JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS VOL. 29, NO. 1, MARCH 1994

The Term Structure of Volatility Implied by Foreign Exchange Options

Xinzhong Xu and Stephen J. Taylor*

Abstract

This paper illustrates regression and Kalman filtering methods for estimating the time-varying term structure of volatility expectations revealed by options prices. Short- and long-term expectations are estimated for four currencies using daily PHLX options prices from 1985 to 1989. Throughout this 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 expectation estimates can be used to value OTC options, to improve hedging strategies, and to test the hypothesis that the options market overreacts.

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 will vary with T whenever volatility is believed to be stochastic. Thus, the rows of the implied volatility matrix may provide information about the term structure of expected future volatility. 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.

We model both 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. Thus, the specification is more general than the single factor approach of

*Department of Accounting and Finance, University of Manchester, England M13 9PL, and the Management School, Lancaster University, England LA1 4YX, respectively. Xu was at the Financial Options Research Centre, University of Warwick, when the paper was written and revised. The authors thank JFQA Managing Editor Jonathan Karpoff, an anonymous JFQA referee, Les Clewlow, Gordon Gemmill, Teng-Suan Ho, Stewart Hodges, Richard Stapleton, participants at seminars held at City University, Lancaster University, Oxford University, and Warwick University, and participants at the 1992 meetings of the European and French Finance Associations for their helpful comments and advice. The authors thank the Philadelphia Stock Exchange for providing their currency options data.

Stein (1989). The term structure specification is particularly appropriate when a satisfactory model for asset prices is GARCH(1,1). This model has often been recommended in empirical studies (see Bollerslev, Chou, and Kroner (1992)). Section III describes two 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.

The 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 modeled 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. Volatility expectations are shown to revert from their short-term level towards their long-term level with a half-life of approximately four weeks. There are considerable fluctuations in the spread between short- and long-term expectations and frequent changes in the slope of the term structure. There are also significant changes in long-term volatility expectations, which can be modeled either by an AR(1) process or a random walk. Section VI presents conclusions.

The term structure of implied volatilities has also been discussed by Poterba and Summers (1986), Stein (1989), Franks and Schwartz (1991), Diz and Finucane (1993), and Heynen, Kemna, and Vorst (1994). 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 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- and long-term expected currency volatility using the Kalman filter.

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 claimed 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. This conclusion has been disputed by Diz and Finucane (1993) following their analysis of similar S&P 100 index data. Furthermore, Heynen, Kemna, and Vorst (1994) 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.

Resnick, Sheikh, and Song (1993) have used stock options data to show that investors perceive monthly differences in return variability and, hence, the implied volatility term structure is not flat. The monthly differences are economically significant.

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 annualized standard deviation of the change in the price logarithm during the same period of time. It is supposed 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 they 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 form expectations of the quantities,

$$\text{(1)} \qquad \operatorname{Var}\left(\operatorname{Ln} P_{t+\tau} - \operatorname{Ln} P_{t+\tau-1}\right), \quad \tau = 1, 2, \ldots,$$

where P refers to the price of the asset upon which options are traded. These expectations can be annualized 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

$$\text{(2)} \qquad \sigma^2_{t,t+\tau} = n \operatorname{Var}\left(\operatorname{Ln}\left(P_{t+\tau}/P_{t+\tau-1}\right) \middle| M_t\right),$$

where M_t is the information set used by the options market.

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,

$$\text{(3)} \qquad \alpha_t = \sigma_{t,t+1}.$$

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

$$\text{(4)} \qquad \mu_t = \lim_{\tau \to \infty} \sigma_{t,t+\tau}.$$

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,

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

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

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

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 seasonal events or the anticipated release of particularly important information.

The preceding equations summarize 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

$$(7) \qquad v_T^2 \;=\; \frac{1}{T}\sum_{\tau=1}^{T}\sigma_{t,t+\tau}^2 \;=\; \mu_t^2 + \frac{1-\phi^T}{T(1-\phi)}\left(\alpha_t^2-\mu_t^2\right),$$

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.

