

The Dynamic of the Volatility Skew: a Kalman Filter Approach

Mascia Bedendo and Stewart D. Hodges
Imperial College London and University of Warwick
Preliminary draft: please do not quote

January 20, 2005

Abstract

In the last few years, a lot of attention has been devoted to the issue of understanding and modeling the dynamic of implied volatility curves and surfaces, which is crucial for both trading, pricing and risk management of option positions. We suggest a simple, yet flexible, model, based on a discrete and linear Kalman filter updating of the volatility skew. From a risk management perspective, we assess whether this model is capable of producing good density forecasts of daily returns on a number of option portfolios, also in comparison with two alternative specifications, the sticky-delta model and the vega-gamma expansion. We find that our model clearly outperforms both alternatives, given its ability to easily account for movements of different nature in the volatility curve.

Keywords: Implied volatility, Kalman filter, density forecasting.

1 Introduction

It is well known that the volatilities implied from observed option prices are not constant across strikes and time to maturity, as the Black-Scholes model would predict. Instead, they exhibit a smile/skew pattern across strikes for a given time to maturity, which extends to an entire volatility surface when different expiries are examined. These implied volatility curves and surfaces also change through time (see Heynen [1993], Cont and da Fonseca [2001]), raising the need for an accurate modeling of their dynamic, which is essential for the purposes of option pricing, trading and risk management.

In order to explain these empirical deviations from the Black-Scholes model, various attempts have been made in the literature, which can be broadly grouped along the following directions.

A first class of models explains the existence of smiles/skews with market frictions, which imply the existence of an entire band of arbitrage-free option prices (see Figlewski [1989a, 1989b], and Longstaff [1995]).

Alternative option pricing models in which the dynamic of the underlying follows a process with jumps have been suggested, amongst the others, by Bates [1996], Carr, Geman, Madan and Yor [2002].

A third class of models explains the existence of the volatility curve with the variability of the volatility over time, which can be either stochastic or deterministic. In stochastic volatility models (see, for example, Hull and White [1987], Scott [1987], Heston [1993]), the instantaneous (or local) volatility itself follows a stochastic process. The main drawback of stochastic volatility and jump models is the inability to express directly the shape of the implied volatility curve (or, equivalently, surface) in terms of the model parameters. Therefore, the calibration of the model parameters to a set of market option prices becomes very difficult, and unrealistic parameters are often required to generate volatility curves or surfaces that are consistent with those implied by observed option prices (see, e.g. Bakshi, Cao and Chen [1997], Andersen and Andreasen [2000], Das and Sundaram [1999]).

The introduction of deterministic local volatility models, where the instantaneous volatility is modeled as a deterministic function of time and stock price, constituted a valid attempt to overcome this problem. By preserving market completeness, these local volatility models are self-consistent, arbitrage-free, and can be easily calibrated to match observed market volatility surfaces and curves. Relevant examples of these models are the non-parametric implied tree approaches suggested by Dupire [1994, 1997], Derman and Kani [1994a, 1994b], Rubinstein [1994] and Derman, Kani and Chriss [1996], and the parametric Normal Mixture Diffusion model of Brigo and Mercurio [2000, 2002] and Brigo, Mercurio and Sartorelli [2003].

Local volatility models perform very well at a static level, by attaining an exact (for non-parametric specifications) or almost exact (for semi-parametric and parametric forms) calibration to the implied volatility surfaces observed on a given day. However, as documented by Dumas, Fleming and Whaley [1998], the dynamic behavior of implied volatilities predicted by these models is inconsistent with the dynamic observed in the option markets. Besides raising the need for constant re-calibration of the model, this drawback leads, from a risk management perspective, to inaccurate and often unstable hedges of the option portfolios.

Motivated by the increasing evidence that option markets have become progressively more autonomous, showing movements in option prices driven not only by the underlying dynamic but also by specific sources of randomness, a new stream of literature has developed in the last few years. This innovative branch focusses on the identification of these extra sources of randomness in the option markets and consequently, on the investigation of their dynamic, also in relation to the dynamic of the underlying, in order to explain and capture the evolution of the empirical implied volatilities.

The implied volatility then becomes a financial state variable by itself. Its dynamic properties have been studied mainly by focussing on either the

term structure of at-the-money (ATM) implied volatility, or the volatility skew for a given maturity. Investigations of the dynamic followed by the entire volatility surface have begun to appear recently.

The most common approach to study the volatility dynamic consists in identifying the number and shapes of the shocks in the implied volatility via Principal Component Analysis (PCA) (see, for example, Skiadopoulos, Hodges and Clewlow [1999], Alexander [2001], Derman and Kamal [1997]).

The most recent contributions involve the specification of a deterministic or stochastic model for the implied volatility smile or surface, which fully describes their evolution through time. The deterministic implied volatility models introduced by Derman [1999] assume that either the per-delta or the per-strike implied volatility surface has a deterministic evolution. Due to their simplicity and tractability, the sticky-delta and the sticky-strike models are currently largely used by practitioners.

A stochastic evolution of the entire smile surface characterizes the stochastic implied volatility models by Schönbucher [1999], Ledoit and Santa-Clara [1998] and Cont, da Fonseca and Durrleman [2002]. In these studies, the prices of liquid vanilla options across a set of strikes and maturities are taken as given, in order to derive the initial volatility surface. The dynamic of the implied volatility is then modeled as a joint diffusion process with the underlying. In the first two models, constraints are imposed on the drift processes followed by the implied volatilities to ensure absence of arbitrage. An arbitrage-free specification of future volatility smiles, when the process for the underlying is unknown, has been investigated by Rebonato and Joshi [2003]. In a much simpler approach, Rosenberg [2000] proposed a stochastic evolution for the ATM implied volatility only, while keeping the shape of the curve fixed.

Our research study fits in the latter class of stochastic implied volatility models. We suggest a very simple model for the temporal evolution of the volatility skew for a given contract based on a linear Kalman filter approach. We first fit cubic polynomials to the empirical volatility curves. We then assume that the cubic's coefficients evolve through time according to a gaussian Ornstein-Uhlenbeck process correlated with the stock price process. The estimates for the skew coefficients are then promptly adjusted as information from new option trades becomes available.¹

Even though the Kalman filter technique seems a natural tool for financial problems of this kind, our study constitutes, to our knowledge, the first application of this powerful and robust econometric technique to the updating of the volatility skew. The specification we suggest is intuitively simple, easy to implement, and, since it accounts for stochastic non-parallel changes in the volatility skew, it is appealing not only for trading, but also for risk

¹In order to keep the analysis simple, we do not impose constraints to ensure an arbitrage-free specification.

management purposes.

In this study we investigate a simple application of the technique to the risk management of portfolios of futures options on the S&P500.² By implementing a standard Monte Carlo technique, we produce density forecasts of the daily changes in the marked-to-market value of four option portfolios, sensitive to shifts of different nature in the volatility curve. The distributions of the returns predicted by our model are then compared with the actual daily profits and losses (P&L) on the option portfolios, across one year. A comparison is also drawn with the portfolio returns estimated according to two methods which are widely used in practice to model changes in implied volatility. The first benchmark is the sticky-delta model by Derman [1999]. The second benchmark applies a Taylor expansion of Black-Scholes option prices to account for first and second order changes in the underlying (delta and gamma), and first order changes in the volatility level (vega). Various evaluation techniques are employed to assess the goodness of the P&L density forecasts. A special emphasis is placed on the accuracy of the tails, which are relevant for VaR computations.

Our main findings can be summarized as follows. First, producing accurate density forecasts of daily returns on option portfolios seems to be a hard task, even for fairly basic portfolios. Not too surprisingly, all our density forecasts show a significant bias in the mean, given that drift components are very difficult to estimate unless data is available over a long period of time (we only have three years of data). Also, daily forecasts can be quite noisy. Once we correct for the bias in the mean, the density forecasts generated from the Kalman filter model display a good fit to the actual portfolio returns, superior to that observed for the sticky-delta and vega-gamma alternatives. This result holds for all the kinds of option portfolios we consider, and indicates that more than one factor is needed to explain the dynamic of the volatility curve itself. The Kalman filter approach suggested in this paper seems to represent a simple, yet effective, way of taking this into account.

The paper is structured as follows. Section 2 presents the data set that

²From the risk management perspective, very little attention has been dedicated to the effects of the dynamic of the volatility skew/surface on the vega risk of a portfolio of options, and on the possible interaction with other risk factors. Most studies simply adjust the traditional Value-at-Risk (VaR) formulation to account for the Black-Scholes vega, without worrying about the existence of a skew/surface. Malz [2001] investigated a linear delta-vega VaR, Cárdenas *et al.* [1997] derived a closed formula for a delta-vega-gamma VaR, Glasserman, Heidelberger and Shahabuddin [2001] included vega risk in an efficient Monte Carlo simulation exercise. The most notable exception to this oversimplified treatment of the volatility risk is given by Malz [2001], who incorporated smile effects on VaR by letting implied volatilities of options with a different delta vary in correlated fashion. Also, Cont, da Fonseca and Durrleman [2002] and Fengler, Härdle and Villa [2001] hinted at the possibility of performing scenario simulations, for VaR purposes, of the joint evolution of the option portfolio and the underlying, on the basis of their principal components.

we use. In Section 3 we describe the Kalman filter model for the dynamic of the volatility skew, whereas the two alternative methods investigated in this study are presented in Section 4. Section 5 deals with the estimation of our model and the two benchmarks. Section 6 illustrates the criteria for building the option portfolios. Section 7 discusses and evaluates the density forecasts obtained from the Monte Carlo exercise. Section 8 concludes.

2 The Data Set

Our data set consists of daily data on quarterly³ futures options on the S&P 500 index traded at the Chicago Mercantile Exchange over the period 1998-2001. The first three years are employed to estimate the parameters of the models, whereas the assessment of their out-of-sample forecasting properties is carried out over the last year.

We only use closing option prices on the quarterly contract closest to expiry, except for the days within two weeks to expiration, when we roll on to the next contract. This is to ensure that we always refer to a very liquid contract with a wide range of strikes.

The usual no-arbitrage restrictions for futures options (see Hull [2002]) are applied to filter the option data. Also, we use the Barone-Adesi and Whaley [1987] approximation for American options to derive pseudo-European prices.

Throughout the entire analysis, we exclude in-the-money (ITM) options, which are less liquid and more sensitive to non-synchronicity pricing errors than out-of-the-money (OTM) options.

We also filter out some options with extreme strikes, which may have very low liquidity. In particular, we eliminate the options traded at the minimum price (tick), as well as those options for which a change in the premium equal to the tick size yields a change in the corresponding implied volatility larger than 15%. After applying all the relevant filters, we end up with an average number of strikes/options of 78 a day on the relevant futures contract, with a minimum of 44 and a maximum of 99. These options on average span a range of deltas (defined as the ratio between strike and futures price, K/F) between 0.66 and 1.18.

We employ daily closing prices of the index futures from the CME over the period 1998-2001 to estimate the model for the returns on the underlying.

