk. Finally, Section VI presents conclusions and suggestions for further research.

II. A model for the term structure

Volatility is defined in our term structure model in the usual way and is always expressed in annual terms. Thus the volatility for some time period is the annualised standard deviation of the change in the price logarithm during the same period of time. We suppose that each year is divided into n smaller intervals of time. These intervals might be calendar days or they might be trading days and so commence when a market closes on one day and end when the market next closes; alternatively the durations of the intervals might be one week.

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

$$\operatorname{var}(\ln P_{t+\tau} - \ln P_{t+\tau-1}), \quad \tau = 1, 2, \ldots. \tag{1}$$

where P refers to the price of the asset upon which options are traded. These expectations can be annualised by multiplying them by n. After doing this, let $\sigma_{t,t+\tau}$ denote the volatility expectation at time t for time interval $t+\tau$, so

$$\sigma^2_{t,t+\tau} = n\, \mathrm{E}[\operatorname{var}(\ln(P_{t+\tau}/P_{t+\tau-1})) \mid M_t] \tag{2}$$

where E[... | ...] denotes an expectation based upon specific information; any dividends are assumed to be paid continuously at a constant rate.

Our term structure model is intended to be as simple as is reasonably possible. The model supposes that the expectations $\sigma_{t,t+\tau}$ are functions of at most three parameters. The first is the short-term expectation α_t for the next time interval :

$$\alpha_t = \sigma_{t,t+1}. \tag{3}$$

The second parameter is the long-term expectation μ_t given by assuming that the expectations converge for distant intervals :

$$\mu_t = \lim_{\tau \to \infty} \sigma_{t,t+\tau}. \tag{4}$$

Expectations are assumed to revert towards the time-dependent level μ_t as τ increases. The third parameter, ϕ, controls the rate of reversion towards μ_t and ϕ is assumed to be the same for all t. It is more practical to suppose that reversion applies to variances than to standard deviations, as follows :

$$\sigma^2_{t,t+\tau} - \mu_t^2 = \phi(\sigma^2_{t,t+\tau-1} - \mu_t^2), \quad \tau > 1. \tag{5}$$

It then follows that the expectation for time interval $t+\tau$ depends upon α_t, μ_t, ϕ and τ, thus :

$$\sigma^2_{t,t+\tau} = \mu_t^2 + \phi^{\tau-1}(\alpha_t^2 - \mu_t^2), \quad \tau > 0. \tag{6}$$

Market agents have mean-reverting expectations when $0 < \phi < 1$. Stein (1989) used an equation similar to the special case of (6) given by constant μ_t. Constant expectations as τ varies, consistent with the Black-Scholes paradigm, are only obtained when $\phi = 0$ or $\phi = 1$. Our preference for a simple model only permits three shapes for a graph of $\sigma_{t,t+\tau}$ against τ. The expectations are either monotonic increasing or monotonic decreasing as τ increases, or they are the same for all τ. Graphs of the expectations cannot contain spikes, perhaps aligned with the anticipated release of particularly important information.

The preceding equations summarise expectations made at time t for unit time intervals commencing at later times. The expected volatility at time t for an interval of general length T, from time t to time $t+T$, is the square root of :

$$v_T^2 = \frac{1}{T}\sum_{\tau=1}^{T} \sigma^2_{t,t+\tau} = \mu_t^2 + \frac{1-\phi^T}{T(1-\phi)}(\alpha_t^2 - \mu_t^2), \tag{7}$$

here assuming that subsequent asset prices, $\{P_{t+\tau}, \tau > 0\}$, follow a random walk. The numbers v_T, $T = 1, 2, 3, \ldots$ define the term structure of expected average volatility at time t; note the units for T are time intervals in (7), not years. We are interested in using implied volatilities from options prices to estimate the time series $\{\alpha_t\}$ and $\{\mu_t\}$ and also the mean-reversion parameter ϕ. This can be achieved because (7) shows that v_T^2 is a linear function of α_t^2 and μ_t^2.

A typical example for the term structure is a 10% expected volatility for the very short-term, a 12% expected volatility 37 calendar days after the current time and a 14% expected volatility for very distant months. These figures use equation (6) and a representative ϕ value equal to 0.972, similar to the estimates documented in Section V. The expected average volatilities which might be used to value near-the-money options for a selection of maturities are then as follows, using equation (7) :

Maturity T	Volatility for pricing
30 days	11.4%
60 days	12.2%
90 days	12.7%
180 days	13.3%
360 days	13.7%

Changes in asset price volatility have been modelled using ARCH and other models by many researchers in recent years. The GARCH(1,1) model has provided a satisfactory description of many asset return series (Bollerslev, Chou and Kroner (1992)). Engle and Mustafa (1992) have used this model to infer volatility expectations from options prices. It is shown in the Appendix that the term structure we are proposing here would be rational for market agents if they believe that future asset returns are generated by the GARCH(1,1) model; $n^{-1}\alpha_t^2$ is then the conditional variance of the asset return from time t to time $t+1$ and $n^{-1}\mu_t^2$ is the limiting conditional variance of the return from $t+\tau$ to $t+\tau+1$ as $\tau \to \infty$.

III. Estimation methods

Two methods have been developed for estimating the term structure model. The first method seeks the best match between the model and a dataset of implied volatilities. This method makes few assumptions about the time series properties of the series $\{\alpha_t\}$ and $\{\mu_t\}$. The method is also very quick. The second method supposes that $\{\alpha_t\}$ and $\{\mu_t\}$ follow autoregressive processes (possibly with unit roots) and then uses the Kalman filter to provide estimates of both the term structure and the parameters of the models assumed for $\{\alpha_t\}$ and $\{\mu_t\}$. This method requires substantially more computer resources.

A. Notation

The time t is now supposed to count trading days. On day t there will be implied volatility information for N_t expiry dates, supposed to be represented by a single number for each expiry date. It is a feature of our datasets that N_t varies from day to day. Let $y_{j,t}$ denote the implied volatility for option expiry date j on day t and suppose the times to expiry are $T_{j,t}$, measured in calendar days, with $T_{1,t} < T_{2,t} < \ldots. < T_{N_t,t}$.

B. A quick method

Forward implied variances $f_{j,t}$ can be calculated from the implied volatilities. At time t, the forward figure for the time interval from $t+T_{j-1,t}$ to $t+T_{j,t}$ is :

$$f_{j,t} = \frac{T_{j,t}y_{j,t}^2 - T_{j,t-1}y_{j,t-1}^2}{T_{j,t} - T_{j-1,t}}. \tag{8}$$

This is an annualised figure. When $j = 1$, $T_{0,t} = 0$ in (8).

The forward implied variance can be compared with the expected value for the appropriate part of the term structure. The forward expected variance $g_{j,t}$ is

$$g_{j,t} = \frac{1}{T_{j,t} - T_{j-1,t}} \left(\sum_{\tau = T_{j,t-1}+1}^{T_{j,t}} \sigma_{C(t),C(t)+\tau}^2 \right), \tag{9}$$

where $C(t)$ is the calendar day count corresponding to the passage of t trading days and τ is measured in calendar days. From (6) it can be seen that the forward expected variance is a linear combination of α_t^2 and μ_t^2. The combination is

$$g_{j,t} = \mu_t^2 + x_{j,t}(\alpha_t^2 - \mu_t^2), \tag{10}$$

with

$$x_{j,t} = \frac{\phi^{T_{j-1,t}} - \phi^{T_{j,t}}}{(1-\phi)(T_{j,t} - T_{j-1,t})}, \tag{11}$$

assuming $\phi < 1$.

Let n now denote the number of days for which there are implied volatilities. We wish to find estimates of

$$\phi, \quad \alpha_1, \alpha_2, \ldots, \alpha_n, \quad \mu_1, \mu_2, \ldots, \mu_n$$

giving small values for the differences

$$e_{j,t} = f_{j,t} - g_{j,t}, \quad 1 \le j \le N_t, \quad 1 \le t \le n.$$

Our estimates are given by minimising sums of terms $e_{j,t}^2$ for various ϕ followed by choosing ϕ to be the value giving the smallest sum across all times t. We could estimate α_t and μ_t using the implied volatilities for period t alone, providing $N_t \ge 2$. These estimates are rather erratic because the differences $e_{j,t}$ are non-trivial. There are many possible explanations for non-trivial differences including bid/ask spreads, incorrectly priced options, mis-specification of the term structure model and non-synchronous implied volatilities. Less erratic estimates for period t can be obtained by using the implied volatilities for a time window $t-k$ to $t+k$. This will be a reasonable method when it can be assumed that the volatility term structure is approximately constant within the time window.

