UvA ⚜ ECONOMETRICS

Discussion Paper: 2004/10

Multifactor volatility models: Evidence from stock and option markets

Antoine P.C. van der Ploeg

UvA ⚜ UNIVERSITEIT VAN AMSTERDAM

Multifactor volatility models:
Evidence from stock and option markets

Antoine P.C. van der Ploeg

Department of Quantitative Economics, UvA Econometrics,
University of Amsterdam,
Roetersstraat 11, 1018 WB, Amsterdam, The Netherlands
e-mail: a.p.c.vanderploeg@uva.nl

1 November 2004

Abstract
Accumulating empirical evidence indicates that stock volatilities are driven by more than one latent factor. In this paper we provide additional evidence by combining FTSE100 stock-index data and three at-the-money option series of various maturities in a Kalman filter-based QML estimation strategy of multifactor affine stochastic volatility option pricing models (see Van der Ploeg et al. (2003)). We find that the volatility dynamics can satisfactorily be described by three independent volatility factors. In line with the literature, the first factor is extremely persistent. The second factor is much quicker mean-reverting. Shocks to this factor have a half-life of 2.5 months. The third factor is very fast mean-reverting, with shocks that have a half-life of about 10 days. We interpret the volatility factors in two ways. One interpretation is that each determines the long-term, middle-long-term and short-term trends in the stock volatility evolution respectively. We show how each of the factors influences the prices of options of different maturities. Our second interpretation concerns their impact on the shape and dynamics of the volatility term structure. The long-memory factor appears mainly responsible for changes in the general level of the term structure. Slope changes are mainly associated with the second (but also the third) factor. Changes in the curvature of the volatility term structure are surprisingly closely related to the third factor.

JEL classification: C13, C33, C52, G10, G12, G13.

Keywords: Multifactor affine stochastic volatility models, Derivative pricing, Kalman filter, State space models, Combining stock and option data for estimation.

Volatility is an important concept in finance. It is central in, e.g., asset allocation decisions, risk management systems and asset pricing. Accumulating empirical evidence indicates that stock volatilities are driven by more than one factor [1]. In this paper we provide further evidence by combining stock and option data in a Kalman filter-based QML estimation approach of multifactor affine stochastic volatility (SV) option pricing models (Van der Ploeg et al. (2003)). We analyze the major stock index from the UK, the FSTE100-index, and use both returns and three at-the-money (ATM) option series jointly for estimation.

So far multifactor volatility models have mainly been estimated using stock return data only [2]. However, an additional and rich source of information on volatility is evidently formed by option prices. Indeed, Tauchen (2004) points out that daily returns (even very long time series) are just not sufficiently informative to discriminate between different competing models, which appear to fit US stock-index returns about equally well [3]. The present paper is not on discriminating

---

[1] See, e.g., Engle and Lee (1999), Gallant et al. (1999), Barndorff-Nielsen and Shephard (2001, 2002), Alizadeh et al. (2002), Cont and Fonseca (2002), Pan (2002), Eraker et al. (2003) and Chernov et al. (2003).

[2] E.g., Engle and Lee (1999), Gallant et al. (1999) and Chernov et al. (2003).

[3] These are the (continuous-path) 2-factor logarithmic (i.e. exponential-linear) SV model considered in Chernov et al. (2003), the 1-factor affine SV model with price jumps considered in Andersen et al.

between different classes of models however. We focus on the affine class of SV models characterized by an arbitrary number of volatility factors (but without jumps). Using a large dataset on the underlying and its ATM options of various maturities for estimation, our main aim is to examine how many factors are needed to obtain an adequate description of the dynamics present in the joint data.

The extension towards multiple volatility factors is well motivated. As Meddahi (2002), Chernov et al. (2003) and Tauchen (2004) for example argue, one of the reasons for the poor fit of 1-factor diffusion models is that this factor cannot simultaneously fit both the fat tails of the return distribution and the volatility persistence. Adding jumps or introducing a second volatility factor breaks this link. As volatility is so persistent Eraker et al. (2003) argue that 1-factor models do not permit fast enough changing volatility, which is unrealistic in times of increasing market stress. They advocate for adding jumps to volatility. A second volatility factor may alternatively be added. Evidence from option markets that points towards multiple factors comes from Cont and Fonseca (2002). These authors show that the daily fluctuations in the volatility surface of S&P500 and FTSE100-index options can satisfactorily be described by three underlying, abstract principal component processes.

Among the various multifactor SV models that have recently been proposed are the 2-factor GARCH model of Engle and Lee (1999), the affine and non-affine (e.g. logarithmic) SV models considered by Gallant et al. (1999), Alizadeh et al. (2002), Chernov et al. (2003), and the Lévy process-driven models of Barndorff-Nielsen and Shephard (2001, 2002). Most of these papers focus on the stock price dynamics and, as such, use stock data for estimation. In a more direct option-pricing context Bates (2000) proposes a 2-factor SV model with jumps in returns. Bates uses calibration to a cross-section of option prices for estimation. Duffie et al. (2000) introduce the very general class of affine jump diffusions for option pricing, of which the model considered here is a special case.

Empirical implementations of multifactor models show the following. Engle and Lee (1999) decompose the volatility into a permanent and transitory component, of which the latter one is mean-reverting towards the former. US and Japanese stock market data support this decomposition and reinforce the common finding in the literature of persistent volatility. The empirical results in Gallant et al. (1999) of a 2-factor non-affine volatility model show an improvement in fit of S&P500 returns over a 1-factor model. Alizadeh et al. (2002) use the FX price range in Kalman-filter QML estimation of non-affine SV models. Using daily data on five major US dollar exchange rates they find strong evidence of two volatility factors, one very persistent and the other fast mean reverting. Chernov et al. (2003) find that a 2-factor logarithmic SV stock price model with leverage but without jumps yields a better fit of daily 1953-1999 DJIA stock-index data, than do 1 or 2-factor affine SV models, or SV models with jumps. They find the first factor to be very persistent and the second quickly mean-reverting. The persistent factor does not seem to feature volatility feedback, whereas the other does.

Important additional evidence from option markets against 1-factor SV models comes from studies that combine stock and option data in the estimation of

---

(2002) and the 1-factor affine SV model with both jumps in prices and volatility as considered in Eraker et al. (2003).

option pricing models; see e.g., Chernov and Ghysels (2000), Pan (2002), Jones (2003) and Van der Ploeg et al. (2003). Besides time-series data on the underlying typically one additional short-maturity ATM option series is used for estimation. The estimated 1-factor models tend to overprice longer-dated options out of sample. Pan (2002) for example, mentions that multiple factors may be necessary to improve on the fit of the volatility term structure, a point also emphasized by Bates (1996b, 2000).

Our main motivation is precisely that: Examining how many volatility factors are needed to obtain an adequate description of the joint dynamics of the underlying and the volatility term structure of ATM options. We are additionally interested in possible interpretations of the factors; both regarding their role in the evolution of the stock volatility, and their impact on option prices.

Our main results are summarized as follows. We find that three volatility-driving factors are needed to represent the dynamics present in the data in a satisfactory way. The first factor is highly persistent with shocks that are virtually permanent. The second factor is much quicker mean-reverting. Shocks to this factor have a half-life of about 2.5 months. The third volatility factor is fastest mean-reverting, with shocks that have a half-life of about 10 days and have moreover largest variance. The third factor clearly governs large volatility changes in relatively short time periods. A first interpretation of the volatility factors is that the first, second, and third factor determine the long-term, middle-long-term and short-term trends in the stock volatility evolution respectively. As such, each factor impacts differently on option prices. We show that the first factor influences all options similarly, irrespective of maturity. The impact of the second factor gradually diminishes the longer the life of the option. The third factor virtually only affects short-maturity options. As this factor reverts so rapidly back to its mean, shocks tend to average out over a sufficiently long period of time. This results in very marginal impact on long-maturity options. Our second interpretation concerns the impact of each factor on the shape and dynamics of the volatility term structure. We show that the first volatility factor mainly governs movements in the general level of the term structure. Changes in the slope are largely associated with the second and third factor factors. However, the third volatility factor is very closely related to dynamic changes in the curvature of the volatility term structure.

The remainder of this paper proceeds as follows. Sections 1 and 2 briefly review the multifactor affine SV option pricing model and associated state space estimation method as considered in Van der Ploeg et al. (2003). Section 3 discusses the data. The various shapes and dynamics of the volatility term structure of at-the-money options is explored in detail. Section 4 presents estimation results for a 1-factor model, section 5 for a 2-factor model. Section 6 extends to three factors and discusses the interpretations of the factors. Section 7 summarizes and presents directions for future research. An appendix concludes.

1.   Model

Van der Ploeg et al. (2003) consider a member of the class of affine-jump diffusions (Duffie et al. (2000)). Here we only provide a brief review of their no-arbitrage derivative pricing model. Under the market measure, the ex-dividend stock price S_t evolves as

$$dS_t = \mu_t S_t dt + \sigma_t S_t dW_{S,t} \; . \qquad\qquad (\mathbb{P}) \qquad\qquad (1.1)$$

The instantaneous stock variance σ_t^2 is modeled as a sum of n latent, possibly correlated, mean-reverting factors $\mathbf{x}' = (x_1,..,x_n)$, i.e.

$$\sigma_t^2 = \mathbf{1}'\mathbf{x}_t \; , \qquad d\mathbf{x}_t = \mathbf{K}_d(\boldsymbol{\theta} - \mathbf{x}_t)dt + \boldsymbol{\Sigma}\boldsymbol{\Lambda}_t d\mathbf{W}_{x,t} \; . \qquad (\mathbb{P}) \qquad\qquad (1.2)$$

Uncertainty is resolved by the $(n+1)$-dimensional standard $\mathbb{P}$-Brownian motion $\mathbf{W}_t = (W_{S,t}, \mathbf{W}_{x,t}')'$. Note that the leverage effect is not modeled (see our comments below in section 1.1). The vector $(nx1)\,\boldsymbol{\theta} > \mathbf{0}$ represents the mean of the factors, $(nxn)\,\mathbf{K}_d, \boldsymbol{\Sigma}$ are matrices of constants with $\mathbf{K}_d = \text{diag}[k_1,..,k_n] > \mathbf{0}$ being diagonal, and $\boldsymbol{\Lambda}_t$ is a diagonal matrix given by

$$\boldsymbol{\Lambda}_t = \text{diag}\left[\sqrt{\alpha_1 + \boldsymbol{\beta}_1'\mathbf{x}_t},..,\sqrt{\alpha_n + \boldsymbol{\beta}_n'\mathbf{x}_t}\right], \qquad\qquad (1.3)$$

in which $\boldsymbol{\alpha} = (\alpha_1,..,\alpha_n)'$ and $\boldsymbol{\beta}_i = (\beta_{i1},..,\beta_{in})'$, $i = 1,..,n$ are $(nx1)$ vectors of positive real-valued constants. The matrix $\mathbf{K}_d$ governs the speed of adjustment of the factors towards their mean $\boldsymbol{\theta}$. We assume the dynamics of the factors to be well defined, which requires $\alpha_i + \boldsymbol{\beta}_i'\mathbf{x}_t \geq 0$ for all i and t [4]. The specification allows for so-called *level-dependent volatility-of-volatility* or *volatility feedback*[5]: If the current volatility is high (low), then the volatility of the volatility is currently high (low). Fluctuations in the volatility level thus depend on the state of the volatility.

The market price of stock risk is given by $\gamma_{S,t} = (\mu_t + q_t - R_t)/\sigma_t$. Here q_t and R_t denote the deterministically time-varying dividend yield and short interest rate respectively. Following the interest rate literature (see e.g. Duffie and Kan (1996), de Jong (2000) and Dai and Singleton (2000)), the market price of risk of factor x_i -denoted by γ_{it}- is modeled as being proportional to its instantaneous standard deviation,

$$\gamma_{it} = \gamma_i\sqrt{\alpha_i + \boldsymbol{\beta}_i'\mathbf{x}_t} \; , \qquad\qquad (1.4)$$

in which $\gamma_i \in \mathbb{R}$. The collective market price of volatility risk may thus be represented by the vector $\boldsymbol{\gamma}_{x,t} = (\gamma_{1t},..,\gamma_{nt})' = \boldsymbol{\Lambda}_t\boldsymbol{\gamma}$ where $\boldsymbol{\gamma} = (\gamma_1,..,\gamma_n)'$. Compensation for volatility risk is thus a fixed multiple of the volatility function of the risk source. It goes to zero if volatility risk goes to zero -as it should by no-arbitrage requirements.

From Girsanov's theorem, changing from the market measure $\mathbb{P}$ to the risk-neutral pricing measure $\mathbb{Q}$ is then governed by the transformation $d\tilde{\mathbf{W}}_t = d\mathbf{W}_t + \boldsymbol{\gamma}_t dt$ where $\boldsymbol{\gamma}_t = (\gamma_{S,t}, \boldsymbol{\gamma}_{x,t}')'$ is the vector of prices of risk. Here, $\{\tilde{\mathbf{W}}_t; t \geq 0\}$ with $\tilde{\mathbf{W}}_t = (\tilde{W}_{S,t}, \tilde{\mathbf{W}}_{x,t}')'$ is an $(n+1)$-dimensional standard Brownian motion under $\mathbb{Q}$. Under $\mathbb{Q}$, the stock price follows $dS_t = (r_t - q_t)S_t dt + \sigma_t S_t d\tilde{W}_{S,t}$. One advantage of the assumed form for the market price of volatility risk is the

---

[4] Parameter restrictions that ensure a unique strong solution to the SDE (2.5) are in Duffie and Kan (1996) and Dai and Singleton (2000).
[5] Recent time series evidence that supports this assumption is in e.g. Pan (2002), Jones (2003), Chernov et al. (2003) and Van der Ploeg et al. (2003).

fact that it delivers the same type of mean-reverting SDE for the factors under $\mathbb{Q}$ as under $\mathbb{P}$ [6]. Specifically, the risk-neutral dynamics of the factors are given by

$$d\mathbf{x}_t = \tilde{\mathbf{K}}(\tilde{\boldsymbol{\theta}} - \mathbf{x}_t)dt + \boldsymbol{\Sigma}\boldsymbol{\Lambda}_t d\tilde{\mathbf{W}}_{x,t} \,, \qquad\qquad (\mathbb{Q}) \qquad\qquad (1.5)$$

where $\tilde{\mathbf{K}} \equiv \mathbf{K}_d + \boldsymbol{\Sigma}\boldsymbol{\Gamma}\mathcal{B}'$, $\tilde{\boldsymbol{\theta}} \equiv \tilde{\mathbf{K}}^{-1}(\mathbf{K}_d\boldsymbol{\theta} - \boldsymbol{\Sigma}\boldsymbol{\Gamma}\boldsymbol{a})$, $\boldsymbol{a} \equiv (\alpha_1,..,\alpha_n)'$, $\mathcal{B} \equiv \left[\boldsymbol{\beta}_1,..,\boldsymbol{\beta}_n\right]$ and $\boldsymbol{\Gamma} \equiv \text{diag}[\gamma_1,..,\gamma_n]$ and where we assume existence of the inverse of $\tilde{\mathbf{K}}$.

Option price

The model-implied time-t value of a European call option C written on the stock S having strike price K and maturity $T > t$ is given by

$$C_t = \mathbb{E}_{\mathbb{Q}}\left[BS(F_t,K,\tau_t,\bar{R}_t,\bar{\sigma}_t^2) \mid \mathcal{F}_t\right], \qquad\qquad (1.6)$$

where $F_t = S_t \exp[(\bar{R}_t - \bar{q}_t)\tau_t]$ is the time-t forward price written in the forward contract that has the same maturity $\tau_t = T - t$ as the option. $\bar{R}_t$ and $\bar{q}_t$ represent the average (deterministic) interest rate and dividend yield respectively, whereas $\bar{\sigma}_t^2$ is the average (random) stock variance, all over the remaining life of the option. $BS(.)$ stands for the Black-Scholes pricing function expressed in terms of the forward price, i.e.,

$$BS(F_t,K,\tau_t,\bar{R}_t,\bar{\sigma}_t^2) = \exp(-\bar{r}_t\tau_t)\left[F_t\,\Phi(d_{1t}) - K\,\Phi(d_{2t})\right], \qquad\qquad (1.7)$$

with

$$d_{1t} = \frac{\ln\dfrac{F_t}{K} + \frac{1}{2}\bar{\sigma}_t^2\tau_t}{\bar{\sigma}_t\sqrt{\tau_t}}, \qquad d_{2t} = d_{1t} - \bar{\sigma}_t\sqrt{\tau_t} \,. \qquad\qquad (1.8)$$

(Notice that if the forward price is known in practice, an estimate of $\bar{q}_t$ -which may be hard to obtain- is no longer needed.) Although we allow for a multifactor SV option pricing model, a similar pricing formula as in Hull and White (1987) results. This is due to the negligence of the leverage effect, and because the stock price follows a geometric Brownian motion-type SDE under $\mathbb{Q}$.

1.1 Comments on simplifications

In order to dramatically facilitate estimation using four joint time series simultaneously, we do some sacrifices in the present paper.