³March, June, September and December expiries.

3 The Kalman Filter Model for the Dynamic of the Volatility Curve

The first step of our analysis is the choice of a method for fitting the volatility curve as a function of the level of moneyness. We opt for a cubic polynomial, such that:

$$\sigma_K = \alpha_0 + \alpha_1 * m + \alpha_2 * m^2 + \alpha_3 * m^3 + \varepsilon \quad (1)$$

m denotes the measure of moneyness which, following Natenberg [1994], is expressed as $\frac{\ln(K/F)}{\sqrt{\tau}}$, i.e. the natural log of the ratio between strike and underlying futures price, normalized by the square root of time to maturity.⁴ Since we cannot work with fixed time to maturity options, the normalization corrects for the effect of τ shrinking over time, and yields more meaningful estimates for the polynomial coefficients.

In Eq. (1), σ_K denotes the market implied volatility for an option with strike K , α_0 is the estimated level of the ATM implied volatility (with log moneyness equal to 0), and the coefficients α_1 , α_2 and α_3 capture, respectively, the slope, curvature and skewness of the volatility skew.

The cubic is easy to implement, has only four coefficients of immediate interpretation, and provides a good fit to the observed volatility curves. In Fig. 1 we report, as an illustrative example, plots of both the market and the fitted implied volatilities for the March 1999 contract, with time to expiration of one, two and three months.

The updating of the skew coefficients is performed by means of a discrete time linear Kalman filter (see Harvey [1989]). The 4×1 state vector of the coefficients, denoted as \underline{x}_t , evolves under the system equation:

$$\underline{x}_t = \underline{a}_t + U_t \underline{x}_{t-1} + C_t \underline{u}_t \quad (2)$$

$\underline{a}_t = A_t * \underline{\mu}_t$, where A_t represents the 4×4 matrix having mean reversion coefficients on the main diagonal, and $\underline{\mu}_t$ is the vector of the long run means. $U_t = I - A_t$, and $C_t C_t' = V_t$ is the covariance matrix of the error terms of the process. $\underline{u}_t \sim N(0, 1)$ are serially uncorrelated disturbances, independent of \underline{x}_{t-1} .

The skew coefficients are not directly observable. Instead, we observe an $n \times 1$ vector of market implied volatilities $\underline{\sigma}_{K,t}$, which is related to \underline{x}_t by the observation equation:

$$\underline{\sigma}_{K,t} = G_t \underline{x}_t + D_t \underline{v}_t \quad (3)$$

where G_t is the $n \times 4$ matrix with elements $[G]_{i,j} = m_i^{j-1}$, $D_t D_t' = R_t$ represents the $n \times n$ measurement noise covariance matrix, and $\underline{v}_t \sim N(0, 1)$ are again serially uncorrelated disturbances, independent of \underline{x}_t . Since the

⁴The time to maturity is computed as the number of calendar days to expiry divided by the total number of days in the calendar year.

number of observed implied volatilities can be small (even one only), this method is particularly suitable for practitioners, who need to update the volatility skew as soon as new trades occur.

We use $\hat{\underline{x}}_t$ and S_t to denote, respectively, $E_{t-1}[\underline{x}_t]$ and $\text{Var}_{t-1}[\underline{x}_t]$, before observing $\underline{\sigma}_{K,t}$. Finally, we assume that the initial distribution of \underline{x}_1 is multivariate Normal with known values for $\hat{\underline{x}}_1$ and S_1 .

According to this state space model, our optimal forecast $\hat{\underline{x}}_t$ for the skew coefficients \underline{x}_t (as well as its covariance S_t) is first adjusted according to the observed $\underline{\sigma}_{K,t}$. This updated quantity then evolves under the system equation, to produce the new optimal forecast $\hat{\underline{x}}_{t+1}$ for the next period. The updating equations for both expected value and covariance are given by:

$$\hat{\underline{x}}_{t+1} = U_{t+1} \left[I - S_t G_t' T_t^{-1} G_t \right] \hat{\underline{x}}_t + K_t \underline{\sigma}_{K,t} + \underline{a}_{t+1} \quad (4)$$

$$S_{t+1} = U_{t+1} \left[S_t - S_t G_t' T_t^{-1} G_t S_t \right] U_{t+1}' + C_{t+1} C_{t+1}' \quad (5)$$

where $T_t = G_t S_t G_t' + D_t D_t'$, and $K_t = U_{t+1} S_t G_t' T_t^{-1}$.

At each time t , after observing the market volatility skew, we produce one-day-ahead forecasts for the expected value of the coefficient vector $\hat{\underline{x}}_{t+1}$ and the covariance matrix S_{t+1} . In fact, at each step we produce complete density forecasts of the evolution of the skew, since $\underline{x}_{t+1} \sim N(\hat{\underline{x}}_{t+1}, S_{t+1})$.

In order to translate the density forecasts of the skew (expressed in terms of log moneyness) into forecasts for the distribution of the daily changes in the value of our option portfolios (expressed in terms of strikes), we also need to produce one-day-ahead forecasts of the distribution of the log returns on the underlying futures. We choose to model the returns on the underlying following an EGARCH(1,1) model with normally distributed errors (correlated with the errors from the coefficients' model), which is both simple to estimate and rich enough to allow for leverage effect and time varying volatility.

4 Two Alternative Models for the Dynamic of the Volatility Curve

The performance of the Kalman filter method at describing the dynamic of the volatility skew is compared with that of two alternative models widely used in practice.

4.1 The sticky-delta model

The sticky-delta model was first discussed by Derman [1999], who described it as the method commonly used by traders to predict the evolution of the volatility curve in a situation of stable trending markets.

Following the sticky-delta model, the volatility curve for a fixed maturity is parametrized as:

$$\sigma_K = \sigma_{ATM} - b(K - F) \quad (6)$$

where $b > 0$ yields the volatility skew. The fixed strike volatilities increase with the futures level, whereas σ_{ATM} is independent of the underlying and fixed. Therefore, the volatility curve changes if measured in terms of strikes, but it is fixed if measured with respect to moneyness (or, equivalently, delta).

The only factor that affects the volatility skew in this very popular model is the evolution of the underlying asset. As before, we model the dynamic of the equity futures with an EGARCH specification with normal errors.

4.2 The vega-gamma model

One of the standard approaches adopted by risk managers in order to measure the risk of an option portfolio, involves applying a Taylor series expansion of Black-Scholes option prices around the risk factors. The variables that affect option prices are the changes in the value of underlying asset and volatility (which are stochastic), and the changes in the time to expiry (which are deterministic).

According to this technique, the change in value of the option portfolio $\Delta\Pi$ is approximated by:⁵

$$\begin{aligned} \Delta\Pi &= \frac{\partial\Pi}{\partial F}\Delta F + \frac{\partial\Pi}{\partial\tau}\Delta\tau + \frac{\partial\Pi}{\partial\sigma}\Delta\sigma + \frac{1}{2}\frac{\partial^2\Pi}{\partial F^2}(\Delta F)^2 \\ &= \delta\Delta F + \Theta\Delta\tau + V\Delta\sigma + \frac{1}{2}\Gamma(\Delta F)^2 \end{aligned}$$

where δ (delta), Θ (theta), V (vega) and Γ (gamma) represent the sensitivities of the portfolio's value to variations in the risk factors.

Since, in our case, we deal with delta-neutral portfolios and we consider profits and losses on the portfolios at a daily level, the approximation can be simplified to:⁶

$$\Delta\Pi = V\Delta\sigma + \frac{1}{2}\Gamma(\Delta F)^2 \quad (7)$$

The computation of both V and Γ relies on specific assumptions on the option valuation model. Given that we work with futures options, we use Black's model for option pricing (the corresponding formulations for V and Γ can be found in Hull [2002]). The vega (gamma) exposure of a single option position is computed as the option vega (gamma) times the number of options in the position. The vega (gamma) exposure of a portfolio of options on the same underlying, which is our case, is simply the sum of the vega (gamma) exposures of the single options.

⁵The remaining higher order terms of the series expansion can be safely neglected.

⁶We have also performed the calculations including the effect of changing time to maturity, but the difference in the results was negligible.

In the vega-gamma model there are two sources of risk whose dynamic needs to be modeled: the underlying and the implied volatility. Only parallel changes in the volatility curve are allowed, with the implied volatilities of all options changing by the same amount $\Delta\sigma$. For our purpose, it is then sufficient to model the evolution of the ATM implied volatility σ_{ATM} . Following a conventional approach, we assume that the ATM implied volatility follows a mean reverting process with Gaussian innovations: $\Delta\sigma_{ATM} = \alpha_{ATM}(\mu_{ATM} - \sigma_{ATM,t-1}) + \sigma_{ATM}^{vol}\xi_t$. For consistency with the other two methods, the evolution of the underlying is modeled by an EGARCH model with normal errors, correlated with the disturbances in the volatility process.

It is worth noticing that both benchmark models can be interpreted as special cases of the more general Kalman filter specification. The sticky-delta model represents an oversimplified version, where all the volatility skew coefficients are constant, and the only factor which evolves is the underlying asset. The vega-gamma model constitutes a special case, where both the ATM level α_0 and the underlying evolve stochastically (and in a correlated fashion). The remaining coefficients of the volatility curve are instead constant. As an alternative to the approach followed in this work (i.e. individual estimation and assessment of the single models), the validity of the different specifications could therefore be assessed by testing the appropriate restrictions in the more general model.

5 The Estimation of the Models

In this section we present the results from the estimation of the parameters of the Kalman filter, the sticky-delta and the vega-gamma models, over the period 1998-2000. We can set the parameters to be constant, or we can let them vary at each time step. In our application, the first assumption would probably be too strong, since we test our model over one year, and the second one would be unnecessarily computationally intensive, given that we do not expect the model parameters to change significantly on a daily basis. Therefore, we choose to re-estimate the parameters of our models every quarter, over a rolling window of three years of data.⁷

For the estimation of the Kalman filter model we first calculate the skew coefficients for each day of the estimation window. Plots of the time series of the four coefficients are reported in Fig. 2, and their autocorrelation functions are displayed in Fig. 3.⁸ For all the coefficients, the autocorrelation function shows an exponential decay which can be conveniently represented

⁷The four estimation windows are: Jan.98 - Dec.00, Mar.98 - Feb.01, Jun.98 - May01, Sep.98 - Aug.01.

⁸The plots refer to the first estimation period, but the patterns stay the same if we consider the other three windows.

by an AR(1)/Ornstein-Uhlenbeck process. Therefore, the assumption made for the system equation in (2) seems appropriate.⁹

The skew coefficients are then employed to estimate $A_t, \underline{\mu}_t$ and V_t from the system equation, by means of a linear simultaneous equation estimation technique. The estimates for the diagonal elements of A_t ,¹⁰ and for $\underline{\mu}_t$, computed on each of the four estimation windows, are reported in Table 1. Table 2 refers to the correlation matrices associated with V_t , which are more intelligible than the corresponding covariance matrices.