The estimation method can be summarised by three steps, supposing k has already been selected; the choice of a sensible value for k is discussed in Section V.B for our data. Step 1 involves selecting a set of plausible values for ϕ, say $\phi_1, \ldots, \phi_m$. Step 2 involves finding the best estimates $\hat{\alpha}_{i,t}, \hat{\mu}_{i,t}$ when $\phi = \phi_i$, $i = 1, \ldots, m$. As $g_{j,t}$ is a linear function of $x_{j,t}$, from (10), these estimates are given for period t by regressing $f_{j,s}$ on $x_{j,s}$, with $1 \le j \le N_s$ and $t-k \le s \le t+k$. From (10), the estimated intercept is $\hat{\mu}_{i,t}^2$ and the sum of the estimated slope and the estimated intercept is $\hat{\alpha}_{i,t}^2$. These estimates are obtained for $t = k+1, \ldots, n-k$ and the sum of the squared regression errors calculated, summing over the three variables j, s and t. Call the sum $S(\phi_i)$ when $\phi = \phi_i$. Step 3 gives $\hat{\phi}$ as the value which minimises $S(\phi_i)$ and the time series of estimates $\{\hat{\alpha}_t)$ and $\{\hat{\mu}_t\}$ as the regression estimates when $\phi = \hat{\phi}$.

C. The Kalman filter method

The expected squared volatility over any period of time is a linear function of the current values of $\alpha_t^2 - \mu_t^2$ and μ_t^2, from (7). This suggests a Kalman filter method is ideal for estimating the term structure, day by day, if a set of squared implied volatilities is considered to be the expected squared volatility (from the term structure model) plus a set of measurement errors which can be attributed to option mispricing, non-synchronous observations and other issues. The Kalman filter formulation has several attractive properties : (i) it permits comparisons of models for the time series behaviour of the state variables, (ii) all the parameters can be obtained by maximising a quasi-likelihood function, (iii) the number of observations N_t can vary from day to day and (iv) it can be extended to give results for several assets simultaneously thus permitting the identification of common factors in the term structures of similar assets.

There are many ways to define the state variables and to model their time series characteristics. One credible example is presented here and further examples are evaluated in Section V.C. We suppose $\{\alpha_t^2\}$ and $\{\mu_t^2\}$ are stationary processes and have the same mean value $\bar{\mu}$. The state variables are taken to be $\alpha_t^2 - \mu_t^2$ and $\mu_t^2 - \bar{\mu}$, which both have zero mean and are unlikely to be highly correlated with each other. This choice is preferred to $\alpha_t^2 - \bar{\mu}$ and $\mu_t^2 - \bar{\mu}$ because these variables will probably have substantial covariation. The simplest plausible model for each of the chosen state variables is an AR(1) process. Independence between the state variables will be assumed. This gives the following <u>state equations</u> :

$$S_t = \begin{pmatrix} \alpha_t^2 - \mu_t^2 \\ \mu_t^2 - \bar{\mu} \end{pmatrix}, \quad \text{a } 2 \times 1 \text{ vector,} \tag{12}$$

$$S_t = \begin{pmatrix} \phi_1 & 0 \\ 0 & \phi_2 \end{pmatrix} S_{t-1} + \varepsilon_t, \tag{13}$$

$$\mathrm{E}[\varepsilon_t] = 0, \quad \mathrm{E}[\varepsilon_t \varepsilon_t'] = \begin{pmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{pmatrix}. \tag{14}$$

The <u>observation equation</u> for the Kalman filter is written as :

$$Y_t = Z_t S_t + \xi_t. \tag{15}$$

Here Y_t is a $N_t \times 1$ vector of squared implied volatilities minus $\bar{\mu}$, S_t is the 2×1 vector of state variables which summarise the term structure of expected volatility, Z_t is a $N_t \times 2$ matrix of state coefficients and ξ_t is a $N_t \times 1$ vector of measurement errors. We have

$$Y_t = \begin{pmatrix} y_{1,t}^2 - \bar{\mu} \\ \cdot \\ \cdot \\ y_{N_t,t}^2 - \bar{\mu} \end{pmatrix} \quad \text{and} \quad Z_t = \begin{pmatrix} z_{1,t} & 1 \\ \cdot & \cdot \\ \cdot & \cdot \\ z_{N_t,t} & 1 \end{pmatrix}, \tag{16}$$

with

$$z_{j,t} = \frac{1 - \phi^{T_{j,t}}}{T_{j,t}(1-\phi)} \tag{17}$$

from (7). The measurement errors are assumed to have zero means. Specification of their covariance matrix H_t is far from straightforward. Our preliminary results were based on the assumption that this matrix is diagonal with all N_t diagonal terms equal to the same number :

$$H_t = \mathrm{E}[\xi_t \xi_t'] = \mathrm{diag}(\sigma_0^2, \ldots, \sigma_0^2). \tag{18}$$

Assuming uncorrelated measurement errors, so $\mathrm{E}[\xi_s \xi_t'] = 0$ when $s \neq t$, concludes the specification of this particular model.

Sequential application of the Kalman filter to increasing information sets $I_t = \{Y_1, Y_2, \ldots, Y_t\}$ yields the minimum mean square linear estimators (MMSLE) of the state variables, $\mathrm{E}[S_t \mid I_t]$, using standard updating equations; these can be found in Harvey (1989, Ch. 3). The MMSLE are 2×1 vectors from which can be calculated the $N_t \times 1$ prediction error vectors :

$$\nu_t = Y_t - Z_t \begin{pmatrix} \phi_1 & 0 \\ 0 & \phi_2 \end{pmatrix} \mathrm{E}[S_{t-1} \mid I_{t-1}] \tag{19}$$

and the term structure estimates :

$$\begin{pmatrix} \hat{\alpha}_t^2 \\ \hat{\mu}_t^2 \end{pmatrix} = \begin{pmatrix} \bar{\mu} \\ \bar{\mu} \end{pmatrix} + \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \mathrm{E}[S_t \mid I_t]. \tag{20}$$

Equations (12)-(18) specify a model having seven parameters, summarised by the vector

$$\theta = (\phi, \phi_1, \phi_2, \sigma_0^2, \sigma_1^2, \sigma_2^2, \bar{\mu}).$$

A quasi-maximum likelihood estimate of θ can be obtained because the likelihood function is the product of conditional densities $f(Y_t \mid I_{t-1})$ and these densities depend, through θ, upon the prediction errors ν_t and their covariance matrices $F_t = \mathrm{E}[\nu_t \nu_t' \mid I_{t-1}]$. Following the arguments of Harvey (1989, p. 126) the quasi-log-likelihood function is as follows, given by assuming the prediction errors are multivariate normal :

$$\ln L(Y_1, Y_2, \ldots, Y_n) = \sum_{t=1}^{n} \ln f(Y_t | I_{t-1}) = -\frac{1}{2} \sum_{t=1}^{n} (N_t \ln(2\pi) + \ln(\det F_t) + \nu_t' F_t^{-1} \nu_t) . \qquad (21)$$

This function can be maximised using standard subroutines. We used the NAG subroutine E04JAF for our optimisations.

The autoregressive models for the term structure defined by (12)-(14) include two innovation terms per trading day, given by the 2×1 vector ε_t. The possibility of creating an options portfolio at time *t-1* whose value at time *t* does not depend on ε_t and the change in the price of the underlying asset is relevant for any appraisal of the economic plausibility of (12)-(14). The construction of a portfolio which generates arbitrage profits appears to be impossible even if transaction costs are ignored, because options prices are non-linear functions of volatility and (12)-(14) are defined in discrete time. Continuous time extensions and analysis may provide more definite conclusions.

IV. Data and computation of implied volatility

A. The market

The Philadelphia currency options market is the world's leading exchange in European and American-style options on spot currencies, with markets in the Deutsche Mark, Japanese Yen, Swiss Franc, British Pound, French Franc, Australian Dollar, Canadian Dollar and European Currency Unit. Total volume in these contracts represented approximately $2 billion in underlying value each trading day in 1990. The expiry months always include March, June, September and December. Two nearby months are also traded so that $N_t = 6$ when trade occurs for all the expiry months.

B. Datasets

The primary source database for this study is the transaction report compiled daily by the Philadelphia Stock Exchange (PHLX). This report contains the following information for each option traded during a day: date of trade (before February 1987; for February 1987 onwards the date on which the report was compiled, usually one day later), the style (call or put, European or American) and currency, expiration month, exercise price, number of trades, number of contracts traded, and the opening, closing, lowest, and highest option prices and the simultaneous spot exchange rate quotes. Only the closing option prices have been used. The database contains options prices for the seven currencies mentioned above and the ECU from November 5, 1984 to November 21, 1989. However, the transaction report is not available for some trading days during the above period; for some others, the report is not complete or in a few cases is in some way 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. The WSJ options prices and the associated spot prices are not simultaneous; we discuss the consequences of this non-simultaneity in detail in Section V.

All the results presented in this paper are for the period commencing January 2, 1985. The prices for November and December 1984 are only used to commence the Kalman filter calculations.