Multifactor OU SV

Notice that if the factors follow a multifactor Gaussian Ornstein-Uhlenbeck (OU) process (i.e., if $\boldsymbol{a} = \mathbf{1}$ and $\boldsymbol{\beta}_i = \mathbf{0}$ $\forall i$) the stock variance is not guaranteed to stay positive at all times [7]. Although this is theoretically inconsistent, the OU assumption yields large analytical tractability. In our estimations this tractability

---

[6] Chernov and Ghysels (2000), Pan (2002) and Jones (2003) make similar assumptions.

[7] In the interest rate literature the Gaussian OU assumption has often been used as a model for the (non-negative) short rate, see e.g. Vasicek (1977), Langetieg (1980), the continuous-time limit of the Ho-Lee (1986) and de Jong (2000). Bakshi and Kapadia (2003) also assume an OU process for the stock volatility in their proposition 2. Barndorff-Nielsen and Shephard (2001) and Nicolato and Venardos (2002) discuss modeling the stock variance as a sum of independent non-Gaussian positive OU processes driven by positive-increment Lévy processes. Cont and Fonseca (2002) show that the autocorrelation structure of the three principal component processes extracted from daily volatility surface fluctuations is well approximated by OU processes.

is more than welcome. We will therefore exploit it in the present paper. The OU assumption should be considered as a first approximation. Notice that in that case there is no volatility feedback. Specifically, in case of multifactor OU SV the factors follow $d\mathbf{x}_t = \mathbf{K}_d(\boldsymbol{\theta} - \mathbf{x}_t)dt + \boldsymbol{\Sigma}d\mathbf{W}_{x,t}$ under $\mathbb{P}$ and $d\mathbf{x}_t = \tilde{\mathbf{K}}(\tilde{\boldsymbol{\theta}} - \mathbf{x}_t)dt + \boldsymbol{\Sigma}d\tilde{\mathbf{W}}_{x,t}$ under $\mathbb{Q}$. The risk-neutral parameters are then given by $\tilde{\mathbf{K}} = \mathbf{K}_d$ and $\tilde{\boldsymbol{\theta}} = \boldsymbol{\theta} - \mathbf{K}_d^{-1}\boldsymbol{\Sigma}\boldsymbol{\Gamma}\mathbf{1}$. The market price of volatility risk reduces to the constant vector $\boldsymbol{\gamma}$. The factors and instantaneous stock variance are Gaussian [8],

$$\mathbf{x}_t \sim \mathcal{N}\left[\boldsymbol{\theta};\ \mathbf{J}\odot\boldsymbol{\Sigma}\boldsymbol{\Sigma}'\right], \qquad \sigma_t^2 \sim \mathcal{N}\left[\mathbf{1}'\boldsymbol{\theta};\ \mathbf{1}'(\mathbf{J}\odot\boldsymbol{\Sigma}\boldsymbol{\Sigma}')\mathbf{1}\right], \qquad (1.9)$$

where $\odot$ is the Hadamard product, i.e., element-by-element multiplication. The correlation matrix of the factors is given by

$$\text{corr}_{\mathbb{P}}\left[\mathbf{x}_t\right] = \left[\mathbf{I}_n \odot \mathbf{J}\odot\boldsymbol{\Sigma}\boldsymbol{\Sigma}'\right]^{-1/2}\left[\mathbf{J}\odot\boldsymbol{\Sigma}\boldsymbol{\Sigma}'\right]\left[\mathbf{I}_n \odot \mathbf{J}\odot\boldsymbol{\Sigma}\boldsymbol{\Sigma}'\right]^{-1/2}, \qquad (1.10)$$

in which $(nxn)\,\mathbf{J}$ with $[\mathbf{J}]_{ij} = 1/(k_i + k_j)$. Notice further that if $\boldsymbol{\Sigma}$ is diagonal the volatility is driven by independent factors.

No leverage
As a second approximation we assume the Brownian motions W_S and $\mathbf{W}_x$ to be independent. Stock price movements thus occur independently from volatility movements in the present model. Modeling the leverage effect (Black (1976)) complicates matters dramatically, mainly with regard to our estimation method [9]. Jones (2003) also assumes a number of approximations in his estimation strategy of combining stock and option data. To validate one of these approximations, Jones uses a Hull-White (1987)-type argument based on independence between the stock and volatility processes, though the models and resulting option prices he considers assume dependence. Given their specific option pricing model, the (mathematical) foundations and justifications for the estimation strategy in Van der Ploeg et al. (2003) are more solid. Jones (2003) does however argue that the leverage effect is an additional reason to invalidate his approximations, but mentions that the practical importance of the leverage effect is unclear for ATM options. Since we only consider such options in our analysis, ignoring the leverage effect may possibly lead to only mild biases, although further research should clarify this. We are mainly focused on the dynamics of volatility.

2. Estimation method

Van der Ploeg et al. (2003) develop a Kalman filter QML-based approach [10] to the estimation of the multifactor model. The benefits of Kalman filter QML are its relative simplicity and consistency of estimates; a drawback is inefficiency of the estimates. The advantages of their approach include the following. First, the state space framework is naturally suited for incorporation of the hidden volatility factors. Moreover, the joint simultaneous inclusion of time series on the

---

[8] E.g., Van der Ploeg (2003) examines the implied statistical properties of the general multifactor SV model in detail.

[9] Nevertheless, we aim at covering it in future research. Notice that although we are concerned with equity markets here, the model and estimation method seem particularly useful in foreign-exchange modeling contexts, in which the assumption of no leverage seems more justified.

[10] Kalman filter QML-based approaches for estimating SV models have been considered by e.g. Harvey et al. (1994) and recently by Alizadeh et al. (2002). Other Kalman filter-based approaches are considered in e.g. Kim et al. (1998) and Sandmann and Koopman (1998). Good textbook treatments on state space models include Durbin and Koopman (2001), Hamilton (1994) and Harvey (1989).

underlying and several options is easily dealt with due to the "panel-data" character of state space models. Third, the method circumvents simulation of option prices during estimation. Convergence is fairly rapidly achieved. Furthermore, the method naturally allows for measurement error in all series, which may be caused by e.g. market microstructure effects or differences in liquidity. The smoothed evolution of the factors and volatility is a direct by-product of the estimation output. (E.g., in contrast, the efficient method of moments (EMM, Gallant and Tauchen (1996)) as considered in e.g. Chernov and Ghysels (2000) requires reprojection besides estimation, to filter the volatilities.) Finally, in the literature on the term structure of interest rates Kalman filter QML estimation methods have proven to be rather robust (Lund (1997), de Jong (2000) and Duffee and Stanton (2001)).

Extracting information from stock prices
The empirical data discussed in section 3 consists of four time series: a stock price series, a short-maturity (SM), a medium-maturity (MM) and a long-maturity (LM) ATM call option series; see section 3. The timing of the data points is denoted by $t = \Delta t, 2\Delta t, .., T\Delta t$ with $\Delta t = 1/260$; i.e. we assume a year to consist of 260 trading days.

Based on an Euler discretization [11] of the stock price SDE (2.1), Van der Ploeg et al. extract information from squared stock returns via

$$\frac{1}{\Delta t}(r_t - \hat{\mu}_{t-\Delta t}\Delta t)^2 = \mathbf{1'\theta} + \mathbf{1'x}_{t-\Delta t}^* + \omega_t; \qquad \omega_t \sim (0, \sigma_\omega^2) \quad (1.11)$$

with $r_t = (S_t - S_{t-\Delta t})/S_{t-\Delta t}$ the daily stock return, and $\mathbf{x}_t^* = \mathbf{x}_t - \mathbf{\theta}$ denoting the factors in deviation from their mean.

Extracting information from option prices
Extracting information from the ATM option series is via Black-Scholes implied variances:

$$\sigma_{implied,it}^2 = \frac{1}{\tau_{it}}\left[A(\tau_{it}) + \mathbf{B}(\tau_{it})'\mathbf{\theta}\right] + \frac{\mathbf{B}(\tau_{it})'}{\tau_{it}}\mathbf{x}_t^* + \varepsilon_{it}; \qquad \varepsilon_{it} \sim (0, \sigma_{\varepsilon i}^2) \quad (1.12)$$

for $t = \Delta t, .., T\Delta t$ and $i = SM, MM, LM$. We allow for possible non-zero contemporaneous correlation between the error series $\{\varepsilon_{SM,t}\}$, $\{\varepsilon_{MM,t}\}$ and $\{\varepsilon_{LM,t}\}$.

The functions $(1x1)$ $A(.)$ and $(nx1)$ $\mathbf{B}(.)$ are deterministic functions of the time to maturity of the option. They satisfy the following system of Ricatti ordinary differential equations

$$\frac{dA(\tau)}{d\tau} = \tilde{\mathbf{\theta}}'\tilde{\mathbf{K}}'\mathbf{B}(\tau) + \tfrac{1}{2}\sum_{j=1}^{n}\left[\mathbf{\Sigma}'\mathbf{B}(\tau)\right]_j^2 \alpha_j \qquad (1.13)$$

$$\frac{d\mathbf{B}(\tau)}{d\tau} = -\tilde{\mathbf{K}}'\mathbf{B}(\tau) + \tfrac{1}{2}\sum_{j=1}^{n}\left[\mathbf{\Sigma}'\mathbf{B}(\tau)\right]_j^2 \mathbf{\beta}_j + \mathbf{1},$$

---

[11] Jones (2003) also uses an Euler approximation for information extraction, though in a Bayesian framework.

with boundary conditions $A(0) = 0$ and $\mathbf{B}(0) = \mathbf{0}$. Notice that $A(.)$ and $\mathbf{B}(.)$ depend on the risk-neutral parameters $\tilde{\mathbf{K}}$ and $\tilde{\boldsymbol{\theta}}$. The multifactor OU SV assumption yields closed-form expressions for these functions [12]. Rewritten in terms of the $\mathbb{P}$-parameters, they become

$$A(\tau) = \mathbf{1}'\left[\tau\,\mathbf{I}_n - \mathbf{D}(\tau)\right]\left(\boldsymbol{\theta} - \mathbf{K}_d^{-1}\boldsymbol{\Sigma}\boldsymbol{\Gamma}\mathbf{1}\right) + \tfrac{1}{2}\mathbf{1}'\left[\mathbf{N}(\tau)\odot\boldsymbol{\Sigma}\boldsymbol{\Sigma}'\right]\mathbf{1} \qquad (1.14)$$

$$\mathbf{B}(\tau) = \mathbf{D}(\tau)\mathbf{1}\,,$$

where $\mathbf{D}(\tau) = \text{diag}[d_1(\tau),..,d_n(\tau)]$ with $d_j(\tau) = [1 - \exp(-k_j\tau)]/k_j;\ j = 1,..,n$, and matrix $(nxn)\mathbf{N}(\tau)$ has $ij-$th element equal to

$$N_{ij}(\tau) = \frac{1}{k_i k_j}\left[\tau - d_i(\tau) - d_j(\tau) + \frac{1 - \exp[-(k_i + k_j)\tau]}{k_i + k_j}\right]. \qquad (1.15)$$

Using these results, the options measurement equation can be rewritten as

$$\sigma^2_{implied,it} = \mathbf{1}'\boldsymbol{\theta} + \mathbf{1}'\left(\frac{\mathbf{D}(\tau_{it})}{\tau_{it}} - \mathbf{I}_n\right)\mathbf{K}_d^{-1}\boldsymbol{\Sigma}\boldsymbol{\Gamma}\mathbf{1} + \frac{1}{2\tau_{it}}\mathbf{1}'\left[\mathbf{N}(\tau_{it})\odot\boldsymbol{\Sigma}\boldsymbol{\Sigma}'\right]\mathbf{1} \qquad (1.16)$$

$$+ \mathbf{1}'\frac{\mathbf{D}(\tau_{it})}{\tau_{it}}\mathbf{x}_t^* + \varepsilon_{it}\,; \qquad\qquad \varepsilon_{it} \sim (0,\sigma^2_{\varepsilon i})\,.$$

Without going too much into detail, the method for extracting information from call prices is based on the result that $\mathbb{E}_{\mathbb{Q}}[\exp(\tau\,\bar{\sigma}_t^2)\,|\,\mathcal{F}_t] = \exp[A(\tau) + \mathbf{B}(\tau)'\mathbf{x}_t]$ (see Duffie and Kan (1996) and Duffie et al. (2000)). Reconsider the call price formula (1.6). Van der Ploeg et al. first linearize the BS pricing function around the exponent of the integrated variance over the options life. They next take the $\mathbb{Q}$-expectation of this linearization such that they end up with a partly analytical expression for the call price. Some further manipulations then result in an equation that is linear in the latent factors. The point of linearization is taken equal to $\exp(\tau\,\sigma^2_{implied,t})$ where $\sigma^2_{implied}$ is the model-implied BS implied variance. They finally neglect the higher-order terms, introduce noise in the form of an additive error term, and then arrive at equation (1.12).

The introduction of noise is well motivated. First, it serves as "compensation" for the negligence of higher-order terms. From an empirical-implementation point of view more important however, is that it recognizes that a model is never a complete description of reality. (See also Renault (1997) for a discussion on option pricing errors.) In contrast to Pan (2002) but in accordance with Jones (2003), we allow each option price on each day to be measured with error.

Evolution of the factors in discrete time
The exact discrete-time evolution of the hidden OU factors in-deviation-from-their-mean is

$$\mathbf{x}_{t+\Delta t}^* = \exp[-\mathbf{K}_d\Delta t]\mathbf{x}_t^* + \mathbf{u}_{t+\Delta t},\quad \mathbf{u}_{t+\Delta t} \sim \mathcal{N}[\mathbf{0};\mathbf{G}(\Delta t)\odot\boldsymbol{\Sigma}\boldsymbol{\Sigma}'] \qquad (1.17)$$

where $\exp[-\mathbf{K}_d\Delta t] = \text{diag}(\exp[-k_1\Delta t],..,\exp[-k_n\Delta t])$ and $(nxn)\mathbf{G}(\Delta t)$ with $[\mathbf{G}(\Delta t)]_{ij} = (1 - \exp[-(k_i + k_j)\Delta t])/(k_i + k_j)$.

The system matrices of the linear state space model are given in the appendix to this chapter. The model is estimated by Kalman filter QML.

Parameter identification
Not all parameters of the n-factor OU SV model can be identified by our state space estimation method. Van der Ploeg (2004b) considers the problem in detail. Here we summarize the results. The parameters $\theta_1,..,\theta_n$ cannot separately be identified; only their sum. We therefore choose to leave θ_1 unrestricted and restrict $\theta_2 = .. = \theta_n = 0$ prior to estimation. The parameters σ_{ij}; $i, j = 1,..,n$ which appear in matrix $\boldsymbol{\Sigma}$ cannot all be identified either. This directly affects the identification of the remaining parameters. However, if $\boldsymbol{\Sigma}$ is diagonal (i.e. independent factors), or if $\boldsymbol{\Sigma}$ is either lower or upper diagonal (which means correlated factors), then all parameters can be identified. We consider both cases in the estimations below.

2.1 Summary of Monte-Carlo results

Van der Ploeg (2004a) presents Monte-Carlo simulation results for the 1-factor OU SV special case of the model. He considers information extraction from realized volatilities (RV) as well, though the empirical results presented below do not take RV (yet) into account. Van der Ploeg estimates the state space model using five different types of data: only squared return data, only RV data, only (short-maturity near-the-money) option data, the combination squared return-option data, and the combination RV-option data. The realized volatilities were computed using 48 intraday returns (i.e. sampled every 10 minutes during a trading day.)

An examination of the performance of the state space method in its ability to recover the true parameters and volatility paths underlying the simulated data reveals the following. The use of only squared return data for estimation performs worst: It leads to the biggest bias and MSEs of the estimates, especially with regard to the mean-reversion and volatility-of-volatility parameters. Moreover, the smoothed volatility series deviates most from the true underlying series. The results confirm that squared returns are generally considered as noisy estimators of the stock variance. A substantial improvement in bias, MSE and volatility evaluation criteria is obtained when using RV data instead. In turn, the use of option data alone outperforms RV data in general on all criteria: both the bias and MSE of the estimates are small, and the volatility evaluation criteria indicate that the smoothed and true volatility series are close. The results further indicate that combining squared return and option data performs in turn better than only option data. Not surprisingly, the use of both RV and option data results again in somewhat more favorable estimation results: The bias is small for all parameter estimates and it performs best on the MSE criterion. Moreover, the smoothed volatility series is particularly close to the underlying volatility series. This combination of data contains the most precise information. Thus, if RV data is available for estimation, its use in combination with option data is preferable.

The results also indicate that the information in option data dominates the estimation results when using either both squared return-option data, or RV-option data for estimation. Option data is clearly very informative.

The Monte-Carlo results finally confirm that market prices of risk are difficult to estimate in practice; see e.g. Pan (2002) and Van der Ploeg et al. (2003).

although the bias in the estimates of the market price of volatility risk is small, the precision is very poor. This is irrespective of the data used for estimation.

3. Explorative data analysis

We examine daily data on the most important stock index of the United Kingdom, the FTSE100 index, and the European option contract traded on that index. The original dataset covers 9 years of data for the period 4 Jan 1993 till 28 Dec 2001. As motivated below however, we only use data for the period 6 Oct 1997 – 28 Dec 2001 for estimation. The input to the state space model consists of daily annualized squared returns (in deviation from their conditional mean) and three ATM Black-Scholes implied variance series; a SM, MM and LM series. Table 3.1 presents summary statistics.