As already suggested by the inspection of the autocorrelation functions, the slope coefficient reverts towards its long run mean very slowly. A slightly faster mean reversion characterizes the ATM volatility level, whereas considerably higher (although slightly decreasing through time) mean reversion coefficients are estimated for both curvature and skewness. The analysis of the correlation matrix of the error terms of the process reveals that the highest correlation occurs between curvature and skewness ($\simeq 0.8$). A correlation of around 0.45 relates curvature and slope. Smaller correlations are found between ATM level and curvature, and ATM level and slope. In both cases, however, the correlation coefficients increase (in absolute terms) through time, from a value of, respectively, -0.20 and -0.27 to a level of -0.36 and -0.34 . Very small correlations are found between slope and skewness, and ATM level and skewness.

Considering that the fit of the cubic polynomial to the market implied volatility skew is better the closer we are to the ATM level, and less accurate as far OTM we move, the estimate for the measurement error covariance matrix is based on a grouping in buckets of log moneyness.¹¹ We select ten buckets of log moneyness, in consideration of the ranges of log moneyness in our data set, and for each of them we calculate the standard deviation of the measurement error across the estimation period. The results, displayed in Table 3, suggest that there is a significant variability in the estimated measurement error volatility which, for deep OTM puts and calls, turns out to be, respectively, four and five times its value for ATM options (0.2%).

The estimates of the coefficients of the EGARCH model for the dynamics of the underlying are presented in Table 4 (with standard errors in brackets). As expected, there is a statistically significant leverage effect in our series, captured by the coefficient γ . In order to investigate whether the dynamic of underlying asset and skew coefficients are related, we calculate empirical correlation coefficients between the futures daily log returns and the daily changes in each of the skew coefficients, over the four estimation

⁹The spikes in both time series and autocorrelation function for the slope coefficient are due to the effect of the change in time to maturity when we roll on to the following contract, not completely captured by the normalization of the log moneyness measure.

¹⁰The off-diagonal elements of A_t are very small and statistically insignificant.

¹¹The measurement error matrix D_t is estimated only once, since its quarterly estimates turned out to be not significantly different from each other.

windows. Only the correlations between log returns and 1) ATM volatility level; 2) slope coefficient; turn out to be statistically significant, and basically constant around the values of -0.82 and 0.32 , respectively, for all the estimation periods.¹²

For the purposes of the sticky-delta model, only the estimates of the EGARCH model are needed, since the underlying is the only stochastic variable.

For the vega-gamma model, the estimates of the EGARCH specification are again relevant for modeling the dynamic of the underlying. We then estimate the coefficients of a mean reverting Gaussian process for the ATM implied volatility (see Table 5). To account for the correlation between the two risk factors, we use again the empirical correlation coefficient of -0.82 .

6 The Option Portfolios

In order to assess the goodness of the linear Kalman filter technique at modeling the dynamic of the volatility curve for risk management purposes, we test how well this method predicts the actual daily variations in the marked-to-market value of option portfolios sensitive to changes of different nature in the volatility curve.

We consider the following four option portfolios:

- A short straddle (short one call and one put ATM), which is sensitive to changes in the level of the ATM implied volatility. The portfolio's value decreases (increases) when the volatility level goes up (down).
- A long risk reversal (short one OTM put and long one OTM call), sensitive to changes in the slope of the volatility smile. A loss (profit) occurs when the slope increases (decreases).
- A long butterfly spread (long one call and one put OTM, short one call and one put ATM), sensitive to changes in the curvature of the volatility smile. The portfolio loses (gains) value when the curvature decreases (increases).
- A long "Mexican hat" (long two calls and two puts OTM, short one call and one put ATM), which is vega-neutral.

Each portfolio is made delta-neutral by assuming the appropriate position in the underlying future.¹³

¹²We have also checked for the presence of residual non-linear dependence between log returns and volatility coefficients not captured by the model. No significant pattern could be identified.

¹³The analysis of the corresponding portfolios of opposite sign (long straddle, short risk reversal, short butterfly spread, short "Mexican hat") would be redundant. We remind

For practical implementation, the ATM options are those with strike price closest to the current level of the underlying. The choice of the OTM options for our portfolios is based on two levels of moneyness (defined, for this purpose, as the ratio between the strike and the underlying futures price), one for the OTM puts and one for the OTM calls, equidistant from the ATM level. A trade-off exists between choosing OTM options too close to the ATM, which are highly liquid also for a short time to expiry but not very sensitive to non-parallel changes in the volatility curve, and selecting too far OTM options, which are very sensitive to different sources of changes in volatility, but become too illiquid as time to maturity approaches. In our case, a ratio K/F of 0.92 for OTM puts and of 1.08 for OTM calls seems to represent a satisfactory compromise.

Each day of the testing period (year 2001) we build the four delta-neutral portfolios from our data set, and we calculate their marked-to-market value on both that day and the following day, in order to compute the actual change in value.

7 Density Forecasts of the Changes in Option Portfolios' Value

For each of the three models, parametrized according to the estimates obtained in Section 5, we produce daily forecasts of the changes in value of the four option portfolios, over the year 2001.¹⁴ In order to assess these distributional forecasts not only in relative, but also in absolute terms, a comparison is drawn with the actual daily changes in the portfolios' value across the year. Particular attention is paid to the tails of the density forecasts, which are especially relevant for risk management purposes.

7.1 Derivation of the density forecasts

The density forecasts of the changes in option portfolios' value are derived by means of a simple Monte Carlo simulation exercise.

In the Kalman filter model, we start with initial estimates \hat{x}_1 and S_1 , at the beginning of the testing period. For simplicity, the coefficients of the cubic fitted to the empirical volatility curve on the last day of the estimation period are chosen as \hat{x}_1 . S_1 is set equal to the covariance matrix of the stationary distribution of the multivariate Ornstein-Uhlenbeck process in Eq. (2), $S_1 = C(A + A')^{-1}C'$. Subsequent estimates \hat{x}_t and S_t are obtained through the updating Eqs. (4.4) and (4.5).

the reader that the upper percentiles of the distributions of the changes in value for the original portfolios correspond to the lower percentiles of the distributions for the portfolios with opposite sign.

¹⁴It is worth emphasizing that in our work we only consider out-of-sample density forecasts.

Each day t we draw 5,000 correlated samples from $N(\hat{x}_t, S_t)$ for the skew coefficients, and $N(0, \sigma_t)$ for the log returns on the underlying future, where σ_t is obtained from the EGARCH model. Once we possess forecasts for the underlying level at time t , we can calculate the log moneyness (again, at t) for the strikes of the options included in the four portfolios at time $t - 1$. We then obtain forecasts for the corresponding implied volatility levels, via the cubic polynomial in Eq. (1), with coefficients equal to the coefficient forecasts for time t . The resulting implied volatilities are placed into Black’s formula to derive the forecasts for the marked-to-market value of the options in the portfolios. Since we repeat this procedure 5,000 times, we obtain an entire density forecast of the changes in the marked-to-market portfolios’ value.

In order to produce density forecasts from the sticky-delta model, each day t we draw 5,000 samples from $N(0, \sigma_t)$, the distribution of the log returns on the underlying. For each sample we compute the forecasted value of the underlying at t , and the new levels of log moneyness for all the options included in the four portfolios at time $t - 1$. We then move along the volatility curve observed at $t - 1$. The implied volatilities corresponding to the updated levels of log moneyness are placed into Black’s formula to obtain forecasts of the new marked-to-market option prices.

In the vega-gamma model, each day we draw 5,000 correlated samples from $N(\alpha_{ATM}(\mu_{ATM} - \sigma_{ATM,t-1}), \sigma_{ATM}^{vol})$ for the changes in the ATM volatility level, and $N(0, \sigma_t)$ for the log returns on the underlying future. We then obtain 5,000 pairs $(\Delta\sigma_{ATM}, (\Delta F)^2)$ which, combined with the vega and gamma exposures (V and Γ) computed for the four option portfolios at $t - 1$, yield density forecasts of the changes in portfolios’ value.

7.2 Assessment of the density forecasts

At this stage, we must assess whether the density forecasts produced by the Kalman filter method represent good forecasts of the actual daily changes in the value of the four option portfolios. These density forecasts are also compared to the corresponding forecasts obtained from the sticky-delta and vega-gamma specifications.

First of all, we find that the density forecasts produced by all models turn out to be significantly biased in the mean. If we believe in market efficiency, our daily density forecasts should be always centered around zero, in order to rule out any significant degree of predictability in the market. Instead, we consistently observe non-zero estimates for the mean of our daily return distributions and this leads to unrealistically high values for the annualized Sharpe ratios. In detail, we estimate average Sharpe ratios ranging between 3.5 (for the butterfly spread) and 6 (for the risk reversal), according to the Kalman filter approach. Even higher Sharpe ratios are estimated from the other two approaches (see Table 7). A bias in the mean can be system-

atically observed in the density forecasts of any option portfolio, whatever approach is used to derive them. Therefore, rather than a model-dependent misspecification problem, this seems to be a more generic issue that affects all approaches, although in a different way. Given the generality of the issue, we believe that it is mainly due to the well-known problem of producing unbiased estimates of the drift component of a model over too short time series. Also, it is worth reminding that here we do not impose no-arbitrage conditions on the drift and we work with predefined portfolios, which are not necessarily arbitrage-free. Finally, we produce density forecasts over a very short time horizon (i.e. one day), and this could introduce additional noise in our estimates. Unfortunately, given the nature of the dataset available, no immediate solution to this problem can be easily found in our case. We simply choose to re-center the densities of the portfolio returns on zero, in order to better focus on the remaining features of our forecasts, which should be much less affected than the drift by the sampling variation. After applying this correction, we then move on to appraise the performance of the different models for each portfolio.

If a sequence of one-step-ahead density forecasts p_t of the variable r_t is correct, then the series of Probability Integral Transforms (PIT) $z_t = \int_{-\infty}^{r_t} p_t(y) dy \sim \text{i.i.d. } U(0, 1)$ or, equivalently, the series of the normalized transforms $x_t = \Phi^{-1}(z_t) \sim \text{i.i.d. } N(0, 1)$. In Figs. 4 - 7 we plot the empirical cumulative distribution functions (CDF) of the z_t series obtained from each model, against the theoretical 45° line, for all option portfolios. A simple visual inspection of the plots highlights the superiority of the Kalman filter forecasts for the daily returns for each of the option portfolios. It is worth noticing that most density forecasts (in particular from the sticky delta and the vega-gamma approach) appear to be still affected by some kind of bias, additional to the one in drift discussed previously. This suggests that forecasting daily changes in value of option portfolios is not a straightforward task and this issue should not be underestimated by risk managers.