The interest rates used are London euro-currency rates, collected from Datastream. This source provides overnight, seven days, one month, three months, six months and one year interest rates. For intermediate times, we simply use linear interpolation. There is unlikely to be a simultaneity problem with the option data as the trading times are similar.

The London euro-currency interest rates were chosen because they consist of the maximum number of different maturities that we could use to make the interest rates used in calculating implied volatility as accurate as possible. Furthermore, they ensure the foreign and domestic interest rates are contemporaneous and are offered by the same institutions.

C. Data selection and revisions

Results have been obtained for American style options on four currencies--British Pound, Deutsche Mark, Japanese Yen and Swiss Franc. Results for the other three currencies have not been sought because trading was thin, in particular during the early part of the period studied.

Two essential changes have been made to the original data. First, we changed all the report compilation dates to the appropriate trading dates. Second, as the options expire on the Saturday before the third Wednesday of the expiration month but settle on the third Wednesday of that month, we have multiplied each option premium by $e^{R_d(4/365.25)}$ with R_d the relevant domestic (i.e. dollar) interest rate. Several exclusion criteria were used to remove uninformative options records from the database. Five criteria are first listed and then explained. We use standard notation, with S the spot rate, X the exercise price, T the time to expiry measured in years and R_f the foreign interest rate.

i) Options with time to expiration less than ten calendar days.

ii) Options violating European boundary conditions:

$$c < Se^{-R_f T} - Xe^{-R_d T}, \qquad p < Xe^{-R_d T} - Se^{-R_f T}.$$

iii) Options violating American boundary conditions:

$$C < S - X, \qquad P < X - S.$$

iv) Options with premia less than or equal to 0.01 cents.

v) Options that are far in- or out-of-the-money: $X < 0.8S$ or $X > 1.2S$.

Criterion (i) was used to eliminate options with small times to maturity as the implied volatilities then behave erratically.

Criteria (ii) and (iii) eliminate options violating the boundary conditions for European and American options. As the American options could be exercised at any time up to expiration, both boundary conditions must be satisfied, otherwise a riskless arbitrage could arise. Where an option price violates a rational pricing bound there are good reasons for suspecting that trades could not be made at this price.

Criterion (iv) is used to exclude options for which the necessarily discrete market prices are particularly likely to distort calculations of implied volatility.

Criterion (v) is used to eliminate those options that are either deep in-the-money or deep out-of-the-money. As their implied volatilities are extremely sensitive to a small change in the option price, they could distort calculations of implied volatility. Furthermore these options trade without much volume and are thus unrepresentative.

Table 1 provides summary information about the distribution of the number of distinct maturities N_t after applying the above exclusion criteria. These figures show that $N_t \geq 3$ for approximately 80% of the days studied.

D. Computation of implied volatility

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

It was decided to calculate the implied volatilities only from the closing prices of the nearest-the-money options; the nearest-the-money option on some day for a specific T is the option whose exercise price minimises $| S - X |$. Nearest-the-money options were chosen because: (1) given the widely reported 'strike bias' (or so-called 'smile effect'), including out-of-the-money and in-the-money options would introduce further noise into the term structure estimates; (2) the approximation that the implied volatility of a rationally priced option will equal the mean expected volatility over the time to expiry is considered better for an at-the-money option than for all other options (Stein (1989), Heynen, Kemna and Vorst (1991)).

V. Results

A. Further comments on data selection

The results from some preliminary calculations made it clear that a few extreme outliers had an excessive influence on the model estimates. Consequently, further exclusion criteria were applied.

For each currency and separately for calls and puts we calculated the sample mean and standard deviation of all the implied volatilities during the five year period. As both the quick and Kalman filter methods are sensitive to extreme outliers, we removed all implied volatilities more than five standard deviations distant from the sample mean. Only 13 observations in total were removed; 5 for Yen calls, 4 for Mark calls, 2 for Franc calls and 1 each for Pound and Yen puts. On a few occasions the forward implied variance, used by the quick method, was negative and then we excluded all the observations on that date; on average, seven days were removed from each dataset.

The majority of the outliers can be explained by one of four facts. First, the nearest-the-money option is occasionally rather in or out-of-the-money when only one exercise price is traded for a particular *T*. Second, non-simultaneity between option prices and the underlying spot prices for the Wall Street Journal data can be serious. Third, although we remove options with time to expiry less than ten days there are still some very large observations for small times to maturity. Fourth, the market may not have been operationally efficient during the first six months we have studied. As noted by Sutton (1988, p. 3), "A genuinely efficient options market in terms of large size and competitive spreads did not develop until the second half of 1985.".

B. Results from the quick method

The quick method described in Section III.B produces an estimate of ϕ, the mean-reversion parameter for volatility expectations, by fitting term structures to implied volatilities over windows of $2k + 1$ trading days. Values of k between 0 and 10 have been considered. The estimates of ϕ change by 0.01 at most moving from $k = 0$ to $k = 1$ and then change little as k increases. The maximum variation in the estimates of ϕ over $1 \leq k \leq 10$ is for Pound puts with a range from 0.970 to 0.976. As the results are insensitive to the choice of k, the choice $k = 5$ has been made for all the subsequent results.

Table 2 shows that the estimates of ϕ are very similar across currencies, ranging from 0.968 to 0.980. The median of the eight ϕ estimates is 0.975 corresponding to a 'half-life' of 27 calendar days by solving the equation $0.975^h = 0.5$. The call and put estimates can be seen to be very similar : the two smallest estimates are for the Yen, the third and fourth in magnitude are for the Pound, the fifth and sixth are for the Mark and the two largest are for the Swiss Franc.

Table 2 also provides estimates of the parameter ϕ for two sub-periods, the first from January 1985 to June 1987 and the second from July 1987 to November 1989. Again the differences between the estimates from call and put options are very similar except for the Yen in the first period. Most of the estimates of ϕ in the second period are higher than their counterparts in the first period although the differences are fairly small. The range is much greater for the first period, from 0.939 to 0.979, than for the second period, from 0.976 to 0.981. The median values are 0.967 and 0.981, respectively for the first and second period, corresponding to 'half-lives' equal to 21 and 36 calendar days. These estimates of the 'half-life' provide the important result that the market does not expect short-term volatility shocks to persist for long, i.e. their effect is expected to disappear quickly.

Figure 1 summarises the term structure of volatility expectations for Mark calls. Very similar numbers are obtained for Mark puts as expected. Figure 1 shows the time series estimates of the 15-day expectations (from eq. (6) with $\tau = 15$) and the long-term expectations

μ_t. The 15-day expectations (eq. 6) are very similar to the 30-day expected average volatilities (eq. 7). We chose to plot estimates for 15-day expectations rather than expectations estimates for the next day to avoid extrapolation beyond the limits implied by our data; recall $T \geq 10$ calendar days in all the calculations.

Five conclusions are suggested by the plots of the time series. First, the difference between 15-day and long-term expectations is often several percent so the implied volatilities reveal a significant term structure. Second, the estimates of the 15-day and long-term expectations frequently crossover, so the slope of the term structure often changes. Crossovers occur, very approximately, at an average rate of once every two to three months. Third, the long-term expected volatility varies significantly. This will become more clear when the results from the Kalman filter are analysed. Fourth, as might be expected, the estimated 15-day volatility expectation is much more variable over time than the estimated long-term expectation. Finally, the implied volatility process may not have been stationary in the sense that the average level appears to have been higher in 1985 than in the later years 1986 to 1989 although historic estimates of volatility are also high in 1985.

In Figures 2 and 3 we plot together for all four currencies the estimated 15-day and long-term expected volatilities, from call options. It is clear that volatility expectations for the three European currencies have been extremely similar and the Yen has come closer to the common trend as time has progressed.

C. Results from the Kalman filter

There are seven parameters in the time-varying term structure model described in Section III.C. The parameter ϕ continues to measure the rate of reversion in volatility expectations towards the long-term level. The spread between short and long-term expected squared volatility is assumed to follow an AR(1) process with AR parameter ϕ_1, mean zero and residual variance σ_1^2. The long-term expected squared volatility is assumed to independently follow an AR(1) process with AR parameter ϕ_2, mean $\bar{\mu}$ and residual variance

σ_2^2.