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 Regression Method

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

$$(8) \qquad f_{j,t} \;=\; \frac{T_{j,t}y_{j,t}^2 - T_{j-1,t}y_{j-1,t}^2}{T_{j,t} - T_{j-1,t}}.$$

This is an annualized number. 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

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

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

(10) $$g_{j,t} = \mu_t^2 + x_{j,t}\left(\alpha_t^2 - \mu_t^2\right), \quad \text{with}$$

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

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 minimizing 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 nontrivial. There are many possible explanations for nontrivial differences including bid/ask spreads and nonsynchronous options and asset prices (Day and Lewis (1988)), incorrectly priced options, and misspecification of the term structure model. 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 summarized by three steps. 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 that minimizes $S(\phi_i)$, and the time series of estimates $\{\hat{\alpha}_t\}$ and $\{\hat{\mu}_t\}$ as the regression estimates when $\phi = \hat{\phi}$. When this method was implemented for our data, it was found that $\hat{\phi}$ is essentially the same for all k between one and 10. Consequently, results are reported in Section V.A with $k = 5$.

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 that can be attributed to option mispricing, nonsynchronous observations, and other issues. The Kalman filter formulation has several attractive properties: i) it permits comparisons of models for the time series behavior of the state variables, ii) all the parameters can be obtained by maximizing 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.B. 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 state equations,

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

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

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

The observation equation for the Kalman filter is written a

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

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 that summarize 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

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

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

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,

$$(18) \qquad H_t = E\left[\xi_t \xi_t'\right] = \operatorname{diag}\left(\sigma_0^2, \ldots, \sigma_0^2\right).$$

Assuming uncorrelated measurement errors, so $E[\xi_s \xi_t'] = 0$ 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, $E[S_t|I_t]$, 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 ,

$$(19) \qquad \nu_t = Y_t - Z_t \begin{pmatrix} \phi_1 & 0 \\ 0 & \phi_2 \end{pmatrix} E\left[S_{t-1}|I_{t-1}\right],$$

and the term structure estimates,

$$(20) \qquad \begin{pmatrix} \hat{\alpha}_t^2 \\ \hat{\mu}_t^2 \end{pmatrix} = \begin{pmatrix} \bar{\mu} \\ \bar{\mu} \end{pmatrix} + \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} E\left[S_t|I_t\right].$$

Equations (12)–(18) specify a model having seven parameters, summarized by the vector

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

A quasi-maximum likelihood estimate of θ can be obtained because the likelihood function is the product of conditional densities $f(Y_t|I_{t-1})$ and these densities depend, through θ, upon the prediction errors ν_t and their covariance matrices $F_t = E[\nu_t \nu_t'|I_{t-1}]$. 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,

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

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

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. Data Sources

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: a date (the trade date 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, 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 studied here 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 percent 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 nonsimultaneity in detail in Section V.B.

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, separately for calls and puts. 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 10 calendar days.
ii) Options violating European boundary conditions, $c < Se^{-R_f T} - Xe^{-R_d T}$, $p < Xe^{-R_d T} - Se^{-R_f T}$.
iii) Options violating American boundary conditions, $C < S - X, P < X - S$.
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.

Preliminary calculations showed that a very small number of extreme outliers ought to be removed because of their excessive influence on the model estimates. After all the exclusions, there are at least three maturities for between 75 percent and 90 percent of the days in each of the eight datasets 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 converges 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 minimizes $|S - X|$. Nearest-the-money options were chosen for two reasons. First, given the widely reported "strike bias" or "smile effect" (Shastri and Wethyavivorn (1987), Sheikh (1991), Taylor and Xu (1993)), including out-of-the-money and in-the-money options would introduce further noise into the term structure estimates; in theory, the smile effect can be a consequence of stochastic volatility (Hull and White (1987), Stein and Stein

((1991), Table 1)). 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 (1992), Heynen, Kemna, and Vorst (1994)).

V. Results

A. Results from the Regression Method

The regression 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 within windows containing $2k+1$ trading days. Table 1 lists the estimates of ϕ when $k = 5$. These 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 1

Regression Method Estimates of the Term Structure Parameter ϕ when $k = 5$

Options	Full Sample (85.01–89.11)	Subsample 1 (85.01–87.06)	Subsample 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

BPC refers to British pound calls, BPP to British pound puts, etc.