Table 3.1: Summary statistics for the period Oct 1997 – Dec 2001

	Returns	Sq.returns	SM series: BS impl.vol.	MM series: BS impl.vol.	LM series: BS impl.vol.
Mean	0.0000	0.0413	0.2325	0.2393	0.2389
Median	0.0003	0.0151	0.2184	0.2296	0.2322
Std.deviation	0.0126	0.0676	0.0573	0.0468	0.0435
Skewness	-0.14	4.23	1.33	1.36	1.28
Kurtosis	3.67	37.73	4.94	5.18	4.68
			Moneyness	Moneyness	Moneyness
Mean			1.000	1.000	1.000
Std.deviation			0.002	0.005	0.005
			Maturity	Maturity	Maturity
Mean			0.116	0.500	0.872
Std.deviation			0.025	0.074	0.072
Minimum			0.077	0.377	0.750
Maximum			0.169	0.638	1.015

The table reports summary statistics for the daily (prewhitened) FTSE100-index returns, the annualized squared daily (prewhitened) FTSE100-index returns, and the SM, MM, and LM option series for the period 6 Oct 1997 – 28 Dec 2001. Moneyness is defined here as the ratio of futures price and strike price. The maturity is reported in years.

3.1 Details on data collection and construction

Let us detail the data collection and construction. The source of the FTSE100-index data is DataStream. Trading at the London Stock Exchange ends at 4.30 pm. DataStream records the daily closing prices of the index, which we use to compute the daily index-return series r_t. According to our model the FTSE100 index evolves according to the SDE $dS_t = \mu_t S_t dt + \sigma_t S_t dW_{S,t}$, which implies that non-overlapping returns are virtually uncorrelated. We did not explicitly specify the drift process μ_t in our theoretical discussion. Our estimation method nevertheless requires a prior estimate of μ_t at each point in time; recall (1.11). We obtain such an estimate as follows. From the Euler discretization of the stock price SDE it holds by approximation that $r_{t+\Delta t} \mid \mathcal{F}_t \sim \mathcal{N}(\mu_t \Delta t, \sigma_t^2 \Delta t)$. Hence, $\mu_t \Delta t$ basically represents the conditional mean daily index return. Estimating an AR(2) model for these returns results in residuals that are not significantly

autocorrelated [13]. We estimate $\mu_t \Delta t$ by the fitted returns from this AR(2) regression [14]. The (uncorrelated) residuals of this regression are known as prewhitened or prefiltered returns [15]. The squared prewhitened returns thus form the input to (1.11).

The source of the option data is the London International Financial Futures and Options Exchange (LIFFE) [16]. The full dataset consists of daily closing prices on a wide range of different call and put options for a total of 902,445 observations. Specifically, for each observation the dataset contains the date, call/put flag, strike price, option price, expiry month, open interest, volume, and the daily settlement price of the FTSE100-index futures contract that has the same maturity as the option contract. The daily settlement time of the futures contract is at 4.30pm; option trading ends at 4.30pm as well. As we select ATM options only as these are most liquid, we expect possible non-synchronicity between daily settlement times to be least a problem. For simplicity we assume non-synchronicity biases between index-futures prices, the FTSE100 index, and the option prices to be negligible. (Recall however that the estimation method implicitly allows for such possible measurement error.)

With regard to the available contracts, on any given trading day index options are traded with expiry months March, June, September, December, and additional months such that the 3 nearest calendar months are always available for trading. The options expire on the third Friday of the expiration month. On each day there are always 6 different-maturity contracts traded, so that the maximum maturity of all contracts on any given day ranges from 9 months to 12 months. For each of the available maturities a wide range of options trades that differ in strike prices.

Extracting option series from the database
We extract three close-to ATM call option series from the database of various maturities: a SM, a MM and a LM series. As the vega of an ATM option is maximal and as it is typically most liquid, we expect it to contain the most valuable volatility information. The maturity of the SM option series ranges from 20 to 44 trading days with an average of 30 trading days, or 1.4 months. The maturity of the MM series ranges from 4.5 to 7.7 months and averages at 6 months. The maturity of the LM option series ranges from 9 to 12 months with an average of 10.5 months. We refer to table 2.1 for details.

Let us explain in more detail how we extracted the three option series from the database. Consider e.g. the SM series. We first separated the calls from the puts. For each day we next selected the calls with a maturity of at least -but at the

---

[13] The fact that we find the index returns to be significantly autocorrelated is generally believed to be caused by non-synchronous trading effects associated with the individual stocks that constitute together the FTSE100 index; see e.g. Campbell, Lo and MacKinlay (1997).

[14] Doing so, our estimation procedure thus essentially boils down to two-step estimation and volatility extraction of the parameters of the model $r_{t+\Delta t} \mid \mathcal{F}_t \sim \mathcal{N}(\phi_0 + \phi_1 r_t + \phi_2 r_{t-\Delta t}, \sigma_t^2 \Delta t)$ in which σ_t is the annualized stochastic volatility. In the first step we estimate the conditional mean parameters ϕ_0, ϕ_1, ϕ_2; in the second step we estimate the parameters governing the stochastic volatility and moreover extract the latent volatility series. Notice the clear analogy with an AR(2)-GARCH(1,1) model for daily stock returns. But there is one important difference: In our model we have true stochastic volatility, whereas in the GARCH model today's volatility is a deterministic function of yesterday's information set.

[15] E.g. Andersen et al. (2002) also prefilter the data prior to estimation to get rid of the autocorrelation.

[16] We thank Joost Driessen from the Finance Department of the University of Amsterdam for providing us the option data. For a complete description of the option data see http://www.liffe.com.

same time closest to- 20 trading days. Given these, we finally selected that call option each day that is closest ATM by minimizing $|\ln(F_{ti}/K_i)|$, where K_i is the strike price of call i and F_{ti} is the current futures price of the index-futures contract that has the same maturity as call i. Following Hull (2003) we assume forward and futures prices to be equal. This procedure eventually left us with the SM-ATM call option series, which counts 2258 daily call observations. Its maturity range turned out to be 20 – 44 trading days. We extracted the MM-ATM and LM-ATM option series in similar ways; both contain 2258 observations as well. We carefully ensured that the characteristic series (i.e. price, strike, maturity, interest rate (see below), FTSE100 index, index futures price etc.) associated with each of the option series contain as many observations, and that all observation dates exactly "match"; i.e. are fully synchronous [17].

To compute the three associated BS implied volatility series we use equations (1.7) and (1.8). This requires for each option on each day an estimate of its associated average risk-free interest rate $\bar{R}_t$ over its remaining life. We obtain such an estimate by performing linear interpolation between the two nearest (in terms of maturity) continuously compounded LIBOR rates on that day. The daily LIBOR rates for terms of 1 month, 2 months up to 12 months are taken from DataStream. For our sample period they fluctuate in between 3.8% and 7.7%.
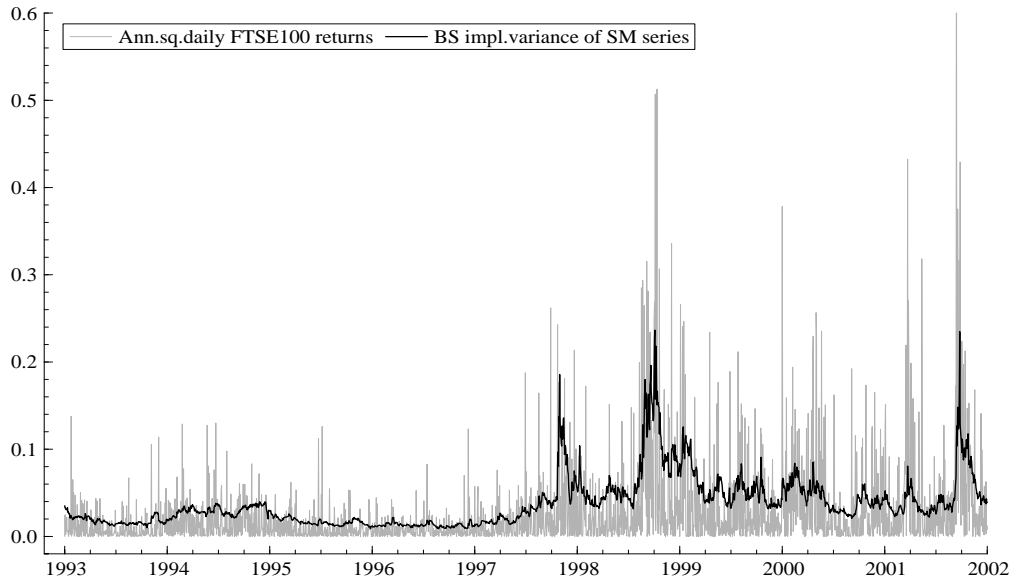
Figure 3.1: Annualized squared daily (prewhitened) FTSE100 returns and the BS implied variances associated with the SM call option series.

3.2   Data exploration and the volatility term structure

Figure 3.1 displays the annualized squared daily (prewhitened) FTSE100-index returns $\frac{1}{\Delta t}(r_{t+\Delta t} - \hat{\mu}_t \Delta t)^2$ and the BS implied variances $\sigma^2_{implied,t}$ of the SM option series in one graph. Notice the clear increase in both the level of the volatility and

---

[17] Data on 20 trading days are missing in the original database. Furthermore, the data for 1-10-1997 is incomplete and the data for 28-5-1998 contains obvious errors, such that we have discarded these dates. This leaves us with three option series of each 2258 observations, whose observation dates exactly match.

the volatility-of-volatility after notably the third quarter of 1997. The effects of both the Asian crisis and the near-collapse of hedgefund Long-Term Capital Management are apparent (fall 1997). Moreover, the impact of the crisis in the Soviet Union, the continuing Asian crisis, the uncertainty about the consequences of the European Monetary Union and the euro, and the uncertainty about president Clinton's retreatment due to the Monica Lewinsky affair –which all started in the fall of 1998- is clearly visible. The influence of September 11, 2001 and the subsequent war on terrorism is also very obvious [18]. These events have led to increased fluctuations in financial markets in the last years.

Figure 3.2 plots the SM, MM, and LM BS implied volatility series and confirms the augmented volatility and volatility-of-volatility from the fourth quarter of 1997 onwards. The series look very different in the first and second part of the sample. To a priori avoid model misspecification as much as possible, we opt to use only data for the period 6 Oct 1997 – 28 Dec 2001 in all our estimations, for a total of 1058 daily observations. For example, for the period till, respectively after, Oct 1997 the estimate of the unconditional stock volatility based on the average of the annualized squared returns equals 11.4%, respectively 20.3%. If we assume the volatility to be driven by one-factor SV which has an unconditional volatility of $\sqrt{\theta}$, this would imply that θ has more than tripled in the second part of the sample. Using all data would lead to obvious biases in that case [19].
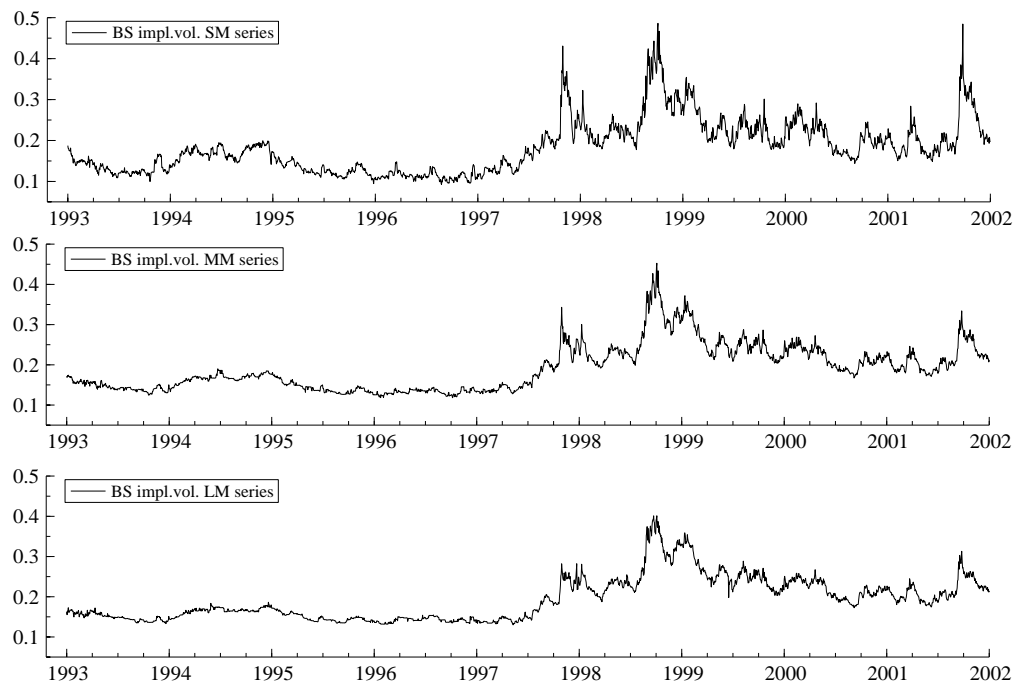

Figure 3.2: The Black-Scholes implied volatility series associated with the short-maturity, medium-maturity and long-maturity call option series.

---

[18] The peak of 0.86 in the annualized squared FTSE100 return at September 11, 2001 has been cut off at 0.6 to enhance the visual quality of the graph.
[19] Stated differently, although using the longer time series would probably (spuriously) increase the precision of the estimates, this is deceptive and hence not to be advocated. As it seems that the DGP has fundamentally changed from 1993:3 onwards (due to global developments), this increased "precision" is likely due to the false assumption (and hence misspecification) of the DGP being the same in both subsamples.

Notice from table 3.1 and figure 3.2 that the SM implied volatility series fluctuates most, whereas the LM series fluctuates least. As BS implied volatilities are often considered as forward-looking in practice (see e.g. Hull (2003)), and assuming there is mean reversion in the volatility, this makes sense.

Volatility term structure
In real-world financial markets one typically finds that BS implied volatilities of options written on the same underlying differ across strike prices and times to maturity (see e.g., Hull (2003)). If the Black-Scholes model were true the *implied volatility surface* should be flat at all times. Instead, for a given maturity one typically observes that the implied volatilities exhibit a smile, smirk or skewed pattern; i.e., the *volatility smile* or *smirk*. For given moneyness one may observe a variety of shapes of the so-called *volatility term structure,* which plots the BS implied volatilities against term to maturity of the options. The volatility surface serves an important function in practice: traders use it as a sophisticated interpolation tool for pricing vanilla options [20] consistently with the market. Moreover, option prices are generally quoted in terms of BS implied volatilities.

In the present paper we focus on the volatility term structure of ATM options. Given the three option series, the (end-of-day) volatility term structure can have one of four possible shapes on each of the 1058 days in our sample. It may either be upward sloping, exhibit a hump shape, be downward sloping, or display an inverted hump shape. Table 3.2 reports the shape frequencies (in percents), and assigns each possible shape a number for later reference. Notice that the ATMF volatility term structure is upward sloping on almost 50% of the days. The inverted hump shape occurs rather seldom.

Table 3.2: Shape frequencies volatility term structure ATMF options (Oct 1997 – Dec 2001)

Shape of vol. term structure	1: upward sloping	2: hump shape	3: downward sloping	4: inverted hump shape
Frequency	48.7%	27.5%	18.3%	5.5%

To explore the daily dynamics in the ATMF volatility term structure, figure 3.3 provides two matrices. The first matrix reports the number of volatility-term-structure shape transitions as a percent of the total number (i.e. 1057) of transitions. For example, of all transitions 9.4% entailed a change in the shape of the volatility term structure from upward sloping to hump shape from one day to the next. The second matrix reports the empirical shape-transition "probabilities"; each row adds to 100%. Given that the volatility term structure has a particular shape today, it is most likely that it has equal shape tomorrow.

---

[20] That is, European call and put options. Exotic derivatives cannot be priced with the surface unless the exotic can be decomposed into a particular combination of vanilla options. Note that European call and put options written on the same underlying with the same strike and maturity have identical implied volatility (this follows from the put-call parity (see Hull (2003)), such that the surface can be used to price both call and put options.

Shape transition frequencies (in %)					Shape transition 'probabilities' (in %)				
t+1 \ t	**1**	**2**	**3**	**4**	t+1 \ t	**1**	**2**	**3**	**4**
1	37.2	9.4	0.9	1.2	**1**	76.4	19.3	1.8	2.5
2	9.9	14.9	1.7	1.0	**2**	36.0	54.0	6.2	3.8
3	0.4	2.1	14.4	1.5	**3**	2.1	11.3	78.4	8.2
4	1.2	1.2	1.3	1.7	**4**	22.4	22.4	24.1	31.1

Figure 3.3: Day-to-day dynamics in the shape of the ATMF volatility term structure (Oct 97 – Dec 01). Left: transition frequencies in % of total number of transitions. Right: empirical transition 'probability' matrix (in %). Shape codes are reported in table 3.2.