The final parameter in III.C is the variance of the measurement errors when the model is fitted to squared implied volatilities. One parameter for the measurement error variances has been found to be insufficient to give a satisfactory model for our implied volatilities data. The magnitude of the measurement errors is larger on average for the WSJ observations because of non-simultaneous spot and options prices. Furthermore, we have noted that the magnitude of the measurement errors increases on average as T decreases, for both data sources. Our preferred model has nine parameters with three parameters (σ_P^2, σ_T^2 and σ_W^2) used to define the dispersion matrix for the measurement errors ξ_t. The following diagonal matrix is preferred :

$$H_t = \mathrm{E}[\xi_t \xi_t'] = \mathrm{diag}\left(\sigma_S^2 + \frac{\sigma_T^2}{T_{1,t}}, \ldots, \sigma_S^2 + \frac{\sigma_T^2}{T_{N_t,t}} \right), \tag{22}$$

$$\sigma_S^2 = \sigma_P^2 \quad \text{for PHLX prices,}$$

$$= \sigma_W^2 \quad \text{for WSJ prices.}$$

Results are discussed in some detail for this nine-parameter model and then for simplifications (e.g. $\phi_2 = 1$) and finally for more general models (e.g. H_t is not diagonal).

1. The preferred model

Table 3 gives the parameter estimates obtained by maximising the quasi-log-likelihood function (eq. 21) defined by the Kalman filter. Panel A presents the estimates and approximate standard errors for the complete five-year period from 1985 to 1989. The standard errors have been calculated from the information matrix using numerical second derivatives, although the reliability of the usual likelihood theory in this context is unknown to us because the matrices of state coefficients, Z_t, are time-dependent. Panel B presents the estimates for the two sub-periods, from January 1985 to June 1987 and from July 1987 to November 1989.

The square root of an estimate $\bar{\mu}$ is an estimate of the median level of volatility expectations. These median estimates are smaller for the Yen than for the European currencies and they decrease from the first sub-period to the second sub-period for all currencies.

The Kalman filter estimates of ϕ are very similar to the estimates for the quick method. The average of the Kalman estimate minus the quick estimate is almost zero and the differences only vary from -0.006 to 0.004. The Kalman filter estimates of ϕ range from 0.967 to 0.980 for the full samples, with median 0.974 and 'half-life' equal to 27 calendar days. The Kalman filter estimates of ϕ, like those for the quick method, are generally larger for the second sub-period. The median and 'half-life' for the first sub-period are 0.966 and 20 days, with a range from 0.947 to 0.981. The corresponding figures are 0.980, 35 days, 0.975 and 0.983 for the second sub-period.

Some models for asset returns imply estimates of ϕ and ϕ_1 should be similar if expectations are formed rationally. A GARCH(1,1) model for returns (with constant long-term expectations) is one example, as noted in the Appendix. The estimates of ϕ_1 are non-trivially smaller than the estimates of ϕ but the former estimates are associated with trading days and the latter estimates with calendar days. The median estimate of ϕ_1 for the full samples is 0.966 and the associated 'half-life' is 20 trading days or approximately 29 calendar days, compared with 27 calendar days for ϕ. The sub-period median estimates of ϕ_1 are very similar : 0.964 and 0.972.

All the estimates of ϕ_2 exceed 0.975 for the complete datasets and half of these estimates exceed 0.99. The 'half-lives' for the median estimates are 66 trading days for the complete period, 51 trading days for the first sub-period and 36 trading days for the second sub-period.

The penultimate column of Table 3 shows estimates of $\sigma_1^2/(1-\phi_1^2)$ which is the variance of the spread term. The variation in the spread term is similar across the sub-periods for three currencies but not for the Pound which has smaller values in the later sub-period. The final column gives estimates of $\sigma_2^2/(1-\phi_2^2)$ which is the variance of long-term expectations. The numbers document a substantial fall over the five years in the variability through time of these expectations. An approximate 95% probability interval for the long-term volatility expectation can be obtained from $\sigma_2^2/(1-\phi_2^2)$ and $\bar{\mu}$. An example is an interval from 10.4% to 13.3% for the Mark, using the call estimates for the later sub-period. A corresponding interval for 15-day volatility expectations can be calculated by additionally using $\sigma_1^2/(1-\phi_1^2)$ and ϕ. This gives 6% to 16% for the same Mark source.

The small estimated values of σ_P^2 and σ_T^2 indicate that the time-varying term structure model fits the PHLX data reasonably well. A very approximate standard deviation for the difference between an observed implied volatility ($y_{j,t}$) and the correct term structure value (v_t, eq. 7) is given by the square root of $(\sigma_P^2 + T_{j,t}^{-1}\sigma_T^2)/(4\bar{\mu})$ for PHLX observations. Typical values are 0.8% for a 15-day option and 0.4% for a 180-day option (from Mark calls, full sample). The relative inaccuracy of the WSJ source is confirmed by the higher estimates for σ_W^2 than for σ_P^2. The illustrative approximate standard deviations for WSJ observations increase to 1.0% and 0.8%, respectively for 15- and 180-day options.

Figure 4 compares the Kalman filter estimates of volatility expectations with the quick method estimates. It can be seen that the estimates of 60-day expected average volatilities (from eq. 7) are very similar and this is also true for 15-day and long-term expectations. Further figures, not presented here, indicate that the plotted series are less smooth for the filter method, particularly for the 15-day expectations, because the quick method uses overlapping eleven-day windows. Also, the differences between the expectations obtained from call and put options are more variable for the Kalman filter.

2. Simpler models

To help evaluate certain simplifications of the preferred specification of the time-varying term structure model, we present comparisons of the maximum quasi-log-likelihoods for the nine parameter model with the corresponding figures for special cases requiring less parameters. The usual likelihood-ratio tests provide some insight. Cautious interpretations of log-likelihood differences are necessary, however, not least because several model parameters may have varied during the five-year period. The results for seven simplifications are summarised in Table 4, panel A.

To emphasise that term structure effects exist the model has been fitted with the restriction that the spread term is always zero ($\sigma_1 = \phi = \phi_1 = \alpha_0^2 - \mu_0^2 = 0$). The maximum log-likelihood then falls by more than 700 for each of the eight datasets. The possibility of constant long-term expectations through time ($\sigma_2 = \phi_2 = 0$) is also not credible as the maximum log-likelihood always falls by more than 500 for this model. Likewise, we can confidently disregard the idea that the two sources provide implied volatilities of equal accuracy ($\sigma_P = \sigma_W$) and can reject the assumption that the model fits with the same accuracy for all times to expiry ($\sigma_T = 0$).

The joint hypothesis that both the spread between short and long-term expectations and the long-term expectation follow random walks ($\phi_1 = \phi_2 = 1$, $\bar{\mu}$ undefined) gives likelihood-ratio test values ranging from 20.52 to 39.02 which could be compared with χ_3^2 if we trust the usual asymptotic theory. The test values strongly suggest that the joint hypothesis is doubtful. The more plausible hypothesis that the long-term expectation alone follows a random walk ($\phi_2 = 1$, $\bar{\mu}$ undefined) can be accepted for the Pound and the Mark using standard theory and a 5% significance level.

The hypothesis that the spread term reverts towards zero through trading time (weekdays less holidays) at the same rate as the term structure displays reversion in calendar time towards long-term expectations ($\phi_1^{4.8} = \phi^7$) is supported by all the datasets with the maximum value of the likelihood-ratio test statistic equal to 1.50.

3. More general models

There are many ways to add a tenth parameter to the preferred model. The results for five generalisations are summarised in Table 4, panel B, although none of them give substantial improvements for a majority of the datasets. The generalisations nearly always change the estimates of ϕ and ϕ_1 by negligible amounts. A few estimates of ϕ_2 change non-trivially, especially for the Franc data.

A variation which deserves evaluation is to remove the assumption that $\{\alpha_t^2\}$ and $\{\mu_t^2\}$

have the same mean value, i.e. on average the term structure is flat. Figure 1 at first sight suggests that on average the term structure slopes upwards. Defining different means for $\{\alpha_t^2\}$ and $\{\mu_t^2\}$ gives a ten-parameter model. The difference between the square root of the estimated long-term mean and the square root of the estimated short-term mean ranges from a minimum of 0.002 for Yen calls to a maximum of 0.011 for Pound calls, implying a positive average slope. However, the increases in the maximum quasi-log-likelihoods are all small and insignificant.

The spread innovation is assumed to be uncorrelated with the long-term innovation in (14) which implies that there is no correlation between the spread and long-term variables. Adding a parameter for correlation between the innovation terms gives small correlation estimates; they vary from 0.03 to 0.28.

The covariance matrix H_t for the measurement errors is assumed to be diagonal in the preferred model. An extra parameter can be added by assuming that all the off-diagonal elements in the associated correlation matrix are equal. Except for the Swiss Franc, the estimated common correlation term is very small (range -0.05 to 0.04) and the changes in the log-likelihood are unimportant. There is far more correlation between the measurement errors for the exceptional currency, 0.28 for the calls and 0.13 for the puts with large changes in the log-likelihood. Three parameters define the diagonal terms of H_t in (22). Increasing this to four, by allowing σ_T^2 to differ for the PHLX and WSJ sources improves some of the model fits but has no discernible effect upon the six parameters which do not appear in H_t.