In order to be more precise about the different nature of such biases, the graphical approach is integrated with a more formal analysis. As suggested by Berkowitz [2001], we use log likelihood ratio (LR) tests on the transformed x_t series to test for the null hypotheses of independence ($LR1$), i.i.d.(0, 1) ($LR2$), zero mean ($LR3$), and unit variance ($LR4$). The results from the implementation of the LR testing techniques are displayed in Table 6. The values for the $LR1$ test indicate that the independence of the PIT series is never a problem. Instead, the joint hypothesis of i.i.d.(0, 1) cannot be rejected only for the density forecasts of changes in value of the straddle and the butterfly spread produced by the Kalman filter model. This evidence highlights the existence of biases in the first two moments of our distributional forecasts. Since the rejection of the null hypothesis in LR tests is only based on the first two moments, we also compute a Jarque-Bera test for normality. The outcomes of the Jarque-Bera test reveal that the normality of

the PIT series is often rejected, with the exceptions of the straddle (Kalman filter) and the mexican hat (Kalman filter and sticky-delta). Therefore, detecting potential misspecifications in higher moments of the distributional forecasts becomes relevant.

For a better comparison between actual and forecasted daily *P&L*, we then compute basic summary statistics for both the time series of the actual daily changes in the values of the portfolios, and the three equally-weighted mixtures (one for each model) of the single density forecasts derived for each day of the testing period.¹⁵ The summary statistics (mean, standard deviation, variance, skewness and excess kurtosis) for the four option portfolios are reported in Table 7.

The analysis of the tails of the density forecasts is also relevant, both as part of a more general and complete assessment of the goodness of our forecasts, and for VaR computations. Following Barone-Adesi, Giannopoulos and Vosper [2002], in order to evaluate whether our density forecasts are appropriate for the calculation of conventional risk measures, we proceed as follows. We compute the 1-day VaR at both 99% and 95% confidence levels as, respectively, the 1st and the 5th percentile of the forecasted distribution of the changes in value of the option portfolios. We then record the number of breaks over the entire testing period, which occur when the actual loss is larger (in absolute value) than the estimated VaR. We also compute the 95th and the 99th percentiles, together with the corresponding number of breaks, which are relevant in terms of VaR calculations for the option portfolios of opposite sign. The number of breaks recorded for both lower and upper percentiles of the density forecasts obtained from the different models, as a percentage of the total number of days in the testing period, are displayed in Table 8, for each of the option portfolios. If the VaR forecasts were correct, we would expect a percentage of breaks of around 1% for both the 1st and the 99th percentiles, and of 5% for both the 5th and the 95th percentiles. The figures indicate that the Kalman filter model produces the best forecasts of both lower and upper tails for all the option portfolios of interest.

Having briefly described the evaluation techniques, we can now focus on the appraisal of the density forecasts of the returns for each of the four option portfolios.

7.2.1 The Short Straddle

The summary statistics in Table 7 indicate that the time series of the actual changes in the value of the short straddle position is highly skewed to the right (skewness of -2.45) and leptokurtic (excess kurtosis of 10.86). The density forecasts obtained from the Kalman filter model seem to provide

¹⁵This method yields more accurate results than the simple averaging of the summary statistics of the individual density forecasts across time.

the best fit (both in the body and in the tails) to the actual portfolio returns, although they show lower variance and less skewness and fat-tailness compared to the time series of the actual returns. The superior performance of these forecasts is evident from both the CDF plot and the number of VaR breaks.

The density forecasts which seem to best reproduce the skewness and the fat-tailness of the actual returns are those generated by the sticky-delta model. Such forecasts, however, present an upward bias in the variance and a severe misspecification in the right tail, as reported in Table 8, and also evident from the CDF plot.

Since the returns on a straddle position only depend on changes in the ATM volatility level, we are not too surprised to observe that the vega-gamma model produces density forecasts whose fit is better than that of sticky-delta forecasts (but still worse than that of Kalman filter forecasts). As for the Kalman filter approach, the forecasts cannot reproduce the amount of asymmetry and leptokurtosis observed in the data. We can also detect an upward bias in the second moment and a slight misspecification of the right tail.

7.2.2 The Long Risk Reversal

The distribution of the realized returns of the long risk reversal portfolio over time turns out to be slightly negatively skewed (skewness of -0.82) and moderately leptokurtic (excess kurtosis of 4.04). According to all our evaluation techniques, the density forecasts obtained from the Kalman filter approach show again the best fit to the actual returns, even though the higher moments are again slightly underestimated.

The sticky delta model instead generates forecasts characterized by severe upward biases in both the variance and the skewness of the distributions. Such biases induce significant misspecifications not only in the body but also in the tails of the density forecasts, as evident from the CDF plot and Table 8.

The density forecasts from the vega-gamma model are even more biased (very low variance and very high skewness and kurtosis compared to the actual returns on the option portfolio), due to the fact that this model cannot account for pure changes in the slope of the volatility curve. The large underestimation of the dispersion measure explains the huge misspecification of both tails, with remarkably high percentages of VaR breaks.

7.2.3 The Long Butterfly Spread

The sample summary statistics of the time series of the actual returns on the long butterfly spread portfolio reveal the presence of a moderate negative skewness (-1.07) and excess kurtosis (2.80). An overall good fit to the

actual portfolio returns is observed for the density forecasts generated from the Kalman filter approach which, however, are affected by an upward bias in the asymmetry and fat-tailness measures.

Much less accurate density forecasts are obtained from either the sticky-delta or the vega-gamma model, which cannot easily account for changes in the curvature of the volatility skew. In this case, significant upward biases affect not only the higher moments of the distributions, but also the variance. The presence of such biases explains the misspecification in the tails, measured by the percentage of VaR breaks.

7.2.4 The Long “Mexican Hat”

The time series of the actual daily returns on the long “Mexican hat” vega-neutral position displays a sample variance of 2.53, and very pronounced higher moments (skewness of 2.55 and excess kurtosis of 22.65). Similar values for the higher moments can be observed only for the mixture of density forecasts generated from the sticky-delta model, which, however, exhibits a significant downward bias in the variance, easily detectable from the CDF plot. This bias in the second moment induces the misspecification in the tails reported in Table 8.

The density forecasts generated from the Kalman filter approach are instead less asymmetric (skewness coefficients around 1.34) and leptokurtic (excess kurtosis around 10.79) but they provide once more the best fit to the actual returns on the option portfolio. The density forecasts from the vega-gamma model also suffer from a severe downward bias in higher moments, which is accompanied by an obvious upward misspecification of the variance (3.42). For both models, the relatively high percentage of VaR breaks in the right tail (around 9% at 95% confidence level) is a consequence of the failure to correctly capture the strong positive asymmetry observed in the actual stream of profits and losses.

Although none of the approaches investigated here seems to be capable of generating correct density forecasts of the daily returns on all our option portfolios,¹⁶ in relative terms we express a strong preference for the discrete Kalman filter as a model for the dynamic of the volatility skew. According to all our evaluation techniques, this method outperforms, both in the body and in the tails of the distribution, the alternative specifications at producing daily forecasts of the changes in the marked-to-market value of some basic option portfolios. This finding holds not only for option portfolios which are sensitive to movements in slope and curvature of the volatility curve and therefore are harder to capture for the other approaches, but also for those sensitive to parallel changes in the curve or vega-neutral. The returns on the latter portfolios, however, are characterized by very pronounced higher

¹⁶See results from the LR2 joint test of i.i.d.(0,1).

moments that the Kalman filter technique seems to underestimate. The sticky-delta model performs better at replicating the highly asymmetric and leptokurtic shape of the return distribution, but this advantage is offset by the presence of severe misspecifications in the second moment and in the tails of the density forecasts.

In the light of our outcomes, we can conclude that, on the whole, the Kalman filter approach constitutes a very promising method for modeling the evolution of the volatility skew, given its ability to easily account for changes in level, slope, curvature and skewness of the curve, especially relevant for risk management purposes. This is consistent with the findings of some recent studies on the dynamic of the volatility curve via component analysis,¹⁷ according to which more than one factor are needed to describe the movements of the curve, and these factors are not perfectly correlated with the underlying asset.¹⁸ In fact, we find that even the 2-factors vega-gamma specification is not sufficient to correctly predict the evolution of the volatility skew.

8 Conclusions and Further Research

Considering the importance of predicting the evolution of the volatility skews/surfaces for pricing, trading and risk management purposes, we have suggested here for the first time a discrete Kalman filter model for the dynamic of the volatility skew.

From a risk management point of view, we have assessed whether this model is capable of generating good density forecasts of the daily returns on a number of option portfolios exposed to variations of different nature in the volatility curve. The results have also been contrasted with those based on density forecasts obtained from two widely used benchmark models, the 1-factor sticky-delta and the 2-factor vega-gamma.

Two main conclusions can be derived from our analysis. First, producing good forecasts of daily changes in the marked-to-market value of option portfolios is much harder than most risk managers believe. None of the methods investigated here yields correct forecasts for all the four option portfolios. A significant bias in the mean, probably due to the excessive sampling variation of the estimates over a short estimation period, affects all the density forecasts. Additional biases in higher moments are also reported for some forecasts. Second, on the whole, the Kalman filter method outperforms the alternative models for all the option portfolios considered here. Our method is easy to implement, of immediate interpretation, and

¹⁷See, for example, Skiadopoulos, Hodges and Clewlow [1999], Alexander [2001], Cont and da Fonseca [2002].

¹⁸In this respect, it could be interesting to re-express our correlated skew coefficients in terms of these orthogonal factors.

very flexible, given that it accounts for various factors/sources of randomness in the option markets.

Further work aimed at improving our Kalman filter specification includes the following: 1) the inclusion of time dependence in the model parameters (long run mean, mean-reversion coefficients, etc.); 2) the inclusion of the specification for the log returns on the underlying into the Kalman filter state space model; 3) additional tests on the effects of the seasonality induced by the decreasing time to maturity on the coefficients, moments estimates, etc. Also, it would be interesting to repeat the analysis for a weekly, instead of daily, forecast horizon, and assess its impact on the bias in the drift component. The generic issue of the bias in the mean also deserves further investigation.

On a different line, it would be interesting to compare the Kalman filter updating of the volatility skew against richer alternatives than the ones explored in this work. For risk management purposes, a meaningful comparison could be drawn, for example, with the filtered historical simulation approach devised by Barone-Adesi, Giannopoulos and Vosper [2002].

References

- [1] C. Alexander. Principles of the skew. *Risk*, 14:29–32, 2001.
- [2] L. Andersen and J. Andreasen. Jump-diffusion processes: volatility smile fitting and numerical methods for pricing. *Review of Derivatives Research*, 4:231–262, 2000.
- [3] G. Bakshi, C. Cao, and Z. Chen. Empirical performance of alternative option pricing models. *Journal of Finance*, 52:2003–2049, 1997.
- [4] G. Barone-Adesi, K. Giannopoulos, and L. Vosper. Backtesting derivative portfolios with Filtered Historical Simulation (FHS). *European Financial Management*, 8:31–58, 2002.
- [5] G. Barone-Adesi and R.E. Whaley. Efficient analytical approximation of American option values. *Journal of Finance*, 42:301–320, 1987.
- [6] D. Bates. Dollar-jump fears, 1984-1992: distributional abnormalities implicit in currency futures options. *Journal of International Money and Finance*, 15:65–93, 1996.
- [7] J. Berkowitz. Testing density forecasts with applications to risk management. *Journal of Business and Economic Statistics*, 19:465–474, 2001.
- [8] D. Brigo and F. Mercurio. A mixed up smile. *Risk*, 13:123–126, 2000.