The frequency of parallel shifts in the volatility term structure is of interest as well. We define a parallel shift to be one in which both the SM, MM and LM BS implied volatilities either jointly go up from one day to the next, or go down (but not necessarily by the same amount). Table 3.3 shows that parallel shifts occur 66% of the time, whereas 34% of the shifts are non-parallel. It is clear that the volatility term structure evolves in complex manners. A realistic model must obviously be capable of reproducing these complex observed dynamics. Evidently, the question is if one SV factor is sufficient in this respect.

Table 3.3: Daily shifts in the ATM volatility term structure (Oct 97 – Dec 01)

	Upward	*Downward*	*Total*
Parallel shifts	31.9%	34.5%	66.4%
Non-parallel shifts			33.6%
		Total:	100%

4. Estimation results 1-factor OU

The first set of estimation results we present is for the 1-factor OU case, for the period 6 Oct 1997 – 28 Dec 2001. In this case, the instantaneous stock variance $\sigma_t^2 = x_t$ evolves as $dx_t = k(\theta - x_t)dt + \sigma dW_{x,t}$ under $\mathbb{P}$, and under $\mathbb{Q}$ as $dx_t = \tilde{k}(\tilde{\theta} - x_t)dt + \sigma d\tilde{W}_{x,t}$. The risk-neutral parameters are given by $\tilde{k} = k$, $\tilde{\theta} = \theta - \sigma\gamma/k$. The unconditional stock return volatility equals $\sqrt{\theta}$. The invariant distribution of the stock variance is given by $\sigma_t^2 \sim \mathcal{N}(\theta, \sigma^2/2k)$, such that the volatility-of-the-variance equals $\sqrt{(\sigma^2/2k)}$. The parameters k and $\tilde{k}$ govern the volatility speed-of-adjustment towards its long-run mean under respectively $\mathbb{P}$ and $\mathbb{Q}$. The persistence in the daily stock variance (see (1.17)) is given by $\exp[-k\Delta t]$ with $\Delta t = 1/260$. The half-life (measured in days) of a shock occurring in the stock variance equals $\ln 2/(k\Delta t)$. The market price of volatility risk is given by γ.

As a benchmark, the first column of table 4.1 reports estimation results when only the squared returns are used for estimation. A Gaussian GARCH(1,1) model for the prewhitened returns is also estimated. As both our model and the GARCH model do not allow for the leverage effect, this permits a consistent comparison. The GARCH conditional stock variance follows $\sigma_{t+\Delta t}^2 = \varphi_0 + \varphi_1(r_t - \mu\Delta t)^2 + \varphi_2\sigma_t^2$. We find

$$\sigma_{t+\Delta t}^2 = \underset{(4.78e-06)}{9.92e-06} + \underset{(0.0281)}{0.0861} (r_t - \hat{\mu}\Delta t)^2 + \underset{(0.0496)}{0.851} \sigma_t^2 \tag{4.1}$$

15

with robust standard errors in parentheses. For this model the persistence in the daily variance is measured by $\varphi_1 + \varphi_2$. The half-life of a shock in the variance is given by $-\ln 2 / \ln(\varphi_1 + \varphi_2)$, whereas the per-annum unconditional stock-return volatility may be approximated by $\sqrt{\varphi_0 / [(1 - \varphi_1 - \varphi_2)\Delta t]}$. To obtain a rough estimate of the per annum volatility-of–the-variance, we compute the standard deviation of the constructed daily GARCH variance series after first having multiplied this series by 260. The second column of table 4.1. reports these quantities for the estimated GARCH model.

The third column shows the results when the combination return - SM option data is used for estimation, as was similarly done in e.g. Chernov and Ghysels (2000), Pan (2002), and Jones (2003). The fourth column presents the results when all four series are jointly used for estimation. These latter results ought to be taken with much care since the model appears heavily misspecified in this case; see below.

Table 4.1: Estimation results for Oct 1997 – Dec 2001

	Return data	GARCH	Return & SM	Return & SM+MM +LM		
θ	**0.0409**		**0.0450**	**0.0475**		
	(0.0052)		(0.0317)	(0.0408)		
k	**16.3**		**1.85**	**0.321**		
	(6.93)		(1.78)	(0.0714)		
σ	**0.164**		**0.151**	**0.0639**		
	(0.049)		(0.022)	(0.0063)		
γ	-		**-0.967**	**0.00389**		
			(0.527)	(0.206)		
σ_ω	**0.0612**		**0.0657**	**0.0704**		
	(0.0056)		(0.0062)	(0.0057)		
$\sigma_{\varepsilon,SM}$	-		**0.0038**	**0.0193**	$\text{corr}(\varepsilon_{SM,t}, \varepsilon_{MM,t})$	**0.950**
			(0.0018)	(0.00124)		(0.0122)
$\sigma_{\varepsilon,MM}$	-		-	**0.00728**	$\text{corr}(\varepsilon_{SM,t}, \varepsilon_{LM,t})$	**0.895**
				(0.00089)		(0.0482)
$\sigma_{\varepsilon,LM}$	-		-	**0.00394**	$\text{corr}(\varepsilon_{MM,t}, \varepsilon_{LM,t})$	**0.781**
				(0.00071)		(0.0770)
Vol. returns	20.2%	20.3%	21.2%	21.8%		
Vol-of-var.	0.0287	0.0175	0.0785	0.0799		
Persistence	0.939	0.937	0.993	0.9988		
Half-life	11	11	98	562		
Std.dev. u_t	0.0099		0.0093	0.0040		
Loglikelih.	1394	3165	4761	13,135		

The table reports parameter estimates (boldface) with robust White (1982) QML standard errors (in parentheses), resulting from estimating the state space model using three types of data: only return data (first column), both return and SM option data (third column), and both return, SM, MM and LM option data (fourth column. The table also reports the unconditional volatility of returns $\sqrt{\theta}$, the volatility-of-the-variance $\sqrt{\sigma^2 / 2k}$, the persistence in the daily stock variance $\exp[-k\Delta t]$, the half-life (in days) of a shock in the stock variance, $\ln 2 / k\Delta t$, the MLE of the standard deviation of u_t, and the maximized loglikelihood value. The second column shows some comparable quantities for the estimated Gaussian GARCH(1,1) model (4.1).

Let us therefore focus on the first three columns of table 4.1. The return volatility is estimated similar as the estimate of 20.3% computed as the root sample average of the squared returns. The persistence is estimated smallest when the model is estimated using return data only. When both return and option data are

used, the estimated persistence dramatically increases. (A value of 1 would imply that shocks have a permanent effect.) The market price of volatility risk is estimated at −0.967 and hence negative, which agrees with, e.g., Coval and Shumway (2001), Buraschi and Jackwerth (2001), Pan (2002), Jones (2003), Bakshi and Kapadia (2003), Driessen and Maenhout (2003) and Carr and Wu (2004) [21].

Section 4.1 first considers diagnostic checking of the model estimated using return – SM data, as this is the typical dataset other researchers (e.g., Chernov and Ghysels (2000) and Jones (2003)) have used when combining stock and option data for estimating 1-factor SV option pricing models. Although in that case the dynamic misspecification may perhaps seem rather modest at first sight, we next confirm that the 1-factor model severely overprices the longer-dated options out of sample, as also reported in Chernov and Ghysels (2000) and Pan (2002). This is a first possible indication of the presence of multiple volatility-driving factors. Section 4.2 next considers diagnostic checking when all four time series are jointly used for estimation. The tests clearly reveal a lack of dynamics in the 1-factor model.
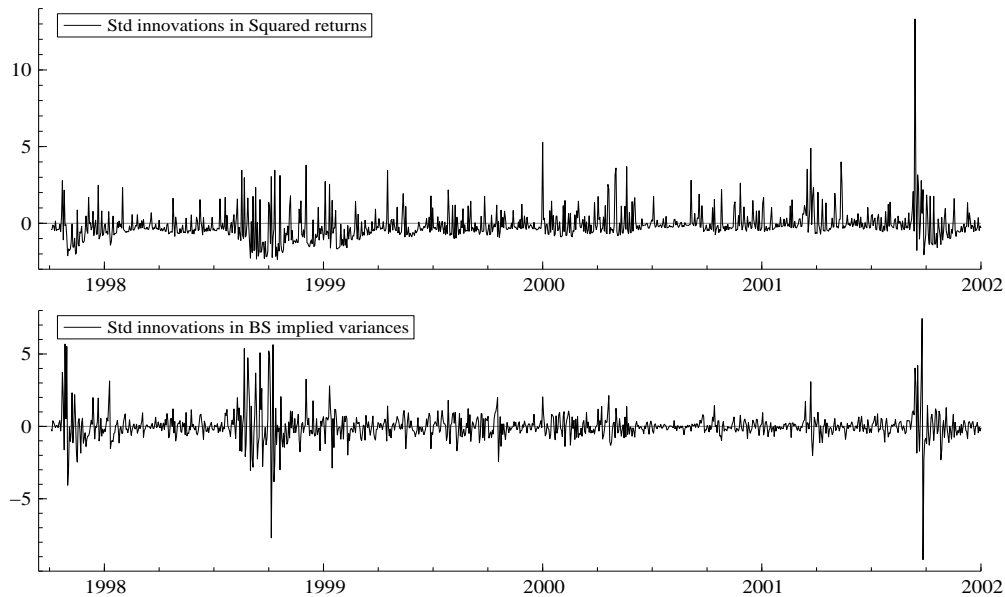
Figure 4.1: Standardized innovations (1-factor OU estimated using return-SM option data).

4.1 Specification tests of the 1-factor OU model estimated with return and SM option data

Given the estimation output, the *standardized prediction errors* or *innovations* [22] of the state space model can be computed. If the model is well specified these

---

[21] Although beyond the scope of the current paper, Van der Ploeg (2004b) shows that this estimate implies an expected return of −3.3% per week on SM-ATM FTSE100-index straddles. This is in close correspondence with the weekly empirical (i.e. "raw") S&P500 SM-ATM straddle return of −3% reported in Coval and Shumway (2001), and the *monthly* empirical SM-ATM FTSE100-index straddle return of −13.1% reported in Driessen and Maenhout (2003).
[22] See, e.g., Harvey (1989) and Durbin and Koopman (2001) for details on specification testing of state space models.

innovations ought to be white noise with a variance equal to 1. Figure 4.1 displays the innovations associated with the squared-return equation (1.11) in the upper graph, and for the BS implied-variance equation (1.12) in the lower graph. Their standard deviations equal 0.993 and 1.000 respectively. Both series are clearly heteroskedastic.

As evidenced by table 4.2, both innovation series are significantly autocorrelated [23], though the autocorrelations do not seem particularly large. This is evidence of misspecification and may be interpreted as a preliminary indication of neglected dynamics.

Table 4.2: Autocorrelations in the standardized innovations

Statistics for the standardized innovations in the squared returns:				*Statistics for the standardized innovations in the SM BS implied variances:*			
Order	AC	Q-Stat	Prob	Order	AC	Q-Stat	Prob
1	0.080	6.835	0.009	1	0.039	1.582	0.208
2	0.056	10.15	0.006	2	-0.151	25.89	0.000
3	0.133	28.88	0.000	3	-0.107	38.09	0.000
4	0.114	42.61	0.000	4	0.043	40.10	0.000
5	0.050	45.24	0.000	5	0.005	40.13	0.000

The table reports the autocorrelation coefficients (AC) up to order 5 and the Ljung-Box Q-statistics (Q-stat) for testing the null of no autocorrelation up to a certain order with associated p-values (Prob).
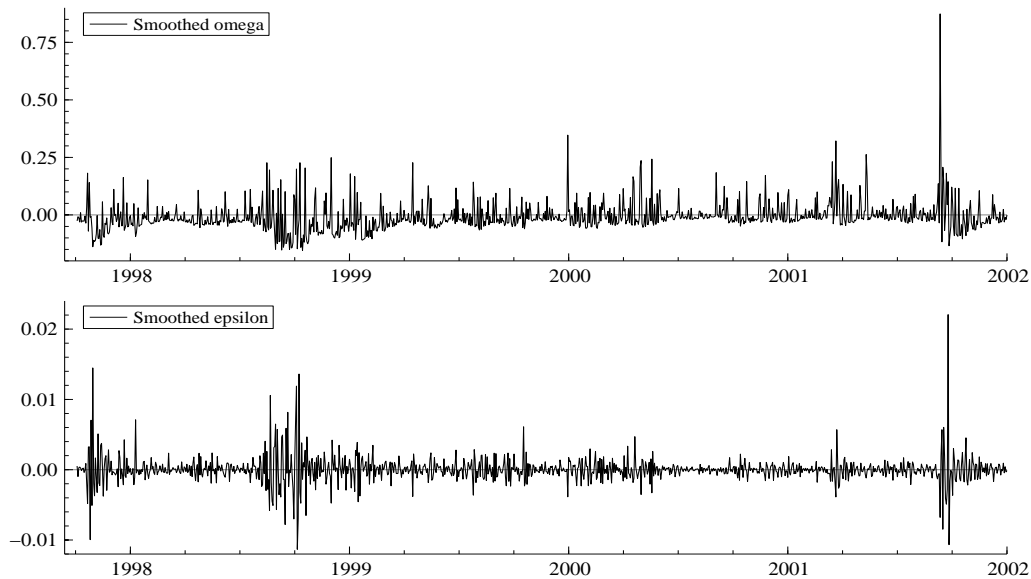
Figure 4.2: Smoothed $\{\omega_t\}$-series (upper plot) and smoothed $\{\varepsilon_{SM,t}\}$-series (lower plot).

The so-called *smoothed disturbances* or *auxiliary residuals* of the state space model yield additional information with respect to model misspecification. They

---

[23] It should be noted that the Ljung-Box Q-test assumes homoskedasticity of the underlying series, which seems obviously violated in this case.

represent the best estimates of the disturbances given the data [24]. Figure 4.2 shows those belonging to the measurement equation. There is clear evidence of conditional heteroskedasticity. The asymmetric pattern in the smoothed $\{\omega_t\}$-series is a consequence of the course of the squared return series over time; recall figure 3.1. When an ARMA(1,1) model is fitted to the squared returns a similar residual picture results.

Figure 4.3 displays the smoothed state $x_t^* = x_t - \theta$ and the smoothed state disturbance u_t of the state space model. When the level of the state is high resp. low, the state disturbance is generally large resp. small (in absolute value). As the state governs the stock volatility this is clear evidence of *volatility-feedback* (see also, e.g., Jones (2003) and Chernov et al. (2003). We already mentioned that the 1-factor OU SV assumption does not allow for this effect, and is thus additionally misspecified in this sense.
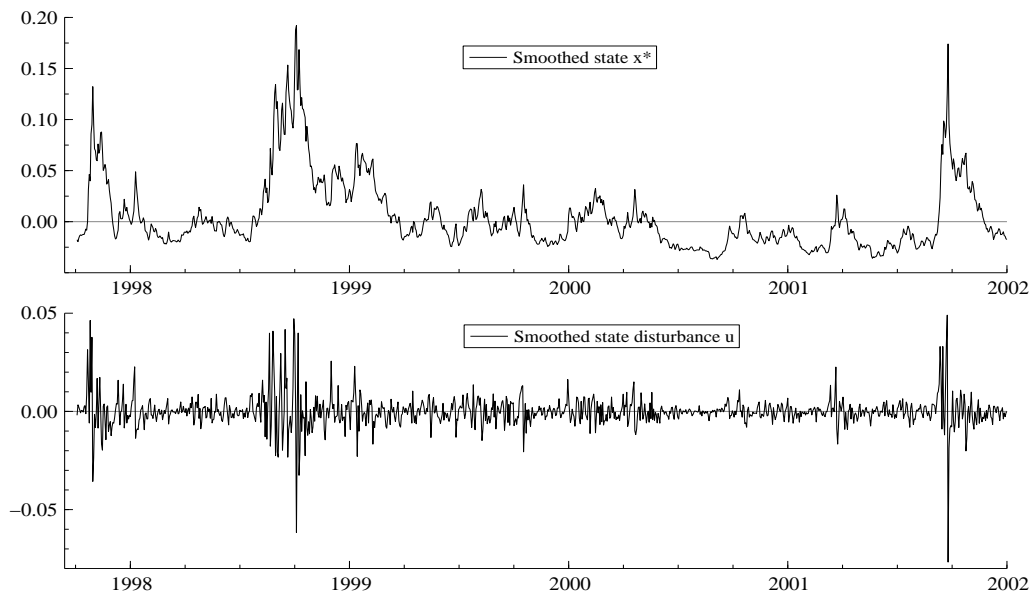
Figure 4.3: Smoothed state x_t^* (upper graph) and smoothed state disturbance u_t.