Figures 1 and 3 and the sub-period estimates of $\bar{\mu}$ suggest that the mean of the process for long-term expectations may have declined as time progressed. Replacing $\bar{\mu}$ by $\mu_0 + \mu_1 t$ leads to negative estimates of μ_1 as expected but the reductions in the log-likelihood are not large.

VI. Concluding remarks

Two ways to estimate the time-varying term structure of volatility expectations have been illustrated. The quick method is easy to apply and produces similar conclusions to the technically more demanding Kalman filter method. The filter method, however, also provides information about satisfactory time series models for short and long-term volatility expectations.

Our study of volatility expectations for four currencies provides five conclusions. First, there are significant term structure effects. Fifteen-day and long-term volatility expectations often differ by several percent which causes implied volatilities to vary significantly across maturities, as illustrated in Section II. Second, the term structure sometimes slopes upwards, sometimes downwards and its direction frequently changes. The direction changes, on average, approximately once every two or three months. Third, there are significant variations in long-term volatility expectations, although these expectations change more slowly than both short-term expectations and the spread between short and long-term expectations. Fourth, the term structures of the Pound, Mark, Swiss Franc and Yen at any moment in time are all very similar. Finally, there are non-stationary elements in the term structure in the sense that some of the parameters of the recommended model changed during the five years investigated. This can be seen in the sub-period estimates presented in Table 3.

Further research can appraise the rationality of observed options prices and implied volatilities since the estimates of the term structure can be used to seek mispriced options. Also, rational implied volatilities should convey more information about future short-term volatility than can be gleaned from the prices of the underlying asset. Estimation of an ARCH model for asset returns based upon historic returns and the short and long-term volatility expectations would be interesting. The conditional variance should then depend primarily on short-term expectations if the options market is efficient. Research into the common element of volatility expectations and individual currency effects should also be informative. It is noted in the Appendix that the changes in long-term expectations are inconsistent with a stationary ARCH model whilst the changes in the spread between short and long-term expectations are inconsistent with an integrated ARCH model. Thus we conjecture that either the expectations are irrational or something more accurate than state-of-the-art ARCH models is needed to describe asset returns. Research into the validity of these conjectures should be interesting.

Appendix : Some implications of GARCH(1,1) models for volatility expectations

First, it is shown that the volatility expectations summarised by (6) in Section II are rational if market agents believe prices and returns are generated by an extension of the stationary GARCH(1,1) model, specified in (A-1) and (A-2) below. At time t suppose it is believed that subsequent prices $P_{t+\tau}$, logarithmic price differences $R_{t+\tau}$ and their conditional variances $H_{t+\tau}$ will be given by

$$\ln P_{t+\tau} - \ln P_{t+\tau-1} = R_{t+\tau} = d + H_{t+\tau}^{1/2} Z_{t+\tau} \qquad \text{(A-1)}$$

and

$$H_{t+\tau} = c_t + \theta(R_{t+\tau-1} - d)^2 + (\phi - \theta)H_{t+\tau-1} \qquad \text{(A-2)}$$

for $\tau > 0$ with

(i) H_{t+1} and c_t given by the current information M_t,

(ii) d, ϕ and θ constants, $0 < \theta < \phi < 1$, and

(iii) the Z_t i.i.d. with mean 0 and variance 1.

Condition (i) is not standard. The standard stationary model has c_t equal to some constant c and H_{t+1} given by the information in the returns history $I_t = \{R_{t-i}, i \geq 0\}$. It is reasonable to assume that M_t is I_t plus additional information. From (A-1) and then (A-2) :

$$\begin{aligned} \text{var}(R_{t+\tau} \mid M_t) &= \text{E}[H_{t+\tau} \mid M_t] \\ &= c_t + \phi \text{E}[H_{t+\tau-1} \mid M_t] \\ &= c_t(1-\phi^{\tau-1})/(1-\phi) + \phi^{\tau-1} H_{t+1}. \end{aligned} \qquad \text{(A-3)}$$

As τ increases the conditional variance converges to $c_t/(1-\phi)$. Suppose one year equals n time intervals, so time $t+n$ is one year hence, and let

$$nH_{t+1} = \alpha_t^2, \quad nc_t/(1-\phi) = \mu_t^2. \tag{A-4}$$

Then (A-3) is the same as (6) since

$$n\,\mathrm{var}(R_{t+\tau}|M_t) = (1-\phi^{\tau-1})\mu_t^2 + \phi^{\tau-1}\alpha_t^2. \tag{A-5}$$

Second, to illustrate the potential for similar estimates of ϕ and ϕ_1 suppose μ_t^2 is a constant $\bar{\mu}$ so c_t is some constant c. Also suppose M_t is simply the returns history I_t. Then (A-1) and (A-2), with $\tau = 1$, imply

$$H_{t+1} = c + \phi H_t + \theta H_t(Z_t^2 - 1). \tag{A-6}$$

From (A-4), this is the same as

$$\alpha_t^2 - \bar{\mu} = \phi(\alpha_{t-1}^2 - \bar{\mu}) + \theta\alpha_{t-1}^2(Z_{t-1}^2 - 1). \tag{A-7}$$

This AR(1) model can be compared with (13), which has AR parameter ϕ_1, except the innovation variance in (13) does not depend on α_{t-1} as it does in (A-7).

Third, to see that the conclusions about expectations derived from the implied volatilities appear to be inconsistent with a standard GARCH(1,1) model ($c_t = c$) separately consider stationary and non-stationary models. In the stationary case, $\mathrm{var}(R_{t+\tau} \mid I_t)$ converges to $c/(1-\phi)$ as $\tau \to \infty$. If M_t is I_t plus information about the quantity of information expected in the near future (cf. Taylor (1992, Sec. 3.5)) then $\mathrm{var}(R_{t+\tau} \mid M_t)$ also converges to $c/(1-\phi)$, contradicting the strong evidence for time-variation in long-term expectations. In the non-stationary case, an integrated ARCH model is obtained by supposing c_t is constant and $\phi = 1$ and then $\mathrm{var}(R_{t+\tau} \mid I_t) = c(\tau - 1) + H_{t+1}$. For large τ it may be supposed that $\mathrm{var}(R_{t+\tau} \mid M_t) \approx c(\tau-1) + H_{t+1}$. The slope of the term structure then always has the same sign as c for sufficiently large τ. This contradicts the evidence that the slope varies significantly around an average value close to zero.

References

Barone-Adesi, G., and R.E. Whaley. "Efficient Analytic Approximation of American Option Values." *Journal of Finance*, 42 (June 1987), 301-320.

Bollerslev, T., Chou, R.Y., and K.F. Kroner. "ARCH Modeling in Finance : a Review of the Theory and Empirical Evidence." *Journal of Econometrics*, 52 (April 1992), 5-59.

Chiras, D.P., and S. Manaster. "The Information Content of Option Prices and a Test of Market Efficiency." *Journal of Financial Economics*, 6 (June/Sept. 1978), 213-234.

Day, T.E., and C.M. Lewis. "The Behaviour of the Volatility Implicit in the Prices of Stock Index Options." *Journal of Financial Economics*, 22 (October 1988), 103-122.

Day, T.E. and C.M. Lewis. "Stock Market Volatility and the Information Content of Stock Index Options." *Journal of Econometrics*, 52 (April 1992), 267-287.

Engle, R.F., and C. Mustafa. "Implied ARCH Models from Options Prices." *Journal of Econometrics*, 52 (April 1992), 289-311.

Franks, J.R., and E.S. Schwartz. "The Stochastic Behaviour of Market Variance Implied in the Prices of Index Options." *The Economic Journal*, 101 (November 1991), 1460-1475.

Gemmill, G.T. "The Forecasting Performance of Stock Options on the London Traded Options Market." *Journal of Business Finance & Accounting*, 13 (Winter 1986), 535-546.

Harvey, A.C. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge, UK : Cambridge University Press (1989).

Heynen, R., Kemna, A.G.Z., and T. Vorst. "Analysis of the Term Structure of Implied Volatilities." Working paper, Erasmus University, Rotterdam, The Netherlands (November 1991).

Hull, J., and A. White. "The Pricing of Options on Assets with Stochastic Volatilities." *Journal of Finance*, 42 (June 1987), 281-300.

Latane, H., and R.J. Rendleman. "Standard Deviation of Stock Price Ratios Implied by Option Premia." *Journal of Finance*, 31 (May 1976), 369-382.

Merville, L.J., and D.R. Pieptea. "Stock Price Volatility, Mean-reverting Diffusion, and Noise." *Journal of Financial Economics*, 24 (September 1989), 193-214.

Patell, J.M., and M.A. Wolfson. "Anticipated Information Reflected in Call Option Prices." *Journal of Accounting and Economics*, 1 (1979), 117-140.

Poterba, J.M. and L.H. Summers. "The Persistence of Volatility and Stock Market Fluctuations." *American Economic Review*, 76 (December 1986), 1142-1151.