- [9] D. Brigo and F. Mercurio. Lognormal-mixture dynamics and calibration to market volatility smiles. *International Journal of Theoretical and Applied Finance*, 5:427–446, 2002.
- [10] D. Brigo, F. Mercurio, and G. Sartorelli. Alternative asset-price dynamics and volatility smile. *Quantitative Finance*, 3:173–183, 2003.
- [11] J. Cárdenas, E. Fruchard, E. Koehler, C. Michel, and I. Thomazeau. VAR: one step beyond. *Risk*, 10:72–75, 1997.
- [12] P. Carr, H. Geman, D. Madan, and M. Yor. The fine structure of asset returns: and empirical investigation. *Journal of Business*, 75:305–332, 2002.
- [13] R. Cont and J. da Fonseca. Deformation of implied volatility surfaces: an empirical analysis. In Takayasu, editor, *Empirical Approaches to Financial Fluctuations*. Springer, Tokyo, 2001.
- [14] R. Cont, J. da Fonseca, and V. Durrleman. Stochastic models of implied volatility surfaces. *Economic Notes*, 31:361–377, 2002.
- [15] S. Das and R. Sundaram. Of smiles and smirks: a term-structure perspective. *Journal of Financial and Quantitative Analysis*, 34:211–240, 1999.
- [16] E. Derman. Regimes of volatility. *Risk*, 12:55–59, 1999.
- [17] E. Derman and M. Kamal. The patterns of change in implied index volatilities. Quantitative Strategies Research Notes, Goldman Sachs, 1997.
- [18] E. Derman and I. Kani. Riding on a smile. *Risk*, (2):32–39, 1994a.
- [19] E. Derman and I. Kani. The volatility smile and its implied tree. Quantitative Strategies Research Notes, Goldman Sachs, 1994b.
- [20] E. Derman, I. Kani, and N. Chriss. Implied trinomial trees of the volatility smile. *Journal of Derivatives*, 3:7–22, 1996.
- [21] B. Dumas, J. Fleming, and R. Whaley. Implied volatility functions: empirical tests. *Journal of Finance*, 53:2059–2106, 1998.
- [22] B. Dupire. Pricing with a smile. *Risk*, 7:18–20, 1994.
- [23] B. Dupire. Pricing and hedging with smiles. In M.A.H. Dempster and S.R. Pliska, editors, *Mathematics of Derivatives Securities*, pages 103–111. Cambridge University Press, Cambridge, 1997.

- [24] M. Fengler, W. Härdle, and C. Villa. The dynamics of implied volatilities: a common principal component approach. Discussion Paper 38/01, Humboldt University, Berlin, 2001.
- [25] S. Figlewski. Options arbitrage in imperfect markets. *Journal of Finance*, 44:1289–1311, 1989a.
- [26] S. Figlewski. What does an option pricing model tell us about option prices? *Financial Analysts Journal*, 45:12–15, 1989b.
- [27] P. Glasserman, P. Heidelberger, and P. Shahabuddin. Efficient Monte Carlo methods for Value-at-Risk. In *Mastering Risk*, volume 2. Financial Times-Prentice Hall, 2001.
- [28] A.C. Harvey. *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press, 1989.
- [29] S. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6:327–344, 1993.
- [30] R. Heynen. An empirical investigation of observed smile patterns. *Review of Futures Markets*, 13:317–353, 1993.
- [31] J. Hull. *Options, Futures and Other Derivatives*. Prentice Hall, 2002.
- [32] J. Hull and A. White. The pricing of options on assets with stochastic volatility. *Journal of Finance*, 42:281–300, 1987.
- [33] O. Ledoit and P. Santa-Clara. Relative pricing of options with stochastic volatility. Working paper, UCLA, 1998.
- [34] F. Longstaff. Option pricing and the martingale restriction. *Review of Financial Studies*, 8:1091–1124, 1995.
- [35] A.M. Malz. Vega risk and the smile. *Journal of Risk*, 3, 2001.
- [36] S. Natenberg. *Option Volatility and Pricing: Advanced Trading Strategies and Techniques*. Probus Publishing, Chicago, 1994.
- [37] R. Rebonato and M. Joshi. Assigning future smile surfaces: Conditions for uniqueness and absence of arbitrage. QUARC Working Paper, Royal Bank of Scotland, 2003.
- [38] J.V. Rosenberg. Implied volatility functions: a reprise. *Journal of Derivatives*, 7:51–64, 2000.
- [39] M. Rubinstein. Implied binomial trees. *Journal of Finance*, 49:771–818, 1994.

- [40] P.J. Schönbucher. A market model for stochastic implied volatility. *Philosophical Transactions of the Royal Society, Ser. A*, 357:2071–2092, 1999.
- [41] L. Scott. Option pricing when the variance changes randomly: theory, estimation and an application. *Journal of Financial and Quantitative Analysis*, 22:419–437, 1987.
- [42] G. Skiadopoulos, S.D. Hodges, and L. Clewlow. Dynamics of the S&P500 implied volatility surface. *Review of Derivatives Research*, 3:263–282, 1999.

Table 1: Estimates of A_t and μ_t - Kalman filter model.

	<i>1st estimation period</i>		<i>2nd estimation period</i>	
	Diag. A_t	μ_t	Diag. A_t	μ_t
Level	0.0541	0.2155	0.0573	0.2192
Slope	0.0169	-0.1229	0.0195	-0.1316
Curvature	0.1625	0.1112	0.1578	0.1111
Skewness	0.1782	0.1033	0.1761	0.1005
	<i>3rd estimation period</i>		<i>4th estimation period</i>	
	Diag. A_t	μ_t	Diag. A_t	μ_t
Level	0.0593	0.2218	0.0508	0.2163
Slope	0.0189	-0.1222	0.0189	-0.1200
Curvature	0.1320	0.1142	0.1208	0.1176
Skewness	0.1699	0.0999	0.1672	0.0992

Table 2: Estimates of correlation matrix and standard deviations of system equation errors - Kalman filter model.

<i>1st estimation period</i>					<i>2nd estimation period</i>				
	Level	Slope	Curvature	Skewness	Level	Slope	Curvature	Skewness	
Level	1	-0.2673	-0.1969	-0.0016	1	-0.2698	-0.1994	0.0001	
Slope		1	0.4875	0.1144		1	0.4756	0.0761	
Curvature			1	0.8346			1	0.8290	
Skewness				1				1	
Std. dev.	0.0132	0.0094	0.0303	0.0319	0.0133	0.0094	0.0281	0.0300	

<i>3rd estimation period</i>					<i>4th estimation period</i>				
	Level	Slope	Curvature	Skewness	Level	Slope	Curvature	Skewness	
Level	1	-0.2966	-0.2670	-0.0384	1	-0.3366	-0.3570	-0.1131	
Slope		1	0.4382	-0.0212		1	0.4491	-0.0399	
Curvature			1	0.7849			1	0.7579	
Skewness				1				1	
Std. dev.	0.0133	0.0092	0.0252	0.0271	0.0110	0.0087	0.0239	0.0251	

Table 3: Estimates of measurement error volatility per buckets of log money-
 eyness - Kalman filter model.

Buckets log moneyness	Std. dev. measurement error
≤ -0.50	0.0097
$> -0.50, \leq -0.40$	0.0086
$> -0.40, \leq -0.30$	0.0059
$> -0.30, \leq -0.20$	0.0043
$> -0.20, \leq -0.10$	0.0030
$> -0.10, \leq -0.01$	0.0028
$> -0.01, \leq 0.01$	0.0023
$> 0.01, \leq 0.10$	0.0037
$> 0.10, \leq 0.20$	0.0072
> 0.20	0.0123

Table 4: Estimates of EGARCH(1,1) model with normal errors.

	<i>1st est. period</i>	<i>2nd est. period</i>	<i>3rd est. period</i>	<i>4th est. period</i>
ω	-0.4201 (0.0971)	-0.4727 (0.1057)	-0.4766 (0.1054)	-0.3693 (0.1035)
α	0.1023 (0.0307)	0.1111 (0.0323)	0.0936 (0.0320)	0.0884 (0.0317)
γ	-0.1779 (0.0177)	-0.1771 (0.0197)	-0.1714 (0.0196)	-0.1240 (0.0188)
β	0.9520 (0.0102)	0.9459 (0.0113)	0.9451 (0.0112)	0.9579 (0.0105)

Standard errors in brackets.

Table 5: Estimates of mean reverting model for ATM volatility.

	<i>Estimation period</i>			
	<i>1st</i>	<i>2nd</i>	<i>3rd</i>	<i>4th</i>
μ_{ATM}	0.2112	0.2171	0.2198	0.2146
α_{ATM}	0.0573	0.0602	0.0647	0.0564
σ_{ATM}^{vol}	0.0133	0.0133	0.0135	0.0112

Table 6: Density forecasts tests - $H_0 : x_t \sim \text{i.i.d. } N(0, 1)$.

		<i>Straddle</i>			<i>Risk-reversal</i>		
		K.F.	Sticky- Δ	V- Γ	K.F.	Sticky- Δ	V- Γ
<i>LR1</i> (independence)	(3.84)	0.01	0.59	0.05	1.28	3.27	0.09
<i>LR2</i> (i.i.d.(0, 1))	(7.81)	*10.83	*131.81	*28.04	5.18	*49.29	*554.15
<i>LR3</i> (zero mean)	(3.84)	0.09	2.34	0.95	0.28	0.81	*7.10
<i>LR4</i> (unit variance)	(3.84)	*10.70	*124.08	*19.39	3.39	*42.58	*519.05
Jarque-Bera	(5.99)	1.63	*64.91	*16.19	*21.70	*186.22	*18.00

		<i>Butterfly spread</i>			<i>Mexican hat</i>		
		K.F.	Sticky- Δ	V- Γ	K.F.	Sticky- Δ	V- Γ
<i>LR1</i> (independence)	(3.84)	1.15	0.10	2.11	0.02	0.02	0.44
<i>LR2</i> (i.i.d.(0, 1))	(7.81)	1.50	*113.25	*18.06	*11.85	*201.21	*44.16
<i>LR3</i> (zero mean)	(3.84)	0.11	2.11	3.08	0.67	0.37	*5.80
<i>LR4</i> (unit variance)	(3.84)	0.19	*107.97	*14.83	*10.95	*199.83	*34.61
Jarque-Bera	(5.99)	*7.08	*85.00	*25.82	4.72	4.58	*32.22

Critical values at 5% in brackets.

* rejected at 5% confidence level.

Table 7: Summary statistics for actual profits and losses and density forecasts.