Fit of the volatility term structure
Let us next investigate to what extent the 1-factor model is capable of fitting the longer end of the volatility term structure out of sample; i.e. the MM and LM option series. Given the one-to-one correspondence between BS implied volatilities and prices, this fit is implicitly linked to option pricing errors. However, as traders typically quote prices in terms of BS implied volatilities, we choose to interpret the fit in terms of deviations between "observed" and fitted implied volatilities. Based on the measurement equation for the options [25] we can compute fitted implied volatilities by

---

[24] Durbin and Koopman (2001) point out that the smoothed disturbances are autocorrelated (in theory/population), but can be useful for detecting outliers and structural breaks, as they are the best estimators of the error terms given the data.
[25] Notice that we avoid computationally intensive Monte-Carlo simulations to first obtain the model-implied call prices and then transform them to BS implied volatilities.

$$\hat{\sigma}_{implied,it} = \sqrt{\frac{1}{\tau_{it}}\left[A(\tau_{it}) + \mathbf{B}(\tau_{it})'\mathbf{\theta}\right] + \frac{1}{\tau_{it}}\mathbf{B}_1(\tau_{it})'\mathbf{x}_t^*} \,, \tag{4.2}$$

for $t = \Delta t, .., T\Delta t$ and $i = SM, MM, LM$, in which we substitute the parameter estimates and smoothed $\{x_t^*\}$-series. The "pricing" errors expressed in percents BS implied volatility are then given by

$$error_{it} = \sigma_{implied,it} - \hat{\sigma}_{implied,it} \,. \tag{4.3}$$

The left panel of figure 4.4 shows the original and fitted SM, MM and LM BS implied volatility series; the right panel displays the error series.
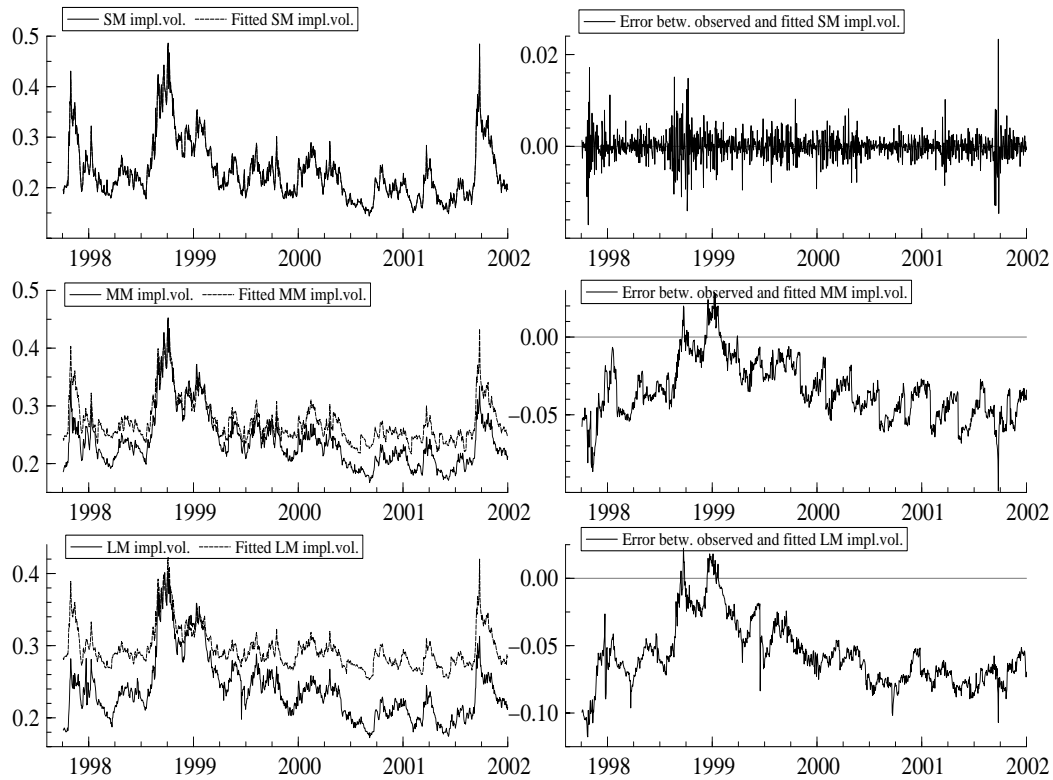
Figure 4.4: Fit of the ATM volatility term structure: In-sample fit of the SM BS implied volatility series and out-of-sample fit of the MM and LM BS implied volatility series. The figure shows the original and fitted BS implied volatility series for each maturity, together with their difference. This difference may be interpreted as a pricing error.

The 1-factor model generally overfits the MM and LM implied volatility series out of sample. Relative to the observed market values, the model thus overprices longer-dated ATM options, and the mispricing seems to get worse the longer the maturity. Table 4.3 reports the mean and standard deviation of each of the three error series. The overfitting can be substantial.

Table 4.3: Summary statistics error series

	Error in SM impl.vol.	Error in MM impl.vol.	Error in LM impl.vol.
Mean	-0.0069%	-3.38%	-5.66%
Std.deviation	0.31%	1.89%	2.38%

The table reports the mean and standard deviation of the error series shown in figure 4.4.

Chernov and Ghysels (2000) reach a similar conclusion (irrespective of moneyness) when assessing the out-of-sample pricing performance of the Heston (1993) model. Pan (2002) also reports the severe out-of-sample overpricing of long-dated near-the-money options by the Heston model. Pan attributes this to the fact that her estimation results imply an explosive risk-neutral volatility process. Interestingly, our results imply a stationary volatility process under $\mathbb{Q}$ as $\tilde{k} = 1.85$ and $\tilde{\theta} = 0.124$, which are both positive.

4.2 Specification tests of the 1-factor OU model estimated with return, SM, MM and LM option data

We now turn to specification testing of the 1-factor OU model when all four series are jointly used for estimation (recall column 4 of table 4.1). As all data will be in used in subsequent multifactor estimations, a comparison with the results presented in this section proves useful in the contribution of each additional SV factor.

Figure 4.5 plots the four standardized innovation series associated with the return and BS implied variance equations (1.11) and (1.12). Table 4.4 reports the relevant summary statistics. The misspecification is obvious. Notice in particular the severe autocorrelation [26] in each series. This seems confirmative of one volatility-driving factor not being sufficient to describe all dynamics present in the data.
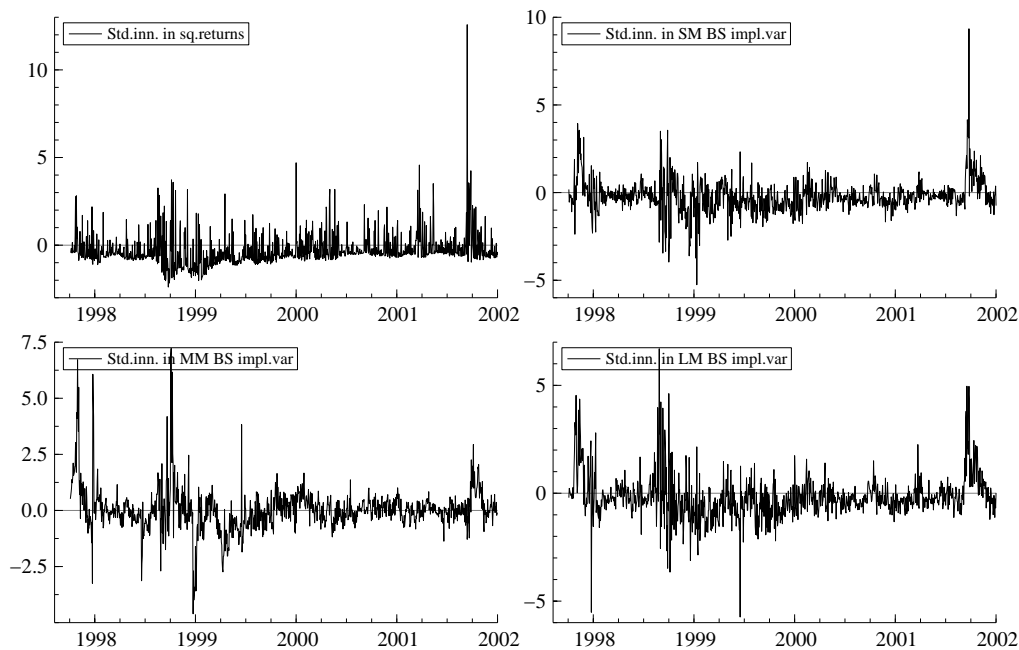

Figure 4.5: Standardized innovations associated with the measurement equations for the squared returns (upper left), the SM (upper right), MM (lower left) and the LM (lower right) BS implied variances (1-factor OU).

---

[26] The p-values associated with Ljung-Box Q-test statistics for testing for zero-autocorrelation up to order 5 all equal 0.000.

Table 4.4: Summary statistics of the standardized innovations (1-factor OU; all data)

	Std.inn. sq. return	Std.inn. SM series	Std.inn. MM series	Std.inn. LM series
Mean	-0.284	-0.209	0.049	-0.168
Std.deviation	0.959	0.946	0.992	1.030
AC(1)	0.138	0.395	0.617	0.433
AC(2)	0.115	0.350	0.534	0.372
AC(3)	0.179	0.301	0.443	0.264
AC(4)	0.167	0.246	0.423	0.339
AC(5)	0.107	0.275	0.373	0.328
Cont.correlation matrix	1.000	0.089	0.052	0.338
		1.000	-0.020	-0.042
			1.000	0.033
				1.000

AC(i) stands for the i-th order autocorrelation coefficient.

Figure 4.6 shows the smoothed disturbance series of the state space model. The misspecification is most apparent from the graphs of the three $\{\varepsilon_t\}$ – series, which ought to look like white noise approximately. They clearly do not.
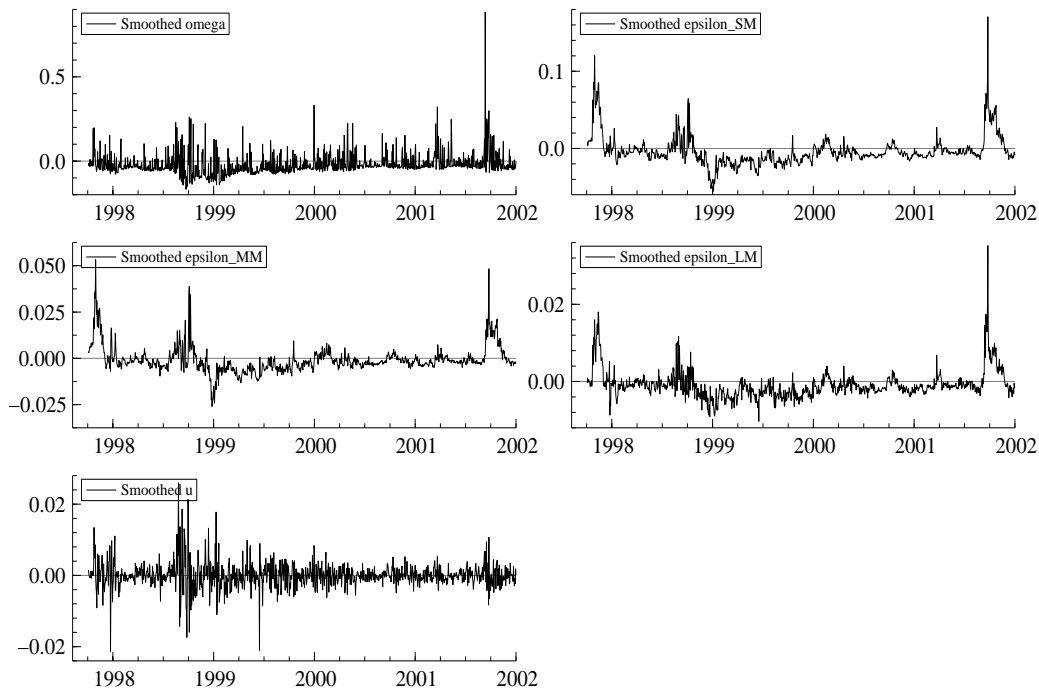

Figure 4.6: Smoothed disturbances of the state space model (1-factor OU, all data).

In-sample fit of the volatility term structure
We also examine the extent in which the 1-factor model estimated with the full dataset fits the ATM volatility term structure in sample. (Obviously, figure 4.6 already provide a clue but then in terms of less easy to interpret BS implied variances instead of volatilities.) Fitted SM, MM and LM BS implied volatilities are again computed by (4.2) and "pricing" errors by (4.3). Figure 4.7 shows the original and fitted series together with their difference. The errors can be substantial.
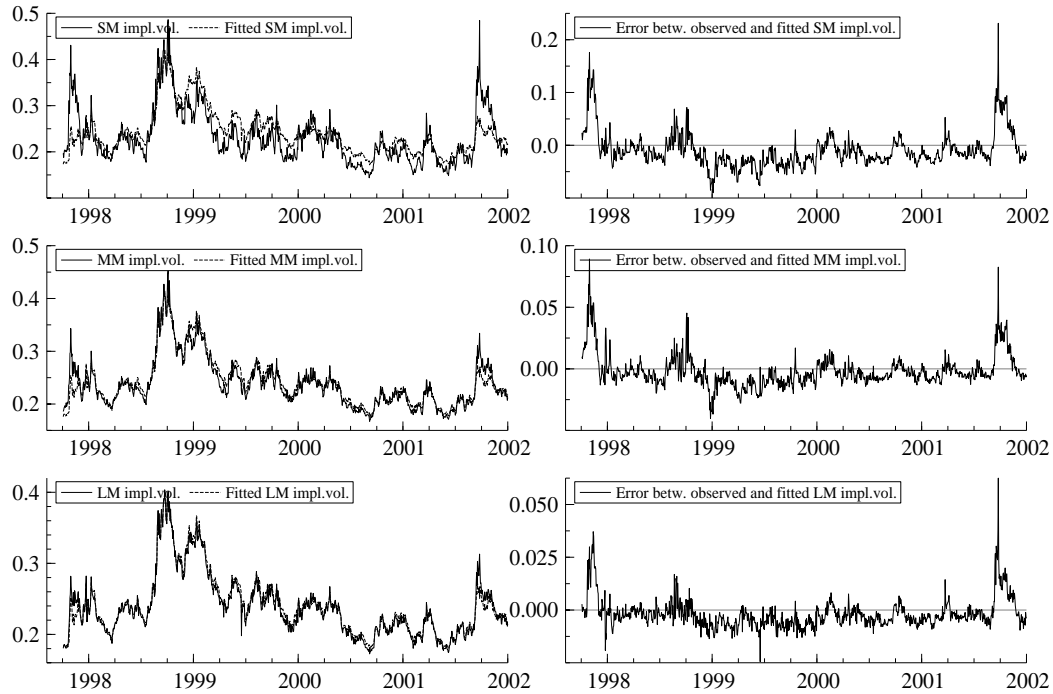
Figure 4.7: In-sample fit of the ATM volatility term structure (1-factor OU, all data). The figure shows the original and fitted BS implied volatilities for each maturity, together with their difference. This difference may be interpreted as a pricing error.

Table 4.5 reports the mean and standard deviation of each of the (absolute) error series. On average, the 1-factor model predicts somewhat higher prices than observed in the market. Notice that the short end of the volatility term structure is generally fitted worst and the long end best.

Table 4.5: Summary statistics error series (1-factor OU)

	Error SM	*Error MM*	*Error LM*	*|Error SM|*	*|Error MM|*	*|Error LM|*
Mean	-1.01%	-0.22%	-0.20%	2.62%	0.93%	0.53%
Std.dev.	3.33%	1.31%	0.70%	2.30%	0.95%	0.51%

The table reports the mean and standard deviation of the error series shown in figure 3.14. The error is defined as the difference between the true and fitted BS implied volatility. Both statistics are also reported for the absolute errors.

1-factor OU SV: Summary of results

When the 1-factor OU model is estimated using stock and SM option data jointly, the misspecification seems rather modest at first sight. Its main deficiency seems negligence of level-dependent volatility-of-volatility (apart from the leverage effect). However, subsequent out-of-sample valuation of longer-dated options generally leads to substantial overpricing. Given similar evidence in Chernov and Ghysels (2000) and Pan (2002) on the Heston (1993) model, this overpricing cannot be overcome by explicit modeling of volatility feedback and the leverage effect it seems.

When the model is estimated using the four joint time series the misspecification of the 1-factor model is even more prominent. The model does not fit the

volatility term structure very well; especially not in the short end. Although its deficiency of not being able to describe volatility feedback is a matter of concern, the most important misspecification seems neglected dynamics. The 1-factor model is not able to adequately describe the complex movements of the term structure, as observed in section 3.2. The various tests strongly reject the 1-factor assumption. We intuitively expect that fitting the Heston (1993) model to a similar dataset incorporating more than one option series would lead to similar conclusions [27]. This point is also made by Cont and Fonseca (2002).