Scott, E., and A.L. Tucker. "Predicting Currency Return Volatility." *Journal of Banking and Finance,* 13 (December 1989), 839-851.

Shastri, K. and K. Tandon. "An Empirical Test of a Valuation Model for American Options on Futures Contracts." *Journal of Financial and Quantitative Analysis*, 10 (December 1986), 377-392.

Stein, J.C. "Overreactions in the Options Market." *Journal of Finance*, 44 (September 1989), 1011-1023.

Stein, E.M., and J.C. Stein. "Stock Price Distributions with Stochastic Volatility : An Analytic Approach." *The Review of Financial Studies*, 4 (Winter 1991), 727-752.

Sutton, W. *The Currency Options Handbook.* Cambridge, UK : Woodhead-Faulkner (1988).

Taylor, S.J. "Modeling Stochastic Volatility". Working paper, Lancaster University, Lancaster, UK (April 1992).

Turvey, C.G. "Alternative Estimates of Weighted Implied Volatilities from Soybean and Live Cattle Options." *Journal of Futures Markets*, 10 (August 1990), 353-366.

Whaley, R.E. "Valuation of American Call Options on Dividend Paying Stock: Empirical Tests." *Journal of Financial Economics*, 10 (March 1982), 29-58.

TABLE 1
The distribution of the number of maturities traded

The numbers tabulated are percentages frequencies and the time to maturity is at least ten days								
N_t	BPC	BPP	DMC	DMP	JYC	JYP	SFC	SFP
1	3.8	5.1	0.1	1.0	1.0	6.8	1.8	4.2
2	17.1	20.6	9.0	14.0	15.0	17.1	16.9	21.8
3	27.4	31.1	21.2	22.4	24.9	19.9	27.6	31.0
4	30.6	24.9	30.2	27.4	29.1	24.9	27.3	26.1
5	15.6	16.0	28.2	27.5	23.4	21.8	20.1	13.7
6	5.4	2.3	11.2	7.7	6.6	9.5	6.4	3.2

Notes: 1. N_t is the number of different expiry months used for estimating volatility expectations.

2. BPC refers to British Pound calls, BPP to British Pound puts, etc.

TABLE 2
'Quick method' estimates
of the term structure parameter ϕ when k=5

Options	Full Sample (85.01-89.11)	Sub-Sample 1 (85.01-87.06)	Sub-Sample 2 (87.07-89.11)
BPC	0.973	0.967	0.980
BPP	0.973	0.965	0.983
DMC	0.976	0.967	0.986
DMP	0.977	0.972	0.983
JYC	0.968	0.939	0.979
JYP	0.972	0.964	0.976
SFC	0.978	0.979	0.978
SFP	0.980	0.978	0.981

TABLE 3

Panel A: Estimated Parameters for the Preferred Term Structure Model

Opt-ions	ϕ	ϕ_1	ϕ_2	$\sqrt{\mu}$	σ_P^2 (10^{-6})	σ_W^2 (10^{-6})	σ_T^2 (10^{-5})	$\sigma_1^2/(1-\phi_1^2)$ (10^{-5})	$\sigma_2^2/(1-\phi_2^2)$ (10^{-5})
BPC	0.9714 (0.0019)	0.9685 (0.0078)	0.9972 (0.0021)	0.1279 (0.0281)	1.14 (0.08)	3.82 (0.59)	9.03 (0.71)	6.39 (1.47)	8.53 (6.43)
BPP	0.9666 (0.0027)	0.9631 (0.0087)	0.9959 (0.0026)	0.1334 (0.0212)	2.08 (0.12)	8.99 (0.95)	6.99 (0.77)	7.79 (1.71)	8.06 (5.08)
DMC	0.9756 (0.0012)	0.9709 (0.0071)	0.9934 (0.0033)	0.1292 (0.0112)	0.63 (0.04)	3.36 (0.38)	5.98 (0.39)	4.92 (1.14)	3.29 (1.59)
DMP	0.9766 (0.0012)	0.9689 (0.0074)	0.9916 (0.0035)	0.1280 (0.0089)	0.37 (0.04)	14.16 (1.25)	7.26 (0.41)	4.77 (1.06)	2.65 (1.10)
JYC	0.9717 (0.0018)	0.9511 (0.0098)	0.9838 (0.0053)	0.1127 (0.0048)	0.99 (0.06)	3.99 (0.47)	4.87 (0.46)	3.90 (0.72)	1.18 (0.37)
JYP	0.9733 (0.0013)	0.9524 (0.0083)	0.9844 (0.0037)	0.1099 (0.0045)	0.33 (0.03)	2.57 (0.32)	5.09 (0.32)	3.56 (0.56)	0.95 (0.21)
SFC	0.9772 (0.0018)	0.9680 (0.0082)	0.9773 (0.0064)	0.1353 (0.0054)	1.35 (0.10)	5.47 (0.76)	10.84 (0.79)	5.66 (1.29)	3.03 (0.81)
SFP	0.9799 (0.0016)	0.9640 (0.0086)	0.9876 (0.0047)	0.1309 (0.0065)	0.53 (0.08)	9.02 (1.18)	13.24 (0.84)	4.82 (1.04)	2.26 (0.81)

Notes: 1. The implied volatilities used are those remaining after deleting observations more than five standard deviations distant from the full sample mean.
2. The numbers in brackets are the estimated standard deviations of the parameter estimates calculated from the information matrix using numerical second derivatives.

TABLE 3

Panel B: Estimated Parameters for the Preferred Term Structure Model

Options	ϕ	ϕ_1	ϕ_2	$\sqrt{\mu}$	σ_P^2 (10^{-6})	σ_W^2 (10^{-6})	σ_T^2 (10^{-5})	$\sigma_1^2/(1-\phi_1^2)$ (10^{-5})	$\sigma_2^2/(1-\phi_2^2)$ (10^{-5})
BPC - Sub 1	0.9651	0.9700	0.9978	0.1248	1.46	1.56	12.00	9.08	16.46
- Sub 2	0.9798	0.9737	0.9807	0.1167	0.47	7.99	6.26	4.16	0.40
BPP - Sub 1	0.9557	0.9652	0.9965	0.1414	2.55	8.85	8.35	14.61	12.94
- Sub 2	0.9811	0.9701	0.9692	0.1173	1.30	10.66	6.23	3.69	0.57
DMC - Sub 1	0.9667	0.9683	0.9921	0.1388	0.78	3.70	6.80	6.45	4.40
- Sub 2	0.9832	0.9764	0.9810	0.1193	0.36	3.24	5.62	4.11	0.31
DMP - Sub 1	0.9705	0.9636	0.9883	0.1384	0.42	7.74	8.45	6.21	3.22
- Sub 2	0.9816	0.9768	0.9803	0.1173	0.25	17.92	6.35	3.66	0.32
JYC - Sub 1	0.9474	0.9300	0.9847	0.1164	1.39	2.14	3.19	5.00	1.92
- Sub 2	0.9814	0.9616	0.9826	0.1084	0.05	3.76	9.95	4.19	0.39
JYP - Sub 1	0.9655	0.9443	0.9830	0.1130	0.35	2.75	5.31	3.34	1.54
- Sub 2	0.9752	0.9564	0.9834	0.1065	0.32	2.62	4.67	4.04	0.30
SFC - Sub 1	0.9783	0.9716	0.9648	0.1463	1.29	6.89	15.53	5.02	3.36
- Sub 2	0.9770	0.9579	0.9769	0.1230	1.07	5.52	7.18	6.34	0.52
SFP - Sub 1	0.9805	0.9596	0.9793	0.1409	0.51	6.82	17.98	4.37	2.42
- Sub 2	0.9790	0.9737	0.9849	0.1206	0.17	11.34	10.70	5.09	0.44

Notes: 1. The implied volatilities used are those remaining after deleting observations more than five standard deviations distant from the full sample mean.
2. The first sub-sample is from January 1985 to June 1987 and the second from July 1987 to November 1989.

TABLE 4

Comparisons of the maximum quasi-log-likelihoods for the preferred time-varying term structure model with the figures for alternative models

	Parameters	Changes in log-likelihood[a]		
		Minimum	Maximum	Significant[b]
Panel A. Simplifications				
Flat term structures	6	-1644.81	-720.71	8
Constant long-term expectations	7	-1413.26	-579.41	8
Measurement error variance :				
same for both sources	8	-466.56	-22.13	8
same for all T	8	-250.64	-52.05	8
Random walks for :				
spread and long-term expectation	6	-19.51	-10.26	8
long-term expectation	7	-6.80	-0.90	4
Same reversion rate in the spread and the term structure	8	-0.75	-0.01	0
Panel B. Generalisations				
Average spread not zero	10	0.05	1.21	0
State variables correlated	10	0.06	5.31	3
Correlated measurement errors	10	0.22	45.08	3
Mean long-term expectation varies with time	10	0.38	9.01	3
σ_T^2 depends on data source	10	0.01	36.33	4

(TABLE 4 continued)

Note : The simplifications and generalisations are defined completely in Section V.C.