	Mean	Std.dev.	Variance	Skew	Exc. Kurt.	Sharpe R.
<i>Straddle</i>						
Sample	0.221	2.862	8.189	-2.451	10.864	-
Kalman filter	0.000	2.449	5.996	-1.713	7.862	4.603
Sticky- Δ	0.000	3.056	9.340	-2.271	8.616	3.471
Vega-gamma	0.000	3.070	9.423	-1.556	5.004	5.883
<i>Risk-reversal</i>						
Sample	0.072	0.515	0.265	-0.816	4.037	-
Kalman filter	0.000	0.507	0.257	-0.491	2.786	6.308
Sticky- Δ	0.000	0.650	0.422	1.543	4.858	10.148
Vega-gamma	0.000	0.170	0.029	1.907	13.154	7.970
<i>Butterfly spread</i>						
Sample	0.122	1.118	1.251	-1.066	2.801	-
Kalman filter	0.000	1.101	1.211	-1.490	5.644	3.376
Sticky- Δ	0.000	1.243	1.545	-2.419	12.104	4.108
Vega-gamma	0.000	1.434	2.057	-1.764	8.810	4.718
<i>Mexican hat</i>						
Sample	-0.078	1.590	2.528	2.548	22.646	-
Kalman filter	0.000	1.576	2.484	1.340	10.792	4.094
Sticky- Δ	0.000	1.500	2.250	2.877	24.710	7.055
Vega-gamma	0.000	1.849	3.420	0.794	7.347	4.734

Table 8: Percentage of VaR breaks.

	Lower percentiles		Upper percentiles	
	1%	5%	99%	95%
<i>Straddle</i>				
Kalman filter	1.08%	5.25%	1.50%	5.25%
Sticky- Δ	0.42%	4.17%	13.33%	13.75%
Vega-gamma	0.42%	5.00%	1.25%	3.33%
<i>Risk-reversal</i>				
Kalman filter	1.67%	5.42%	0.83%	4.58%
Sticky- Δ	9.58%	9.58%	0.00%	1.25%
Vega-gamma	31.25%	37.08%	14.58%	27.08%
<i>Butterfly spread</i>				
Kalman filter	1.50%	5.25%	1.25%	3.75%
Sticky- Δ	0.83%	4.58%	12.08%	12.08%
Vega-gamma	0.42%	2.92%	0.83%	2.08%
<i>Mexican hat</i>				
Kalman filter	1.25%	5.83%	3.75%	8.33%
Sticky- Δ	14.58%	16.25%	7.92%	13.33%
Vega-gamma	2.50%	5.00%	6.67%	10.00%

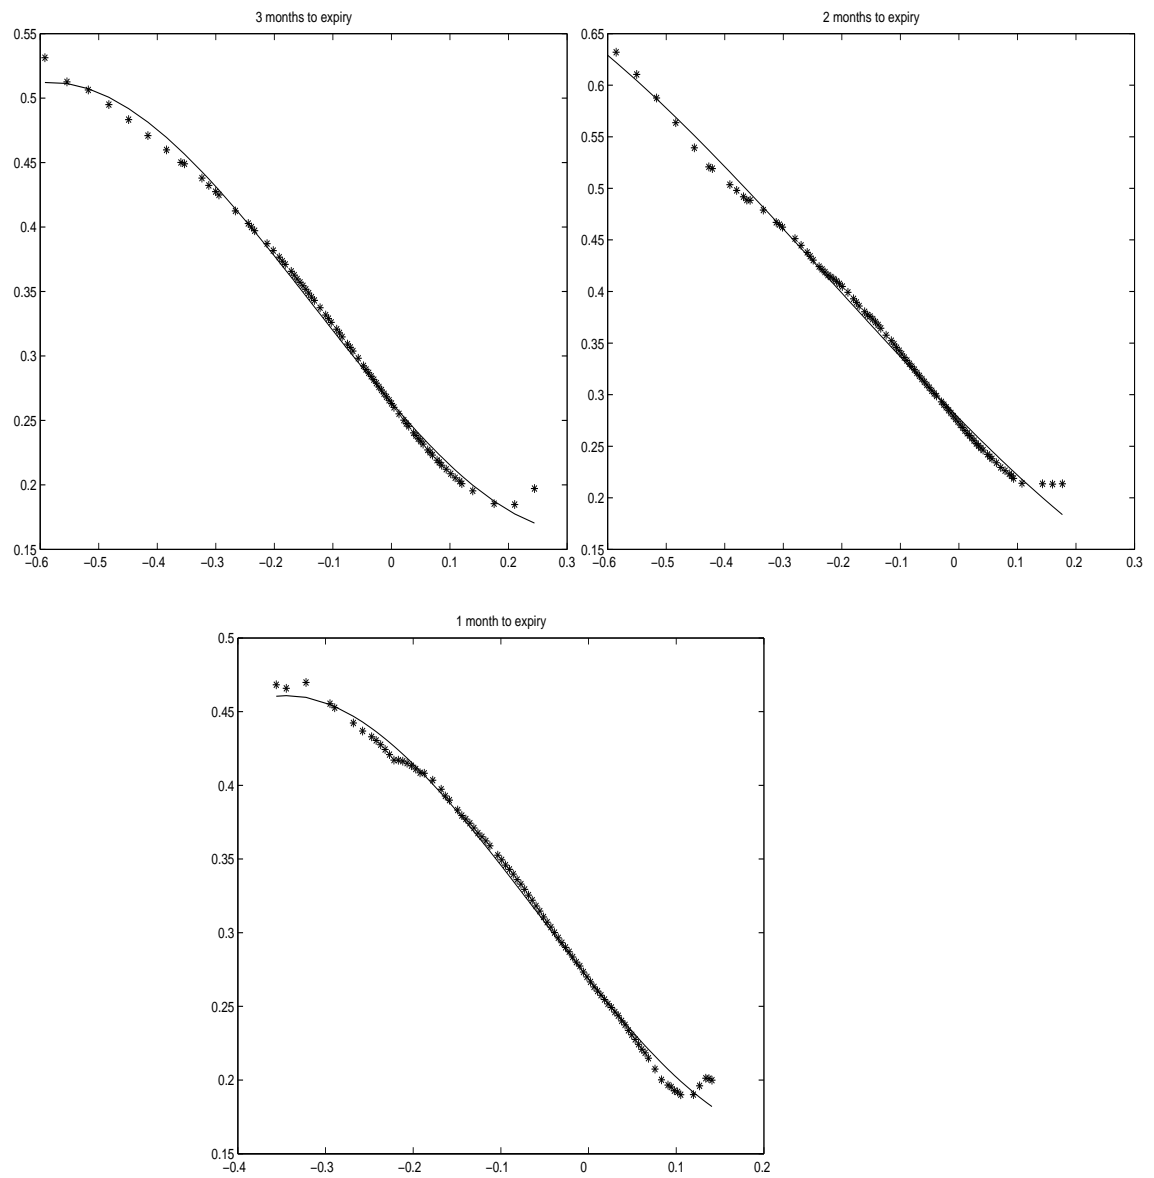


Figure 1: Cubic fit to the implied volatility curve for three expiries of the March 99 contract.

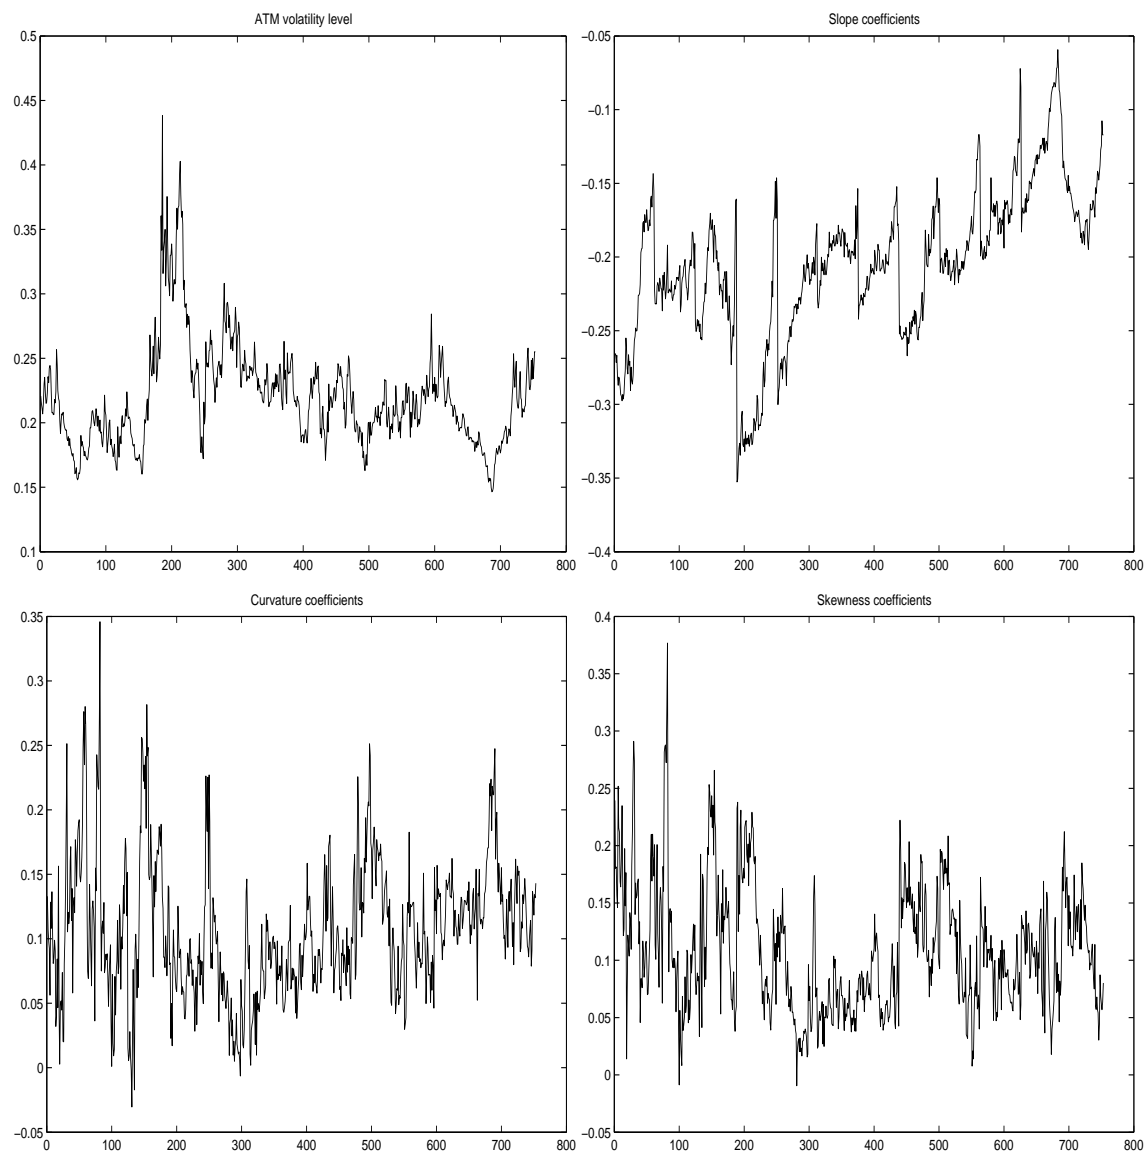


Figure 2: Time series of the volatility skew coefficients (Jan.98 - Dec.00).