Table 1 also provides estimates of the parameter ϕ for two subperiods, 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 0.986. The median values are 0.967 and 0.981, respectively, for the first and second periods, 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 volatility shocks to persist for long, i.e., their effect is expected to disappear quickly.

Figure 1 summarizes 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 Equation (6) with

$\tau = 15$) and the long-term expectations μ_t. The 15-day expectations (Equation 6) are very similar to the 30-day expected average volatilities (Equation 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. Further figures, available from the authors, make clear that volatility expectations for the three European currencies have been extremely similar with the yen close to the European currencies after mid-1987.

FIGURE 1
Estimated Volatility Expectations
(Regression Method, DM Calls)

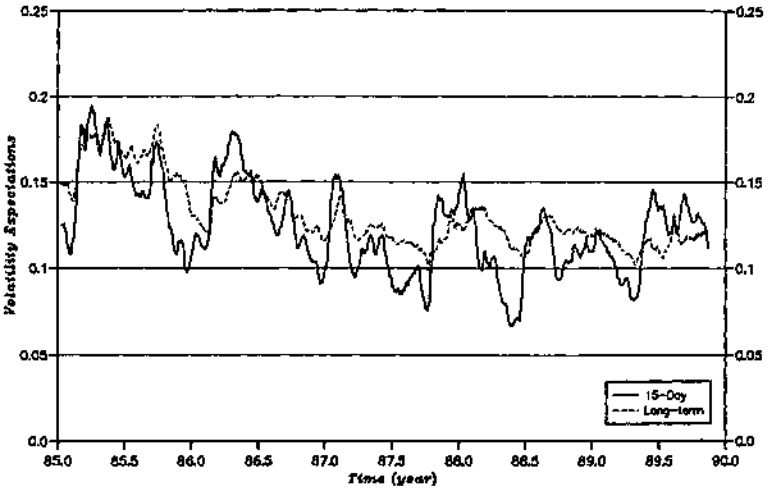

Five conclusions are suggested by Figure 1. First, the difference between 15-day and long-term expectations is often several percent so the implied volatilities reveal a significant term structure. Second, the estimates of the 15-day and long-term expectations frequently cross over, so the slope of the 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 clearer when the results from the Kalman filter are analyzed. 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.

B. 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 nonsimultaneous 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,

$$
\begin{aligned}
(22) \qquad H_t &= \mathrm{E}\left[\xi_t \xi_t'\right] = \mathrm{diag}\left(\sigma_S^2 + \frac{\sigma_T^2}{T_{1,t}}, \ldots, \sigma_S^2 + \frac{\sigma_T^2}{T_{N_t,t}}\right), \\
\sigma_S^2 &= \sigma_P^2 \quad \text{for PHLX prices,} \\
&= \sigma_W^2 \quad \text{for WSJ prices.}
\end{aligned}
$$

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 2 gives the parameter estimates obtained by maximizing the quasi-log-likelihood function (Equation 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 subperiods, 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 subperiod to the second subperiod for all currencies.

The Kalman filter estimates of ϕ are very similar to the estimates for the regression method. The average of the Kalman estimate minus the regression 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 regression method, are generally larger for the second subperiod. The median and "half-life" for the first subperiod are 0.966 and 20