Table 5.1: Estimation results 2-factor OU using return, SM, MM and LM option data jointly

	Restrictions (a)	Restrictions (b)		Restrictions (a)	Restrictions (b)
θ_1	**0.0431** (0.0342)	**0.0408** (0.0386)	$\text{corr}(\varepsilon_{SM,t}, \varepsilon_{MM,t})$	0	**0.548** (0.342)
θ_2	0	0	$\text{corr}(\varepsilon_{SM,t}, \varepsilon_{LM,t})$	0	**-0.089** (0.548)
k_1	**0.516** (0.080)	**0.445** (0.076)	$\text{corr}(\varepsilon_{MM,t}, \varepsilon_{LM,t})$	0	**-0.516** (0.658)
k_2	**6.72** (0.731)	**8.71** (1.37)			
σ_{11}	**0.0602** (0.0046)	**0.0669** (0.0058)	Vol. returns	20.8%	20.2%
σ_{12}	0	0	Vol-of-var.	0.0804	0.0858
σ_{21}	0	**0.0612** (0.0216)	Std.dev. x_{1t}	0.059	0.071
σ_{22}	**0.199** (0.018)	**0.146** (0.035)	Std.dev. x_{2t}	0.054	0.038
γ_1	**0.116** (0.258)	**0.090** (0.241)	$\text{corr}(x_{1t}, x_{2t})$	0	0.166
γ_2	**-0.330** (0.467)	**-0.661** (0.595)	Persistence x_1	0.9980	0.9983
σ_ω	**0.0671** (0.0057)	**0.0672** (0.0056)	Persistence x_2	0.9745	0.9671
$\sigma_{\varepsilon,SM}$	**1.59*10⁻⁵** (2.10*10⁻⁵)	**0.00306** (0.00158)	Half-life x_1	349	405
$\sigma_{\varepsilon,MM}$	**0.00201** (0.00016)	**0.00205** (0.00069)	Half-life x_2	27	21
$\sigma_{\varepsilon,LM}$	**0.00206** (0.00027)	**0.00165** (0.00038)	Std.dev. u_{1t}	0.0037	0.0041
Loglik.	14,305	14,363	Std.dev. u_{2t}	0.0122	0.0097

The table reports (restricted) parameter estimates (in boldface) with robust White (1982) QML standard errors in parentheses, resulting from estimating the state space model associated with the 2-factor OU SV assumption using the combination of return, SM, MM and LM option data under two sets of restrictions (see main text), together with the MLEs of some other quantities of interest.

5.   Estimation results 2-factor OU

In this section we present estimation results for 2-factor OU SV. As it appears that the model is still not able to capture most dynamics observed in the data, we keep the discussion rather short.

---

[27] To the best of our knowledge this has not been pursued so far in the literature.

As not all parameters can be identified we consider two sets of a priori imposed restrictions. In *Restrictions (a)* we consider the case of two independent volatility-driving factors and do not allow for contemporaneous correlation between the error series $\{\varepsilon_{it}\}, i = SM, MM, LM$. Restrictions (a) imposes $\theta_2 = 0$, $\sigma_{12} = \sigma_{21} = 0$, and $\text{corr}(\varepsilon_{SM,t}, \varepsilon_{MM,t}) = \text{corr}(\varepsilon_{SM,t}, \varepsilon_{LM,t}) = \text{corr}(\varepsilon_{MM,t}, \varepsilon_{LM,t}) = 0$. In *Restrictions (b)* we allow both for correlation between the factors and contemporaneous correlation between the error series $\{\varepsilon_{it}\}, i = SM, MM, LM$. In this case we impose $\theta_2 = 0, \sigma_{12} = 0$ such that $(2x2)\mathbf{\Sigma}$ is lower diagonal. Table 5.1 presents the estimation results.

Allowing for two OU factors instead of one raises the quasi-loglikelihood from 13,135 to more than 14,300. Allowing for two correlated factors does not improve the fit dramatically as compared to the independent-factors case. The maximum likelihood estimate of the factor correlation equals 0.166, which, given that σ_{21} differs significantly from zero, seems significant. (The model is however misspecified.) The contemporaneous correlations between the $\{\varepsilon_{it}\}$ series do not seem to differ significantly from zero. The volatility of returns, $\sqrt{\mathbf{1'\theta}}$, is estimated near the 20.3% obtained from averaging the squared returns. The two volatility factors differ in their characteristics. The first factor has similar properties as the factor distilled from the joint data in the 1-factor case. Its persistence $\exp(-k_1\Delta t)$ is close to the random walk value of 1, with a half-life $\ln 2 / k_1 \Delta t$ of about 1.45 years. Shocks in this factor die out very slowly. The second factor shows much quicker mean reversion with a half-life of about a month.
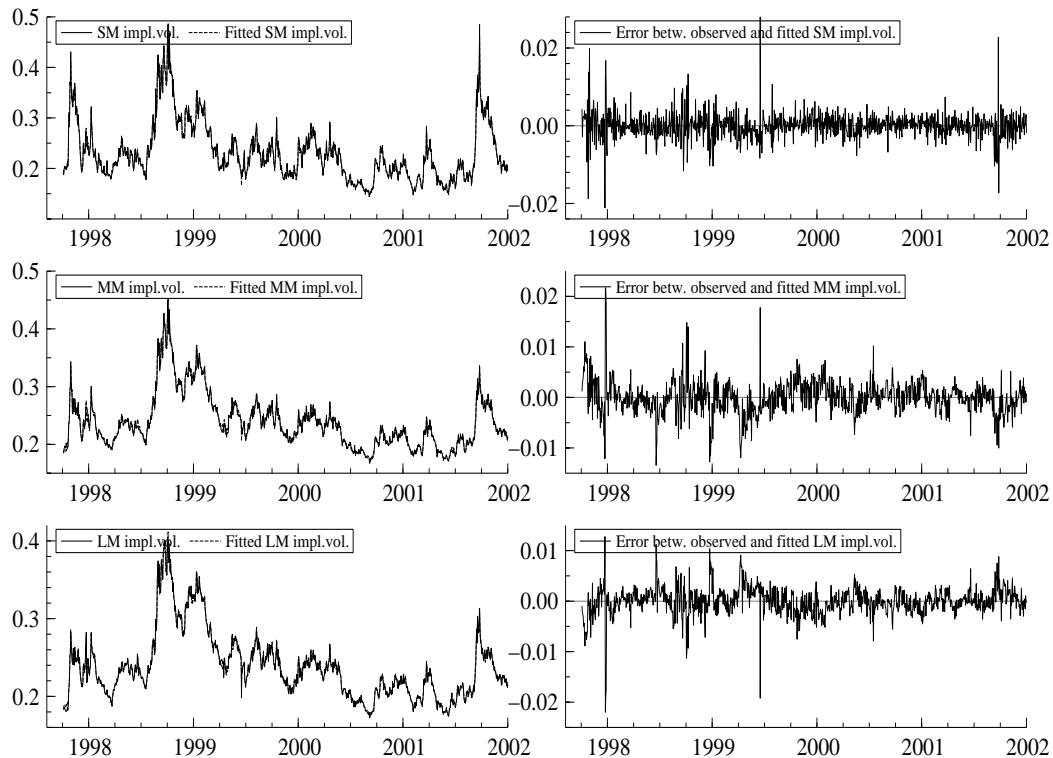

Figure 5.1: In-sample fit of the ATM volatility term structure (2-factor OU, Restrictions (b)). Left: original and fitted BS implied volatilities for each maturity. Right: their difference (i.e. "pricing" error).

In-sample fit of the volatility term structure

For further investigation and diagnostic checking we concentrate on the least restricted model (Restrictions (b)). Consider the in-sample fit of the volatility term structure. The left panel of figure 5.1 shows the "observed" and fitted SM, MM and LM BS implied volatilities in one graph. The errors are displayed in the right panel. Compare this figure to figure 4.7. Notice that the error graphs have improved in the sense that they seem more randomly distributed.

Table 5.2 reports the mean and standard deviation of each of the (absolute) error series. Compare this table to table 4.5. The substantial increase in fit achieved by allowing for two volatility factors instead of one is obvious. The biggest improvement is for the SM option series. The "pricing" errors concentrate around zero now, irrespective of maturity. The average fit is similar for all series, although the LM series is still fitted somewhat best.

Table 5.2: Summary statistics error series (2-factor OU, Restrictions (b))

	Error SM	Error MM	Error LM	|Error SM|	|Error MM|	|Error LM|
Mean	0.00%	0.00%	0.00%	0.23%	0.24%	0.19%
Std.dev.	0.34%	0.33%	0.28%	0.25%	0.24%	0.20%

The table reports the mean and standard deviation of the error series shown in figure 5.1. Both statistics are also reported for the absolute errors.

Specification tests

We now turn to specification testing. The upper panel in figure 5.2 shows the smoothed factors in deviation from their mean, $x_{it}^* = x_{it} - \theta_i$; $i = 1, 2$. The lower panel shows the smoothed transition equation errors u_{1t} and u_{2t}. Volatility feedback is present in both series; a feature not accounted for by the OU SV models. Interestingly, analyzing daily 1953-1999 data on the DJIA stock index, Chernov et al. (2003) find one very persistent and one quickly mean-reverting factor in a 2-factor logarithmic (i.e. exponential linear) SV stock price model. The persistent factor does not seem to feature feedback, whereas the other does.
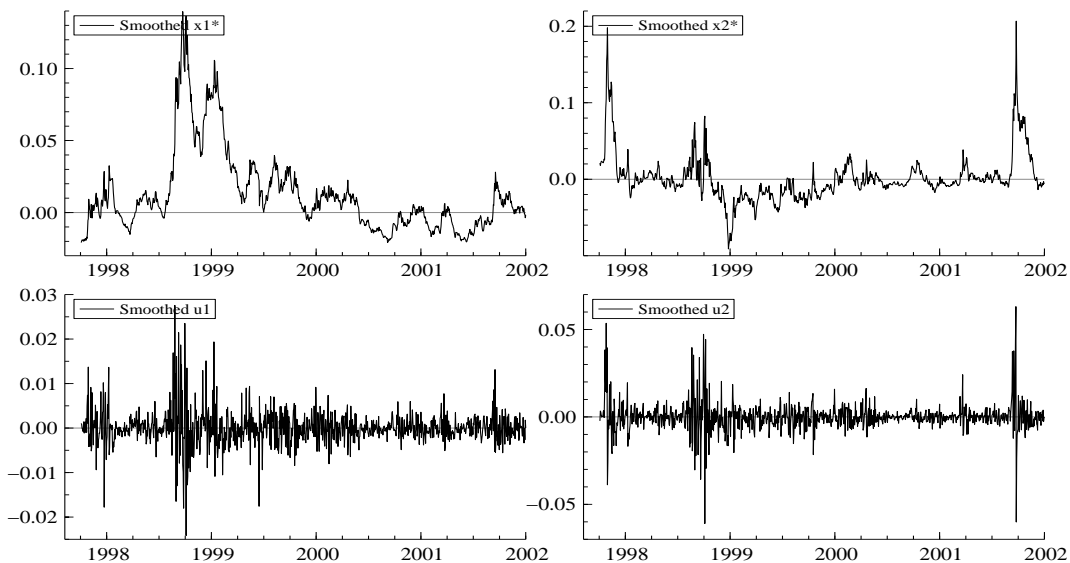

Figure 5.2: Smoothed factors in deviation from their mean (x_{1t}^* and x_{2t}^*), and smoothed state equation errors u_{1t} and u_{2t} (2-factor OU, Restrictions (b)).

Figure 5.3 shows the standardized innovations. Table 5.3 reports summary statistics for these series.
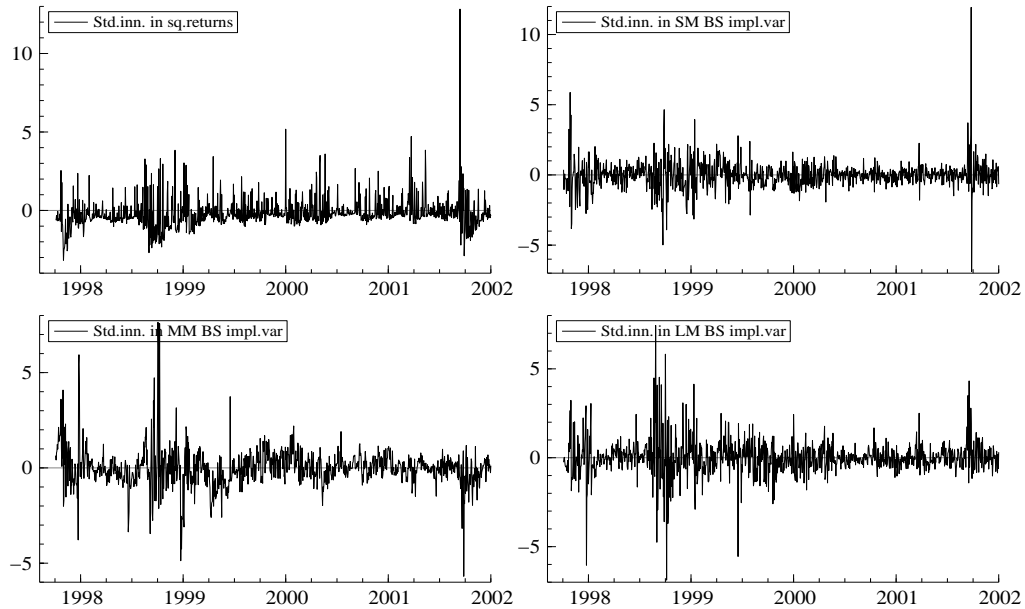

Figure 5.3: Standardized innovations (2-factor OU, Restrictions (b)).

Recalling figure 4.5 and table 4.4 for the 1-factor OU case, it is clear that the 2-factor assumption provides a much better description of the data. The innovations are closer to behaving as unit-variance white noise than in the 1-factor OU case. Especially prominent is the dramatic reduction in autocorrelation in three of the four series. The dynamics in the medium-maturity option series are however still not well captured by this model [28].

Table 5.3: Summary statistics standardized innovations (2-factor OU; Restrictions (b))

	Std.inn. sq. return	Std.inn. SM series	Std.inn. MM series	Std.inn. LM series
Mean	-0.193	-0.008	-0.006	0.010
Std.deviation	0.983	0.995	1.012	0.992
AC(1)	0.086	0.075	0.399	0.045
AC(2)	0.059	-0.016	0.231	0.021
AC(3)	0.136	-0.025	0.200	-0.100
AC(4)	0.113	-0.016	0.186	0.070
AC(5)	0.050	0.014	0.212	0.024
Cont.correlation matrix	1.000	0.074	0.004	0.221
		1.000	-0.004	0.015
			1.000	0.022
				1.000

AC(i)) stands for the i-th order autocorrelation coefficient.

---

[28] Based on the Ljung-Box Q-statistic we reject the null hypothesis of zero autocorrelation up to order 5 at 5% significance for each innovation series, except for the SM series.

Figure 5.4 shows the smoothed disturbances of the state space model. A comparison with figure 4.6 confirms the much improved data description once more. Nonetheless, like table 5.3, figure 5.4 still seems to indicate a lack of dynamics.
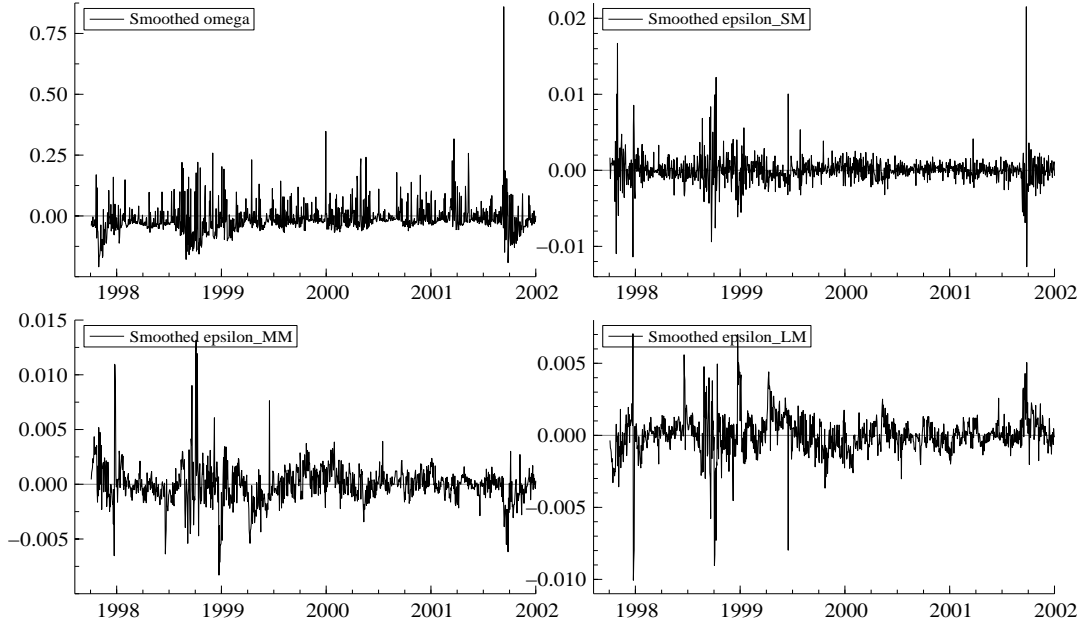

Figure 5.4: Smoothed disturbances of the state space model; $\omega_t, \varepsilon_{SM,t}, \varepsilon_{MM,t}$ and $\varepsilon_{LM,t}$ (2-factor OU, Restrictions (b)).