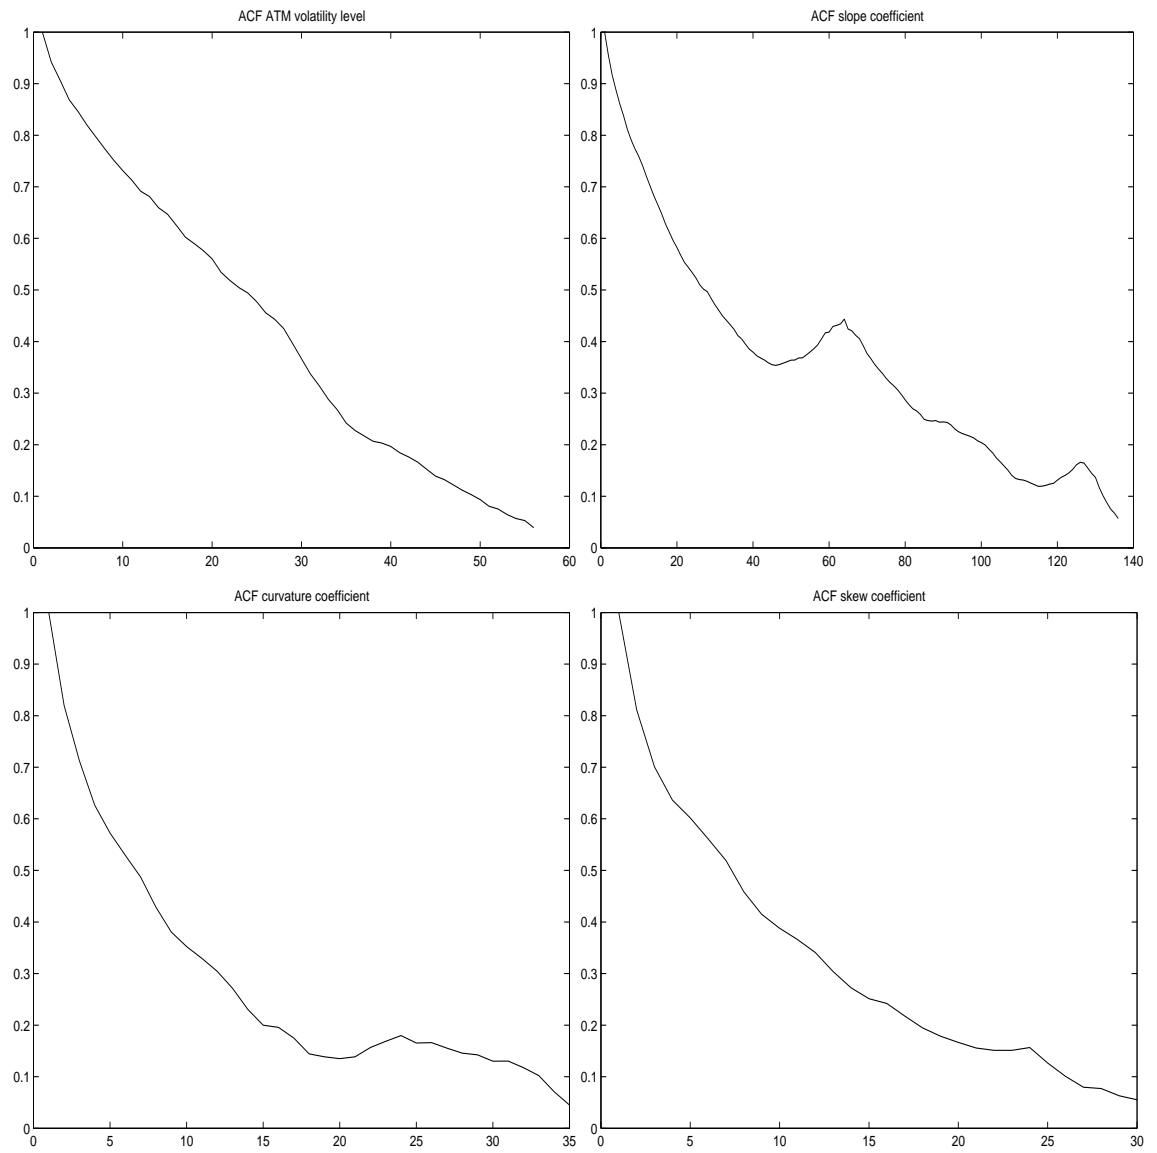


Figure 3: Autocorrelation functions of the volatility skew coefficients (Jan.98 - Dec.00).

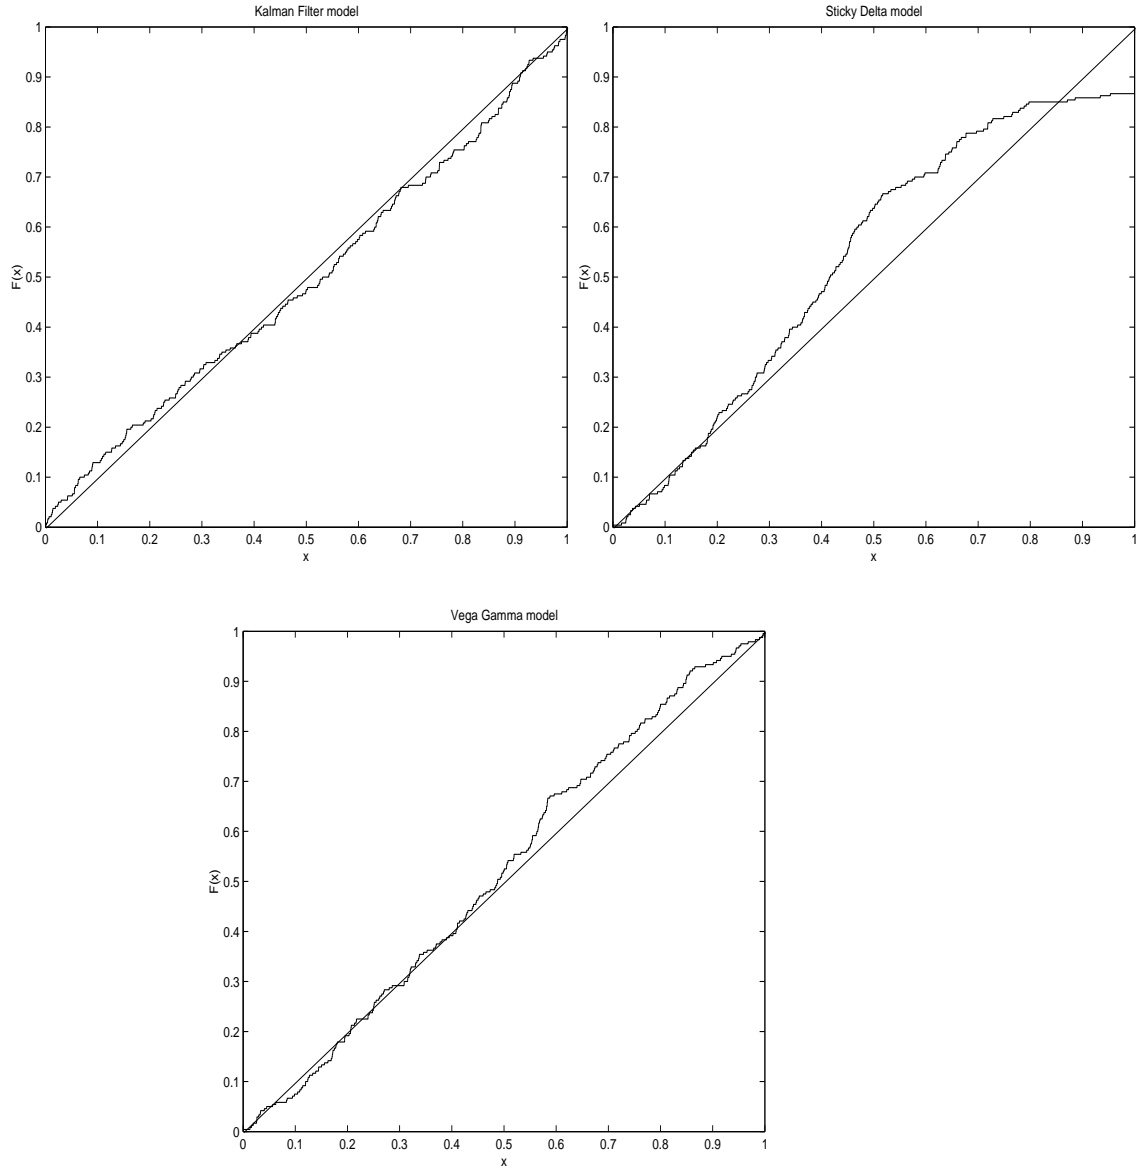


Figure 4: CDF plots of goodness-of-fit - short straddle.