TABLE 2

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})
Panel A. Estimates, 1985–1989[a]									
BPC	0.9714	0.9685	0.9972	0.1279	1.14	3.82	9.03	6.39	8.53
	(0.0019)	(0.0078)	(0.0021)	(0.0281)	(0.08)	(0.59)	(0.71)	(1.47)	(6.43)
BPP	0.9666	0.9631	0.9959	0.1334	2.08	8.99	6.99	7.79	8.06
	(0.0027)	(0.0087)	(0.0026)	(0.0212)	(0.12)	(0.95)	(0.77)	(1.71)	(5.08)
DMC	0.9756	0.9709	0.9934	0.1292	0.63	3.36	5.98	4.92	3.29
	(0.0012)	(0.0071)	(0.0033)	(0.0112)	(0.04)	(0.38)	(0.39)	(1.14)	(1.59)
DMP	0.9766	0.9689	0.9916	0.1280	0.37	14.16	7.26	4.77	2.65
	(0.0012)	(0.0074)	(0.0035)	(0.0089)	(0.04)	(1.25)	(0.41)	(1.06)	(1.10)
JYC	0.9717	0.9511	0.9838	0.1127	0.99	3.99	4.87	3.90	1.18
	(0.0018)	(0.0098)	(0.0053)	(0.0048)	(0.06)	(0.47)	(0.46)	(0.72)	(0.37)
JYP	0.9733	0.9524	0.9844	0.1099	0.33	2.57	5.09	3.56	0.95
	(0.0013)	(0.0083)	(0.0037)	(0.0045)	(0.03)	(0.32)	(0.32)	(0.56)	(0.21)
SFC	0.9772	0.9680	0.9773	0.1353	1.35	5.47	10.84	5.66	3.03
	(0.0018)	(0.0082)	(0.0064)	(0.0054)	(0.10)	(0.76)	(0.79)	(1.29)	(0.81)
SFP	0.9799	0.9640	0.9876	0.1309	0.53	9.02	13.24	4.82	2.26
	(0.0016)	(0.0086)	(0.0047)	(0.0065)	(0.08)	(1.18)	(0.84)	(1.04)	(0.81)
Panel B. Subperiod Estimates[b]									
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

[a]The numbers in parentheses are the estimated standard deviations of the parameter estimates calculated from the information matrix using numerical second derivatives.
[b]The first subsample is from January 1985 to June 1987 and the second from July 1987 to November 1989.

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 subperiod.

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 is one example. The estimates of ϕ_1 are nontrivially 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 subperiod 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 subperiod, and 36 trading days for the second subperiod.

The penultimate column of Table 2 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 subperiods for three currencies but not for the pound, which has smaller values in the later subperiod. 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-percent 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 percent to 13.3 percent for the mark, using the call estimates for the later subperiod. 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 percent to 16 percent 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, Equation 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 percent for a 15-day option and 0.4 percent 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 percent and 0.8 percent, respectively for 15- and 180-day options.

Figure 2 compares the Kalman filter estimates of volatility expectations with the regression method estimates. It can be seen that the estimates of 60-day expected average volatilities (from Equation 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 regression method uses overlapping 11-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 fewer 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 summarized in Table 3, Panel A.

To emphasize 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

FIGURE 2

Estimated 60-Day Expected Average Volatilities

(Regression Method and Kalman Filter, DM Calls)

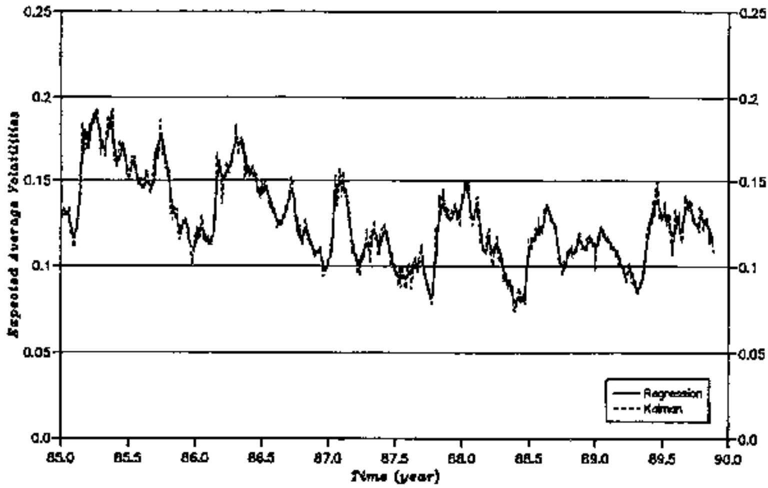

also not credible as the maximum log-likelihood always falls by more than 500 for this model. Likewise, we can confidently disregard the idea that the two sources provide implied volatilities of equal accuracy ($\sigma_P = \sigma_W$), and we can reject the assumption that the model fits with the same 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 expectations 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-percent 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 generalizations are summarized in Table 3, Panel B, although none of them give substantial improvements for a majority of the datasets. The generalizations nearly always change the estimates of ϕ and ϕ_1 by negligible amounts. A few estimates of ϕ_2 change nontrivially, especially for the Swiss franc data.