6.   Estimation results 3-factor OU

In this section we therefore extend to three OU volatility factors. As not all parameters can be identified we consider two sets of a priori imposed restrictions, which are similar as in the 2-factor case. In *Restrictions (a)* we consider the case of three independent factors and do not allow for contemporaneous correlation between $\{\varepsilon_{it}\}, i = SM, MM, LM$. Restrictions (a) imposes $\theta_2 = \theta_3 = 0$, $\boldsymbol{\Sigma} = \mathrm{diag}[\sigma_{11}, \sigma_{22}, \sigma_{33}]$ and $\mathrm{corr}(\varepsilon_{SM,t}, \varepsilon_{MM,t}) = \mathrm{corr}(\varepsilon_{SM,t}, \varepsilon_{LM,t}) = \mathrm{corr}(\varepsilon_{MM,t}, \varepsilon_{LM,t}) = 0$. In *Restrictions (b)* we allow both for correlation between the factors and contemporaneous correlation between $\{\varepsilon_{it}\}, i = SM, MM, LM$. In this case we impose $\theta_2 = \theta_3 = 0$, $\sigma_{12} = \sigma_{13} = \sigma_{23} = 0$ such that $(3x3)\boldsymbol{\Sigma}$ is lower diagonal. Table 6.1 presents the estimation results.

Allowing for three factors instead of two increases the quasi-loglikelihood from more than 14,300 to more than 14,500. Permitting for correlated factors and contemporaneous correlation between the $\{\varepsilon_{it}\}; i = SM, MM, LM$ series does not raise the loglikelihood very much. It does however introduce additional uncertainty in most parameter estimates, as evidenced by the increased standard errors (relative to the estimates). There are clear opportunity costs associated with the increased flexibility obtained from an extra introduction of six additional parameters when going from Restrictions (a) to (b). We feel most confident with

the results of Restrictions (a), the independent-factors case, such that in the remainder we focus on those [29].

Table 6.1: Estimation results 3-factor OU using return, SM, MM and LM option data jointly

	Restrictions (a)	Restrictions (b)		Restrictions (a)	Restrictions (b)
θ_1	**0.0392** (0.0417)	**0.0412** (0.0990)	corr($\varepsilon_{SM,t}, \varepsilon_{MM,t}$)	0	**0.59** (0.47)
θ_2	0	0	corr($\varepsilon_{SM,t}, \varepsilon_{LM,t}$)	0	**0.14** (0.50)
θ_3	0	0	corr($\varepsilon_{MM,t}, \varepsilon_{LM,t}$)	0	**-0.16** (0.84)
k_1	**0.023** (0.029)	**0.167** (0.388)			
k_2	**3.32** (0.29)	**1.22** (3.40)	Vol. returns	19.81%	20.30%
k_3	**15.7** (2.5)	**13.6** (2.98)	Vol-of-var.	0.2000	0.2355
σ_{11}	**0.0402** (0.0036)	**0.171** (0.626)	Std.dev. x_{1t}	0.189	0.295
σ_{12}	0	0	Std.dev. x_{2t}	0.057	0.181
σ_{13}	0	0	Std.dev. x_{3t}	0.031	0.041
σ_{21}	0	**-0.266** (0.679)	corr(x_{1t}, x_{2t})	0	-0.61
σ_{22}	**0.147** (0.013)	**0.097** (0.132)	corr(x_{1t}, x_{3t})	0	0.13
σ_{23}	0	0	corr(x_{2t}, x_{3t})	0	-0.27
σ_{31}	0	**0.129** (0.149)	Persistence x_1	0.9999	0.9994
σ_{32}	0	**0.047** (0.185)	Persistence x_2	0.9873	0.9953
σ_{33}	**0.176** (0.035)	**0.161** (0.037)	Persistence x_3	0.9413	0.9490
γ_1	**-0.034** (0.177)	**-0.165** (0.231)	Half-life x_1	7990	1080
γ_2	**0.266** (0.481)	**0.243** (0.475)	Half-life x_2	54	147
γ_3	**-1.47** (0.77)	**-1.29** (0.73)	Half-life x_3	11	13
$\sigma_{\varepsilon,SM}$	**0.00096** (0.00090)	**0.00274** (0.00151)	Std.dev. u_{1t}	0.0025	0.0106
$\sigma_{\varepsilon,MM}$	**0.00140** (0.00013)	**0.00129** (0.00095)	Std.dev. u_{2t}	0.0090	0.0175
$\sigma_{\varepsilon,LM}$	**0.00182** (0.00024)	**0.00137** (0.00043)	Std.dev. u_{3t}	0.0106	0.0128
σ_{ω}	**0.0669** (0.0059)	**0.0673** (0.0059)	Loglikelihood	14,500	14,560

The table reports (restricted) parameter estimates (boldface) with robust White (1982) QML standard errors in parentheses, resulting from estimating the state space model associated with the 3-factor OU SV assumption using the combination of return, SM, MM and LM option data under two sets of restrictions (see main text), together with the MLEs of some other quantities of interest.

---

[29] The outcome and interpretation of the diagnostic checks only marginally differ for both cases. Moreover, the evidence in Cont and Fonseca (2002) seems to support this choice further.

6.1 Decomposing stock volatility: long-memory, middle-long-term and short-term movements

We start with an examination of the stock volatility. The upper graph of figure 6.1 plots the smoothed stock volatility obtained from the 3-factor model. In the middle graph this series is plotted again, together with the smoothed stock volatility obtained from the 1-factor model estimated using return - SM option data. The lower graph shows that their difference can be substantial. A close inspection reveals that, e.g., in times of sudden rises in the volatility (i.e. around the third quarter of 1997 and September 11, 2001), the volatility extracted from the 3-factor model responds quicker than the volatility obtained from the 1-factor model. This is as expected. The unconditional volatility of returns, $\sqrt{\mathbf{1'\theta}}$, is estimated around 20% as was previously the case.

The 3-factor model models the dynamics of the stock variance as a sum of three hidden processes. The evolution of the stock volatility is thus decomposed into three different dynamic components. Table 6.1 shows that the three volatility-driving factors distilled from the data differ greatly in their features. The first two factors behave more or less similar as the ones obtained from the 2-factor model, although each has further increased persistence. The first factor is extremely persistent and almost behaves like a random walk; it has very long memory. Shocks to this factor have close-to-permanent effects on its future values (which also explains why the MLE of the standard deviation of x_{1t} is much larger than for both other factors). Although the second factor reverts much faster to its mean with a half-life of around 54 days or 2.5 months, the third factor is quickest mean-reverting. It takes about 11 days for a shock in this factor to lose half its impact.
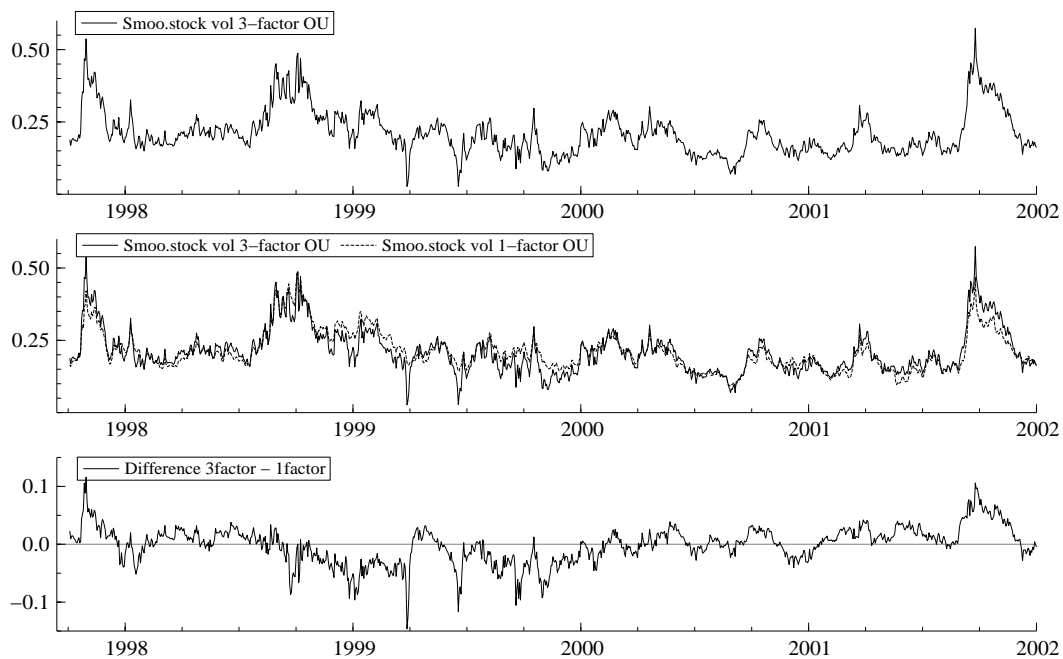

Figure 6.1: The upper graph shows the smoothed stock volatility obtained from 3-factor OU, Restrictions (a). The middle graph plots this series again, together with the smoothed stock volatility obtained from 1-factor OU estimated using return – SM option data. Their difference is plotted in the lower graph.
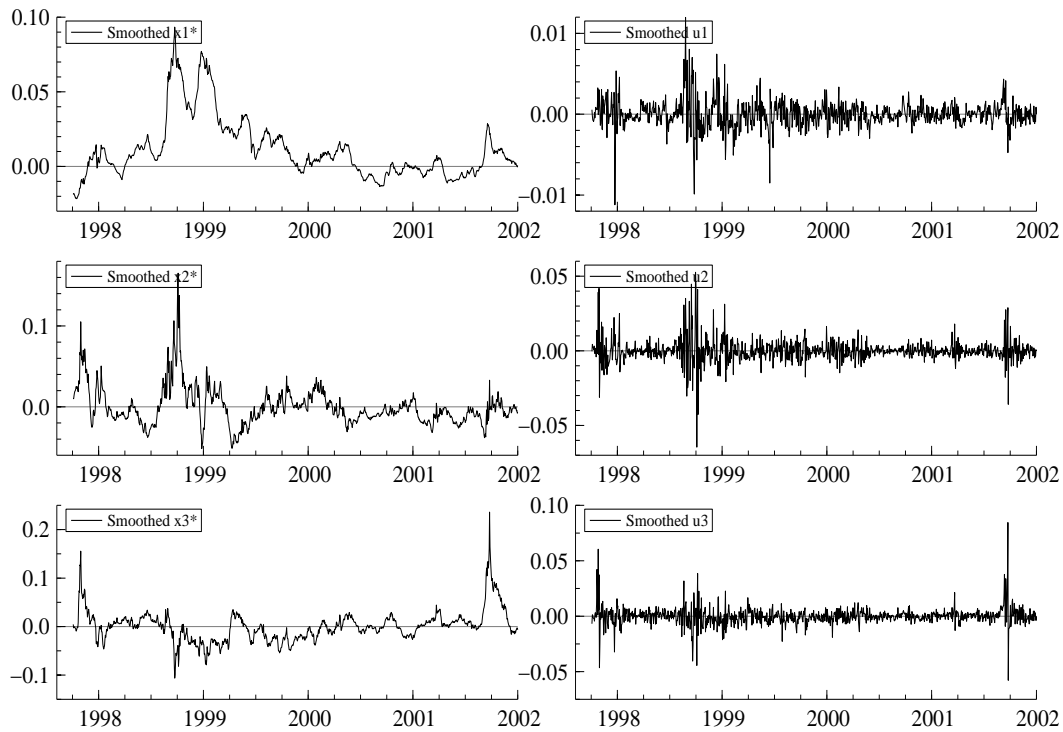
Figure 6.2: Smoothed factors in deviation from their mean x_{1t}^*, x_{2t}^* and x_{3t}^*, and smoothed state equation errors u_{1t}, u_{2t} and u_{3t} (3-factor OU, Restrictions (a)).

The estimates indicate that the first factor may be interpreted as determining the long-term trend in the stock-volatility evolution. The second factor seems more associated with middle-long-term movements. The third factor determines the very short-term volatility fluctuations.

The left panel of figure 6.2 plots the smoothed factors x_i^*. The right panel plots the daily shocks to these factors, i.e., the smoothed disturbances u_i. Notice the scales on the vertical axes, and in particular the graph of u_3. The third factor is clearly responsible for large changes in the volatility in short periods of time. This is confirmed by the fact that its instantaneous standard deviation, σ_{33}, is largest, and the fact that the daily shocks to this factor, u_3, have largest variance [30].

6.2   Another interpretation: Level, slope and curvature factors

Our first interpretation of the hidden factors concerned an interpretation in terms of what each factor contributes to the dynamic evolution of the stock volatility. Our second interpretation concerns the impact of each factor on the prices of options; that is, on the shape and dynamics of the volatility term structure.

Reconsider the measurement equation for the options in the general multifactor model; i.e.,

---

[30] If one is willing to believe that there are true jumps in volatility present in the data (see, e.g., Eraker et al. (2003)), then this (continuous-path) factor is clearly trying to accommodate these jumps.

$$\sigma^2_{implied,t} = \frac{1}{\tau_t}\left[A(\tau_t) + \mathbf{B}(\tau_t)'\boldsymbol{\theta}\right] + \frac{\mathbf{B}(\tau_t)'}{\tau_t}\mathbf{x}_t^* + \varepsilon_t; \; \varepsilon_t \sim (0, \sigma^2_\varepsilon) \qquad (6.1)$$

The functions $B_j(.)/(.); j = 1,..,n$ can be interpreted as reaction coefficients or factor loadings. Each measures the instantaneous ceteris-paribus response of the BS implied variance to a change in one of the factors, i.e., $\partial\sigma^2_{implied,t}/\partial x_{jt} = B_j(\tau_t)/\tau_t$. These reaction coefficients depend on the maturity of the option. Options of different maturity (and hence the volatility term structure) respond differently to changes in each of the volatility-driving factors. To investigate how a certain shock affects each option, it seems natural to take into account the fact that the factors have different variances. Since $\mathrm{var}_\mathbb{P}[\mathbf{x}_t] = \mathbf{J}\odot\boldsymbol{\Sigma\Sigma}'$, the direct impact of a one-standard-deviation shock in each of the factors on the BS implied variance of an option with maturity τ is given by the vector $(nx1)$ $(\mathbf{I}_n \odot \mathbf{J} \odot \boldsymbol{\Sigma\Sigma}')^{1/2}\mathbf{B}(\tau)/\tau$. In the $n-$factor OU case, the factor loadings are given by $B_j(\tau)/\tau = (1 - \exp[-k_j\tau])/k_j\tau$ for $\tau \geq 0$. If the OU factors are independent, then $\mathrm{var}_\mathbb{P}[x_{jt}] = \sigma^2_{jj}/2k_j$.

The upper graph of figure 6.3 shows the reaction coefficients as a function of maturity, as implied by our estimation results. The lower graph displays the instantaneous ceteris-paribus response of the BS implied variances of all maturities to one-standard-deviation shocks to each of the factors. The pictures yield clear insight in how the volatility term structure is affected by such shocks.
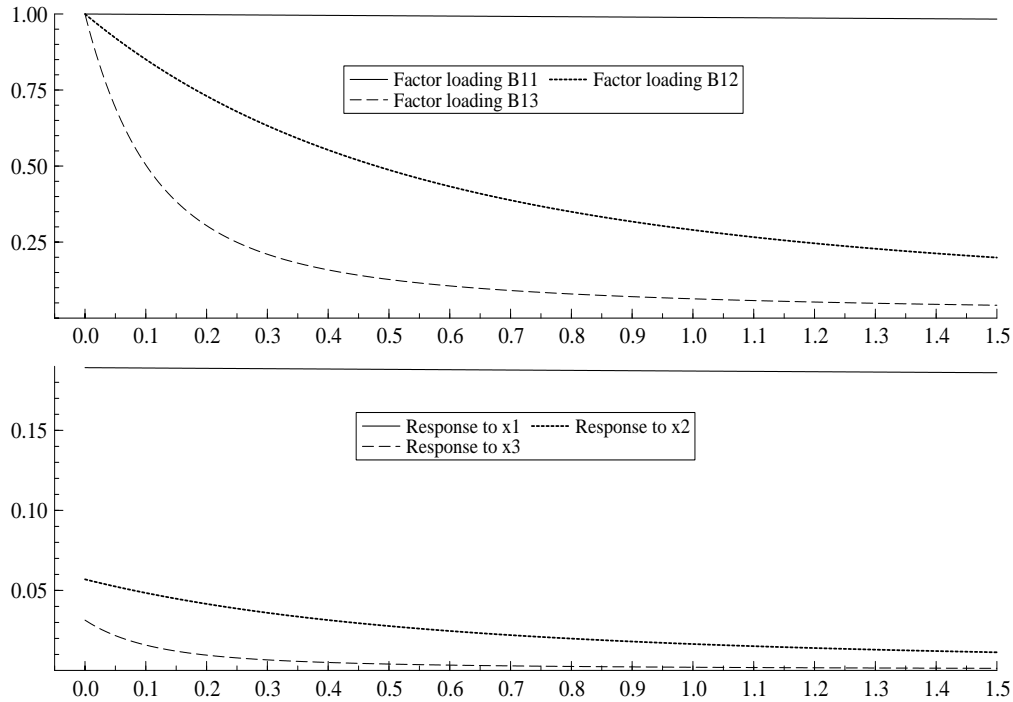

Figure 6.3. Upper plot: Reaction coefficients $B_1(.), B_2(.)$ and $B_3(.)$ as a function of option maturity (in years). Lower plot: Instantaneous responses of the BS implied variances to one-standard-deviation shocks in each of the factors, as a function of maturity.