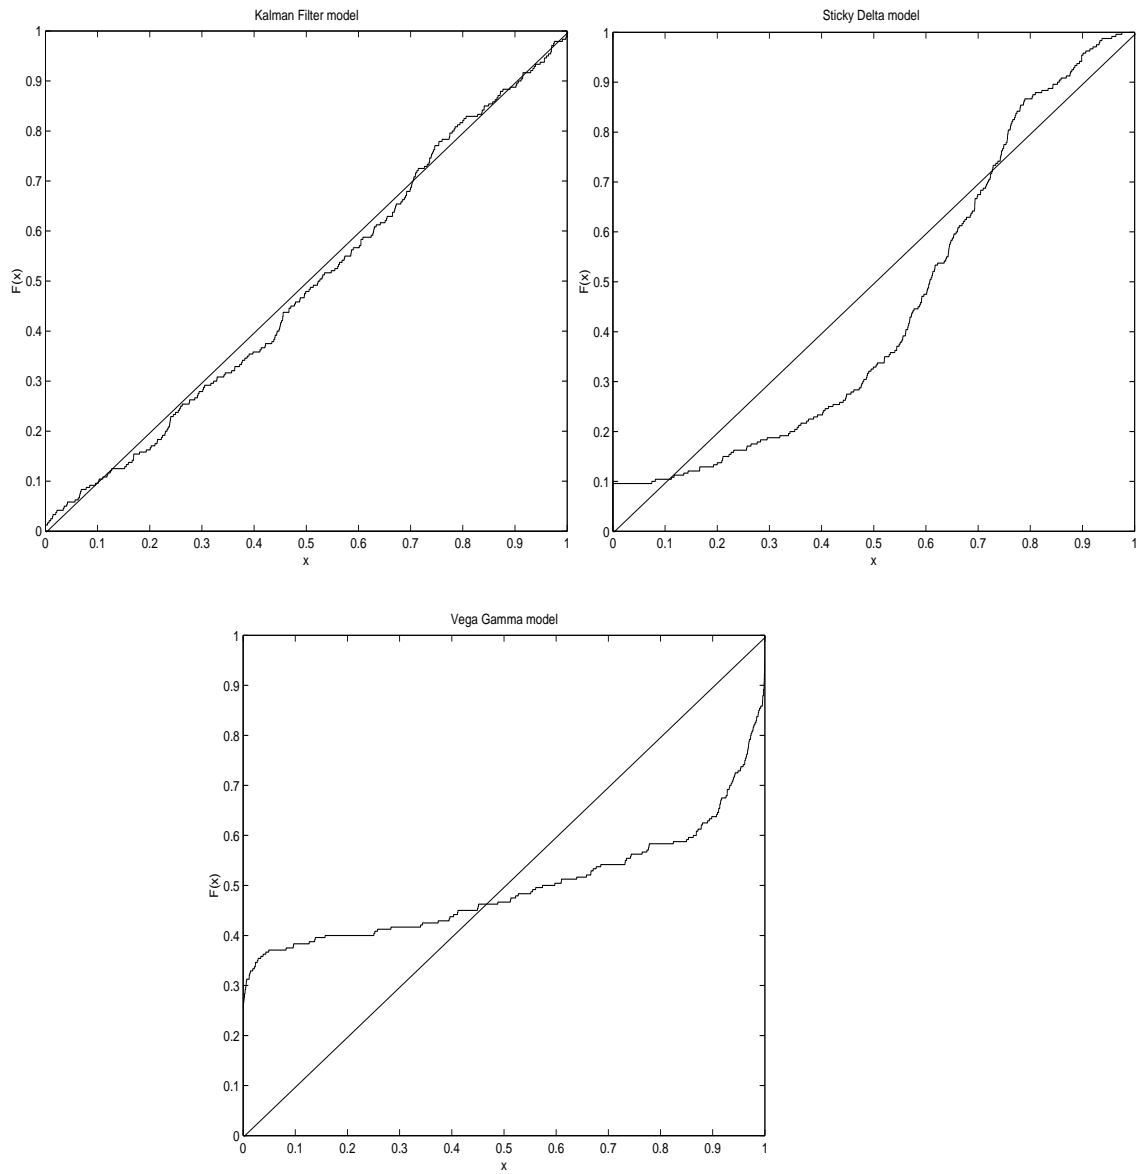


Figure 5: CDF plots of goodness-of-fit - long risk reversal.

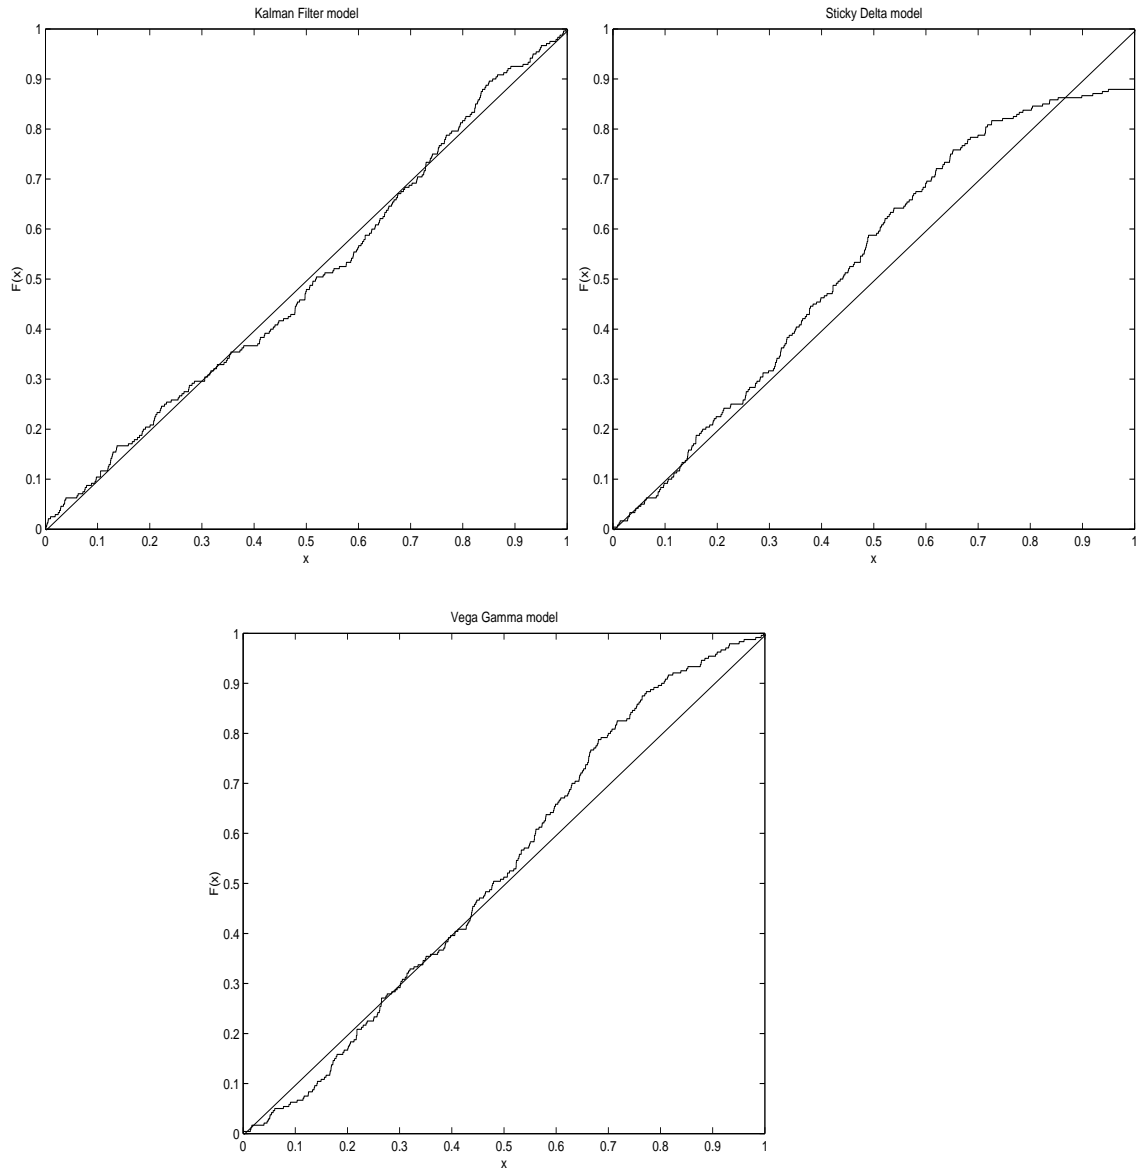


Figure 6: CDF plots of goodness-of-fit - long butterfly spread.

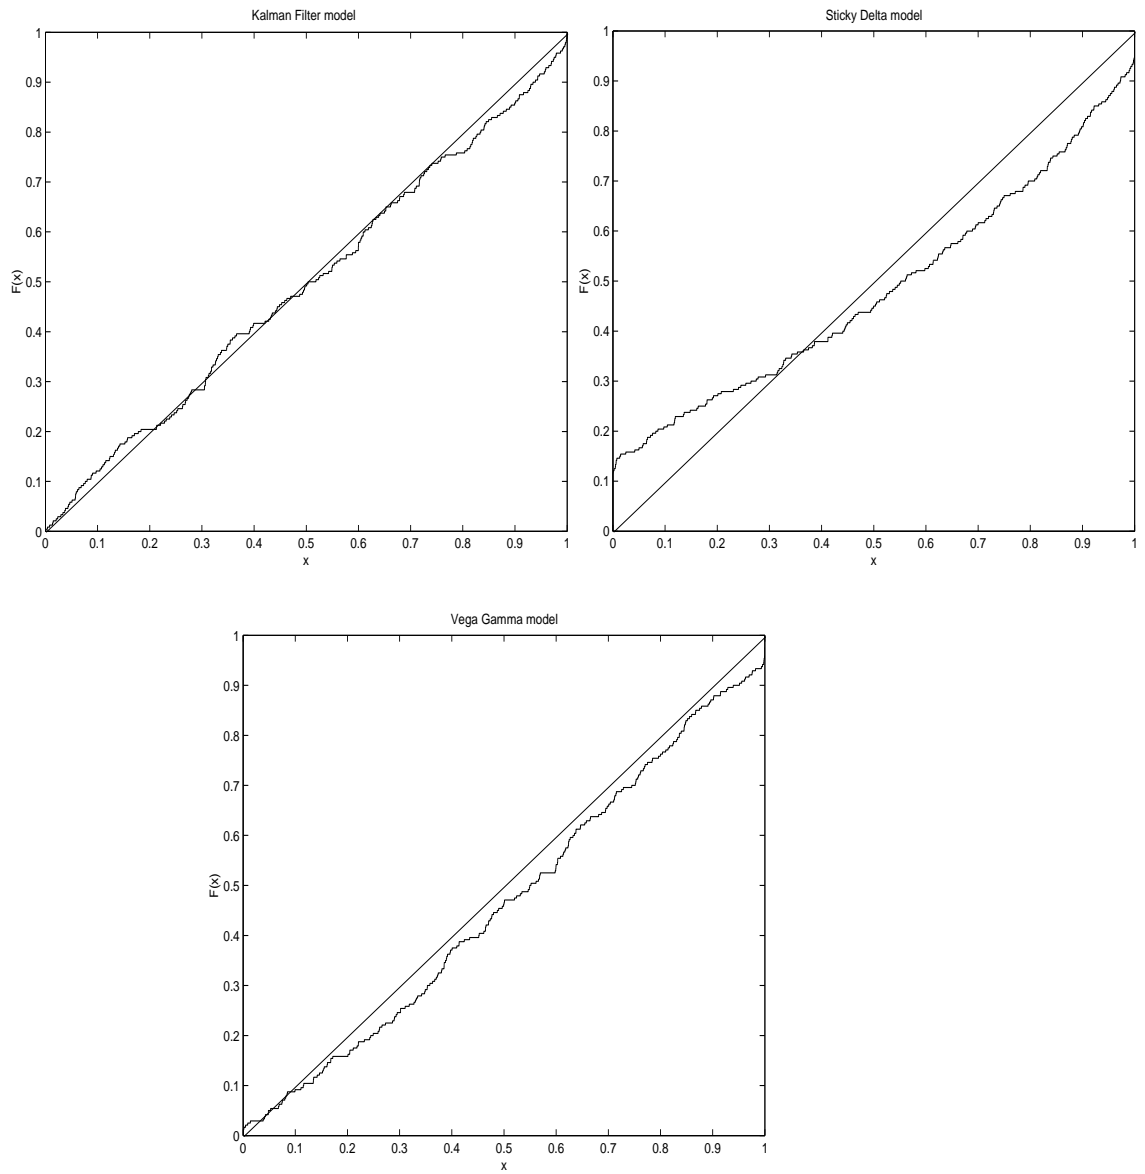


Figure 7: CDF plots of goodness-of-fit - long Mexican hat.