A variation that 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,

TABLE 3

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:				
Spreads and Long-Term Expectations	6	−19.51	−10.26	8
Long-Term Expectations	7	−6.80	−0.90	4
Same Reversion Rate in the Spread and the Term Structure	8	−0.75	−0.01	0
Panel B. Generalizations				
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

The simplifications and generalizations are defined completely in Section V.B.

[a]The change in the quasi-log-likelihood function is the maximum of the function for the particular simplification or generalization minus the maximum for the preferred nine-parameter model. Each row of the table summarizes 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-percent 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 generalization 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.

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 that do not appear in H_t.

Figure I and the subperiod 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 increases in the log-likelihood are not large.

VI. Conclusions

Two methods for estimating the time-varying term structure of volatility expectations have been illustrated. These methods assume that expectations revert monotonically from a short-term value towards a long-term level as the horizon of the expectations increases. The regression method is elementary and obtains many of the conclusions provided by the technically more demanding Kalman filter method. The filter method, however, also permits estimation of time series models for volatility expectations and comparisons between models. Independent autoregressive models for long-term expectations and the spread between short- and long-term expectations are preferred for currency markets.

Our study of volatility expectations for four currencies provides five further 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. Second, the term structure sometimes slopes upwards, sometimes downwards, and its direction (up or down) 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 have been very similar. Finally, there are nonstationary elements in the term structure in the sense that some of the parameters of the preferred autoregressive models changed during the five years investigated.

The volatility expectations provide insights into how the currency options market behaves. A constant volatility assumption is not made by the market. Volatility shocks are assumed to be transitory with an estimated half-life of approximately only one month. There is no evidence that the currency options market overreacts because this half-life is indistinguishable from the half-life for the mean-reverting spread between short- and long-term expectations. This is contrary to the equity results of Stein (1989), who assumed constant long-term expectations.

The volatility term structure estimates summarize the market's beliefs about volatility for all future periods. These estimates are expected to be more informative than forecasts obtained from historic prices alone. The estimates can be used to enhance hedging strategies and to value options for all maturities T including those that are not traded at exchanges.

Day and Lewis (1992) and Lamoureux and Lastrapes (1993) have estimated ARCH models for returns using implied volatility information, but disregarding term structure effects. Further research should estimate ARCH models for asset returns using the information in historic returns and short- and long-term volatility expectations. The conditional variance should then depend on short-term expectations alone if the options market is efficient. Xu and Taylor (1993) use this ARCH methodology and conclude that the PHLX currency options market is informationally efficient.

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

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 (Oct. 1988), 103–122.

__________________. "Stock Market Volatility and the Information Content of Stock Index Options." *Journal of Econometrics*, 52 (April 1992), 267–287.

Diz, F., and T. J. Finucane. "Do the Options Markets Really Overreact ?" *Journal of Futures Markets*, 13 (June 1993), 299–312.

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

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

Heynen, R.; A. G. Z. Kemna; and T. Vorst. "Analysis of the Term Structure of Implied Volatilities." *Journal of Financial and Quantitative Analysis*, 29 (March 1994), 31–56.

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

Lamoureux, C. G., and W. D. Lastrapes. "Forecasting Stock Return Variance: Toward an Understanding of Stochastic Implied Volatilities." *Review of Financial Studies*, 6 (Summer 1993), 293–326.

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

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

Resnick, B. G.; A. M. Sheikh; and Y. Song. "Time Varying Volatilities and Calculation of the Weighted Implied Standard Deviation." *Journal of Financial and Quantitative Analysis*, 28 (Sept. 1993), 417–430.

Shastri, K., and K. Wethyavivorn. "The Valuation of Currency Options for Alternate Stochastic Processes." *Journal of Financial Research*, 10 (No. 2, 1987), 283–293.

Sheikh, A. M. "Transaction Data Tests of S&P 100 Call Option Pricing." *Journal of Financial and Quantitative Analysis*, 26 (Dec. 1991), 459–475.

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

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

Taylor, S. J., and X. Xu. "The Magnitude of Implied Volatility Smiles: Theory and Empirical Evidence for Exchange Rates." *Review of Futures Markets* (forthcoming).

Xu, X., and S. J. Taylor. "Conditional Volatility and the Informational Efficiency of the PHLX Currency Options Market." *Journal of Banking and Finance*, 19 (forthcoming Feb. 1995).