[a] The change in the quasi-log-likelihood function is the maximum of the function for the particular simplification or generalisation minus the maximum for the preferred nine-parameter model. Each row of the table summarises eight changes, four for Pound, Mark, Yen and Franc call options and four for put options.

[b] Number of significant test values out of eight at the 5% level. In panel A the test value is minus twice the change and the null hypothesis is that the preferred model is no better than the simplification. In panel B the test value is twice the change and the null hypothesis is that the generalisation is no better than the preferred model. Test values are compared with a chi-squared distribution with degrees-of-freedom given by the number of extra parameters in the alternative hypothesis. The test values must be interpreted with caution.

Figure 1: Estimated Volatility Expectations
(*DMC – Quick Method*)

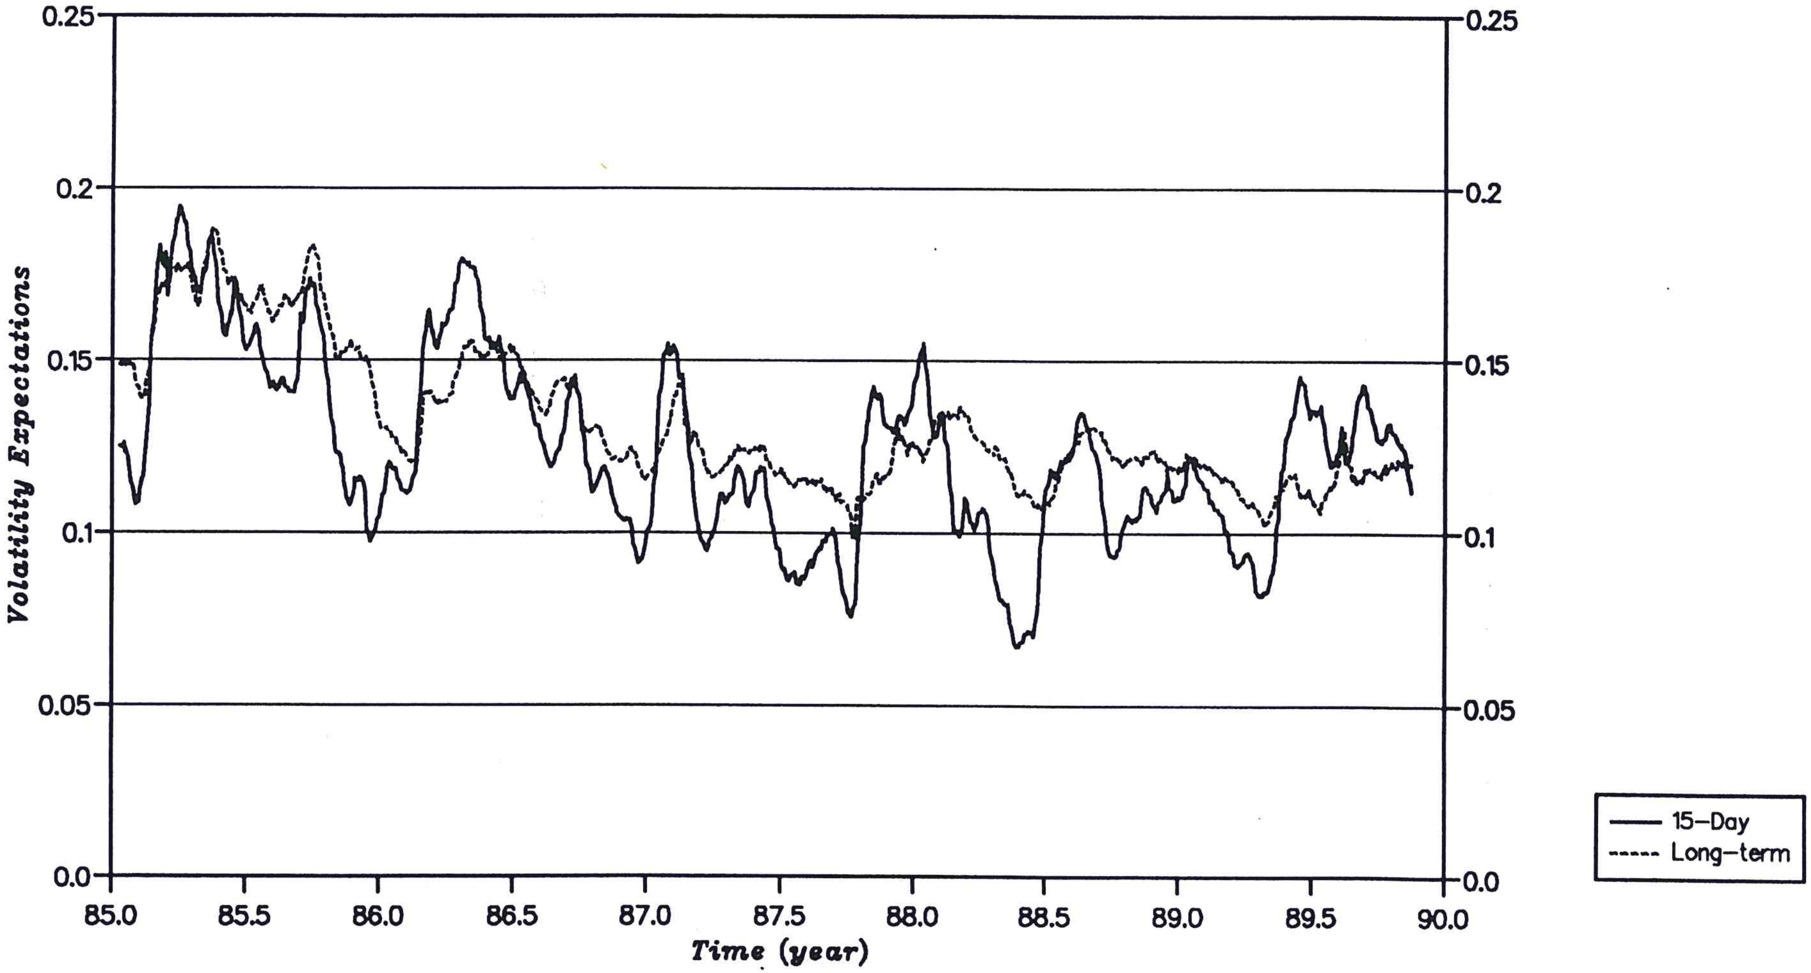

Figure 2: Estimated 15-day Volatility Expectations
(*Quick Method*)

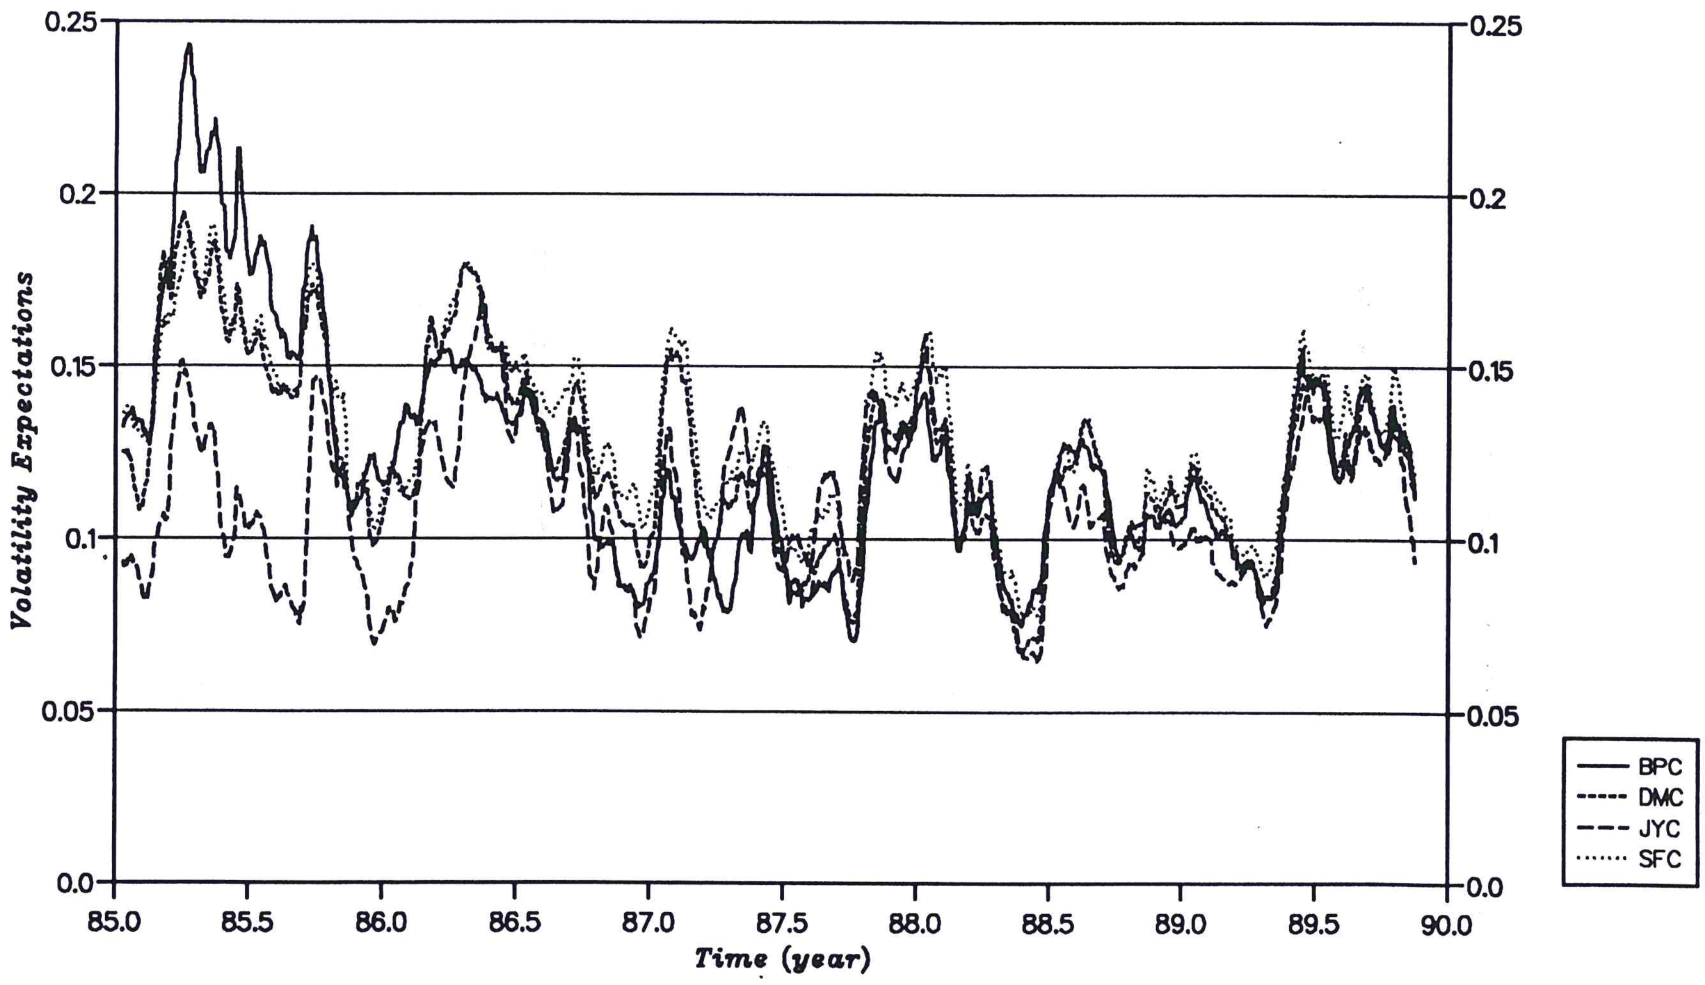

Figure 3: Estimated Long-term Volatility Expectations
(Quick Method)

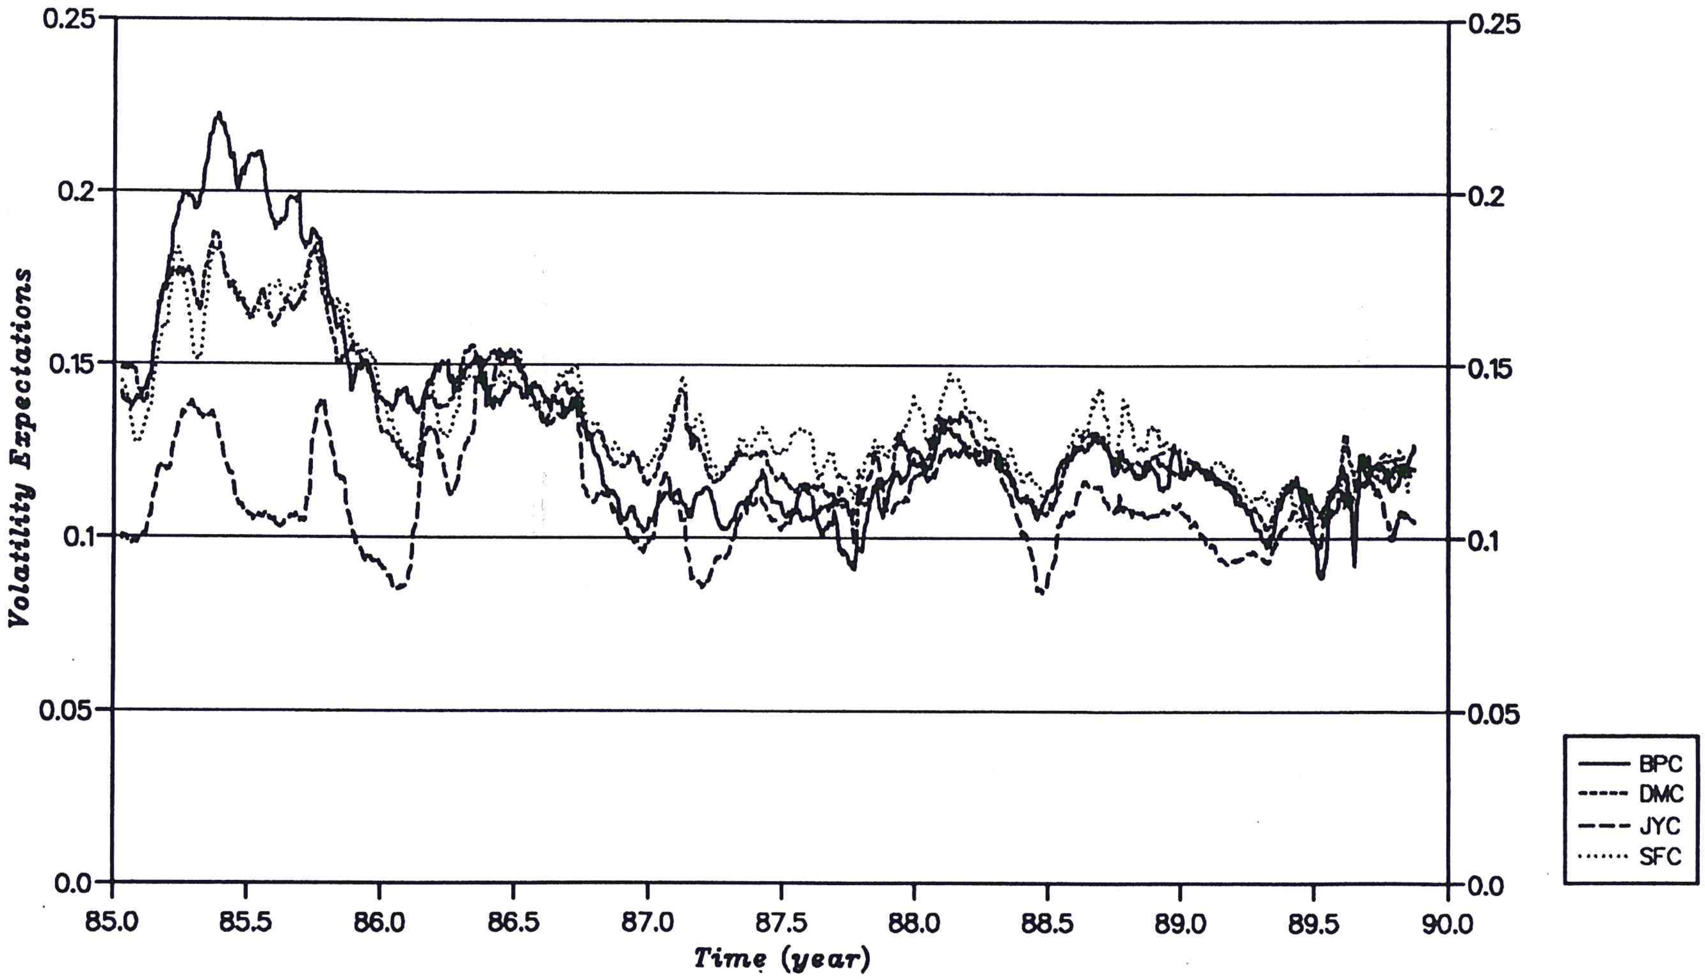

Figure 4: Estimated 60–Day Expected Average Volatilities
(Quick Method and Kalman Filter, DM Calls)