A shock to the first factor has similar impact on all options, irrespective of maturity. Ceteris paribus, this factor seems to cause parallel shifts in which all implied volatilities increase with approximately the same amount. This is

attributed to the fact that this factor is so highly persistent with inherent shocks that die out very slowly. As such it mainly influences the general level of the volatility term structure. Therefore, the first factor may be interpreted as a *level* factor. A shock to the second factor affects options of all maturities as well, but by different amounts. Short-maturity options respond most, and the response gradually diminishes the longer the maturity of the option. Shocks to the third factor have biggest impact on the implied volatilities (and hence prices) of short-maturity options. The effect quickly wears off as the maturity of the option increases. These shocks do not virtually affect options with a maturity longer than half a year. As this factor is so quickly mean-reverting this makes sense: The longer the maturity, the more often this factor will have reverted back to its mean, the closer its average value over the lifetime of the option will be to this mean. As this average value is a major determinant of the option price (i.e. recall equation (1.6) where $\bar{\sigma}_t^2 = \bar{x}_{1t} + \bar{x}_{2t} + \bar{x}_{3t}$), it is clear that shocks to this factor have virtually no effect on sufficiently long-dated options.
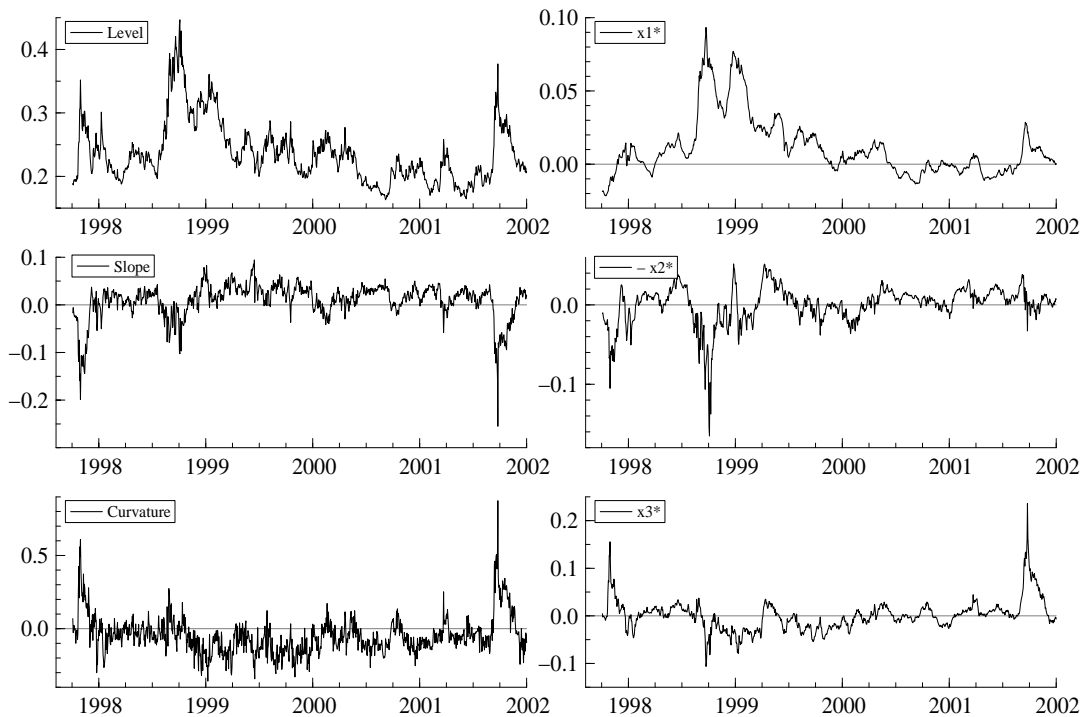

Figure 6.4. Left panel: Evolution of level, slope and curvature of the volatility term structure through time. Right panel: Smoothed x_t^*, $-x_2^*$ and x_3^*.

Level, slope and curvature of the volatility term structure
Given this analysis, and given that the first factor seems mainly associated with fluctuations in the general level, it makes intuitively sense that both the second and third factor will largely be responsible for dynamic changes in the *slope* and *curvature* (or *convexity*) of the volatility term structure over time. Let us examine if this is indeed the case, and what contribution each factor makes to these dimensions of the volatility term structure. First, define the level of the term structure at time t as the average of the SM, MM and LM BS implied volatilities: $level_t \equiv (\sigma_{implied,SM,t} + \sigma_{implied,MM,t} + \sigma_{implied,LM,t})/3$. Next, define the slope as the difference between the LM and SM BS implied volatility, divided by the maturity

difference; i.e. $slope_t \equiv (\sigma_{implied,LM,t} - \sigma_{implied,SM,t})/(\tau_{LM,t} - \tau_{SM,t})$. We expect a negative correlation between the second and third factors, and the slope. Finally, define the curvature of the volatility term structure at time t as

$$curvature_t \equiv a_t\, \sigma_{implied,LM,t} + b_t\, \sigma_{implied,MM,t} + c_t\, \sigma_{implied,SM,t}, \qquad (6.2)$$

where $a_t = 2/[h_{1t}(h_{1t} + h_{2t})]$, $c_t = 2/[h_{2t}(h_{1t} + h_{2t})]$, $b_t = -a_t - c_t$, $h_{1t} = \tau_{LM,t} - \tau_{MM,t}$ and $h_{2t} = \tau_{MM,t} - \tau_{SM,t}$. As explained in the appendix, this curvature measure essentially represents a numerical approximation of the second-order derivative of a smooth function in a certain point computed using three coordinate pairs.

The left panel of figure 6.4 plots the dynamic evolution of the level, slope and curvature of the volatility term structure through time. The right panel shows the smoothed factors x_1^*, $-x_2^*$ and x_3^*. (Notice that we have multiplied the second factor with –1 due to the expected negative relation.) The first factor may indeed be interpreted as mainly being a level factor: the correlation between x_1 and the general level of the term structure equals 0.86. Given that the impact of a shock to the second factor only gradually becomes less the longer the option maturity (recall figure 6.3), it is not surprising that x_{2t} and $level_t$ are also rather correlated; their correlation coefficient equals 0.68. In contrast, the correlation between x_{3t} and $level_t$ is only –0.08. The second factor tracks changes in the slope of the term structure rather well. As expected, the slope is negatively correlated with x_2: their correlation equals –0.60. In contrast, the correlation between x_{1t} and $slope_t$ is only –0.02. Not surprisingly, the correlation between x_{3t} and our slope measure is also rather large, -0.62. However, the third factor even more closely follows the movements of the curvature of the volatility term structure. The correlation between $curvature_t$ and x_{3t} is substantial, 0.82, whereas the correlation between $curvature_t$ and x_{1t} (resp. x_{2t}) is only –0.07 (resp. 0.19).

We may conclude the following. The first, most persistent factor is mainly responsible for the dynamics of the general level of the volatility term structure. The dynamics of the second (but also the third) factor are largely associated with changes in the slope. Dynamic changes in the curvature of the term structure are mainly driven by the third, quickest mean-reverting factor [31] [32].

6.3   Fit of the volatility term structure

We next consider the in-sample fit of the ATM volatility term structure obtained by allowing for three OU SV factors. Figure 6.5 shows the original and fitted implied volatilities together with their difference, the "pricing" errors. Compare this figure to figures 4.7 and 5.1 for the 1 and 2-factor models respectively, and notice the subsequent improvement. In contrast to the earlier pictures, there does not seem any systematic pattern left in the error series.

---

[31] It should be noted that analogical findings are found in empirical implementations of affine models of the term structure of interest rates (Duffie and Kan (1996), Dai and Singleton (2000)). Recall that the term structure of interest rates plots the yield on default-free zero-coupon bonds as a function of the maturity of the bonds. Empirical studies towards these models show that more than one short-interest-rate-driving factor is necessary to obtain a good fit of empirical zero-coupon bond data. Three (correlated) factors seem needed; see e.g. de Jong (2000), Dai and Singleton (2000) and Andersen and Lund (1997). These factors have similar interpretations as in our stock-option-pricing setting under SV.
[32] Cont and Fonseca (2002) attach a level, slope and curvature interpretation to the shapes of the three eigenmodes (i.e., "principal component surfaces") extracted from the daily implied volatility-surface fluctuations.
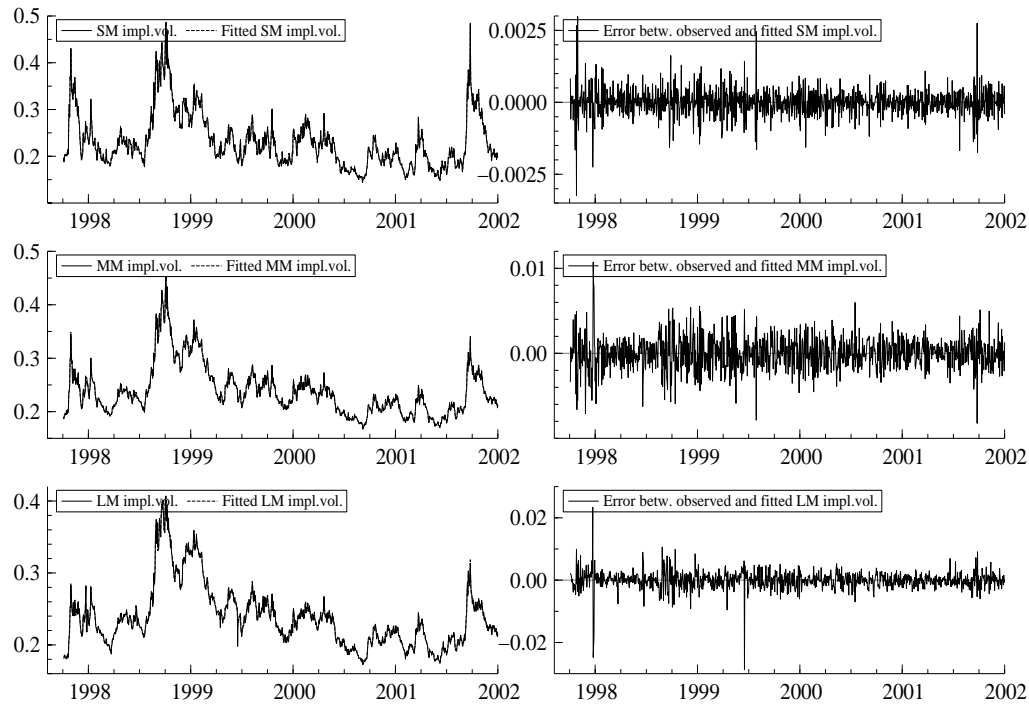
Figure 6.5: In-sample fit of the ATM volatility term structure (3-factor OU, Restrictions (a)). Left: original and fitted BS implied volatilities for each maturity. Right: their difference (i.e. "pricing" error).

The improvement in fit is also apparent from table 6.2. All pricing errors concentrate around zero and are typically small. The average absolute errors seem small as well. The SM option series is now fitted best. If we subsequently compare the mean absolute errors in table 4.5 with table 5.2 and table 6.2, the observed pattern is logically understood given our previous analysis and interpretation of the factors.

Table 6.2: Summary statistics error series (3-factor OU, Restrictions (a))

	Error SM	Error MM	Error LM	|Error SM|	|Error MM|	|Error LM|
Mean	0.00%	0.00%	0.00%	0.03%	0.15%	0.19%
Std.dev.	0.05%	0.19%	0.29%	0.03%	0.13%	0.22%

The table reports the mean and standard deviation of the error series shown in figure 6.5. Both statistics are also reported for the absolute errors.

Fitting the 1-factor model to the joint data results in distillation of the overall volatility movement from the data; i.e., the long-memory or volatility-term-structure level factor. As this factor is so highly persistent it cannot respond quickly enough to sudden changes in short periods of time, which (given the knowledge from figure 6.3 which indicates that it is especially SM derivatives that are prone to these changes), results in the SM series fitted worst. Not surprisingly, the LM series if fitted best: Intuitively, short-term fluctuations tend to "average out" the further one looks into the future.

Moving from one factor to two factors yields the biggest improvements in fit for the MM and especially the SM option series. The second factor in the 2-factor model is much quicker mean-reverting, and therefore more easily picks up

sudden changes. The increase in fit of the LM series was rather modest, which may again be explained by the averaging-out effect. Going from two to three factors resulted in distillation of an additional, very fast mean-reverting factor. It will be clear why we observe the biggest improvement for the SM series. Taking the averaging-out effect into account once more, it also makes sense why the fit of the LM series remained virtually constant when going from two to three factors.

6.4   Diagnostic checking

Figure 6.6 shows the standardized innovation series obtained from the estimated 3-factor model. Compare this figure to figures 4.5 and 5.3 from the 1 and 2-factor models respectively. The innovations seem to resemble white noise closest in the 3-factor case. However, conditional heteroskedasticity is still present in all series with biggest fluctuations in times of largest volatility (recall figure 6.1). This may well be attributed to the fact that OU stochastic volatility does not model level-dependent volatility-of-volatility. Other evidence that supports the presence of volatility feedback is in figure 6.2: Fluctuations in the smoothed disturbances u_j are generally largest when the (absolute) levels of the volatility factors x_j are largest. If the model were correctly specified these smoothed disturbance series ought to (more or less) look like white noise instead.
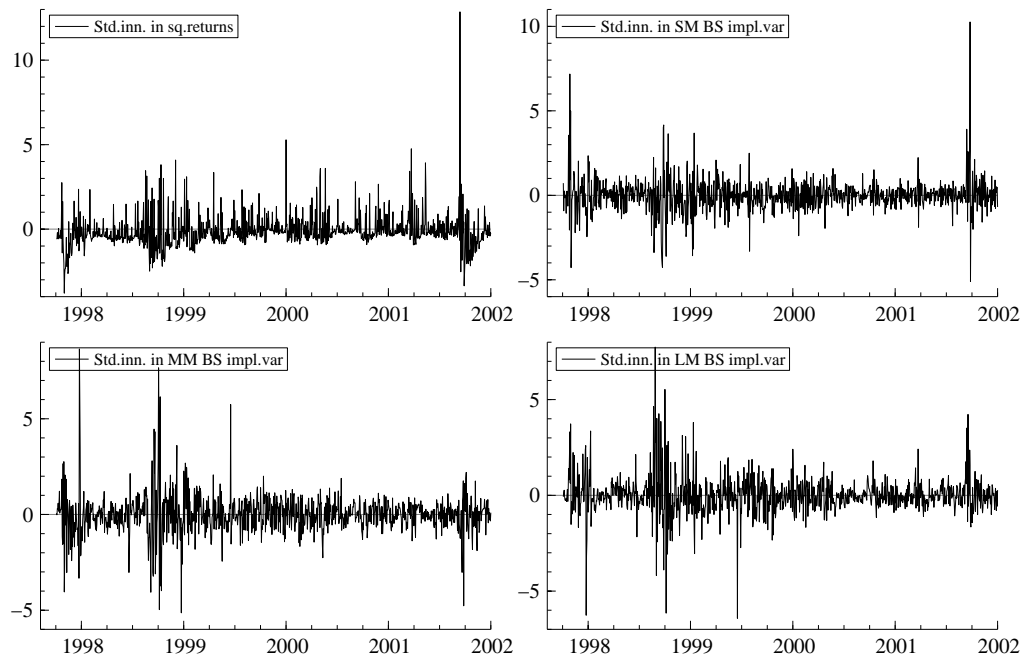

Figure 6.6: Standardized innovations (3-factor OU, Restrictions (a)),

Table 6.3 provides the mean, standard deviation and autocorrelation coefficients of each innovation series. Allowing for three factors removes most autocorrelation from the innovations. (Compare to tables 4.4 and 5.3.) Notice the substantial decrease in correlation in the MM series as compared to the 2-factor case. Although most autocorrelations are close to zero, the Ljung-Box Q-test still rejects the null hypothesis of no serial correlation up to order 5 for all but the SM series, with respective p-values of 0.00, 0.40, 0.00 and 0.00. Keep in mind

however that a proper application of this test requires homoskedasticity of the underlying series, and this is clearly violated.

Table 6.3: Summary statistics standardized innovations (3-factor OU Restrictions (a))

	Std.inn. sq. return	Std.inn. SM series	Std.inn. MM series	Std.inn. LM series
Mean	-0.132	-0.009	-0.001	0.003
Std.deviation	0.995	0.968	1.043	0.985
AC(1)	0.099	0.043	0.130	0.086
AC(2)	0.064	-0.006	-0.016	0.015
AC(3)	0.141	-0.033	-0.044	-0.101
AC(4)	0.114	-0.041	-0.037	0.063
AC(5)	0.053	-0.013	0.031	0.022
Cont.correlation matrix	1.000	0.085	0.022	0.195
		1.000	-0.064	0.098
			1.000	-0.026
				1.000

AC(i) stands for the i-th order autocorrelation coefficient.

Figure 6.7 draws the smoothed disturbances $\omega_t, \varepsilon_{SM,t}, \varepsilon_{MM,t}$ and $\varepsilon_{LM,t}$ of the state space model. Again, their sample paths resemble white-noise sample paths more closely than in the 1 and 2-factor cases. Nonetheless, also in these pictures some effect of neglected level-dependent volatility-of-volatility seems present.

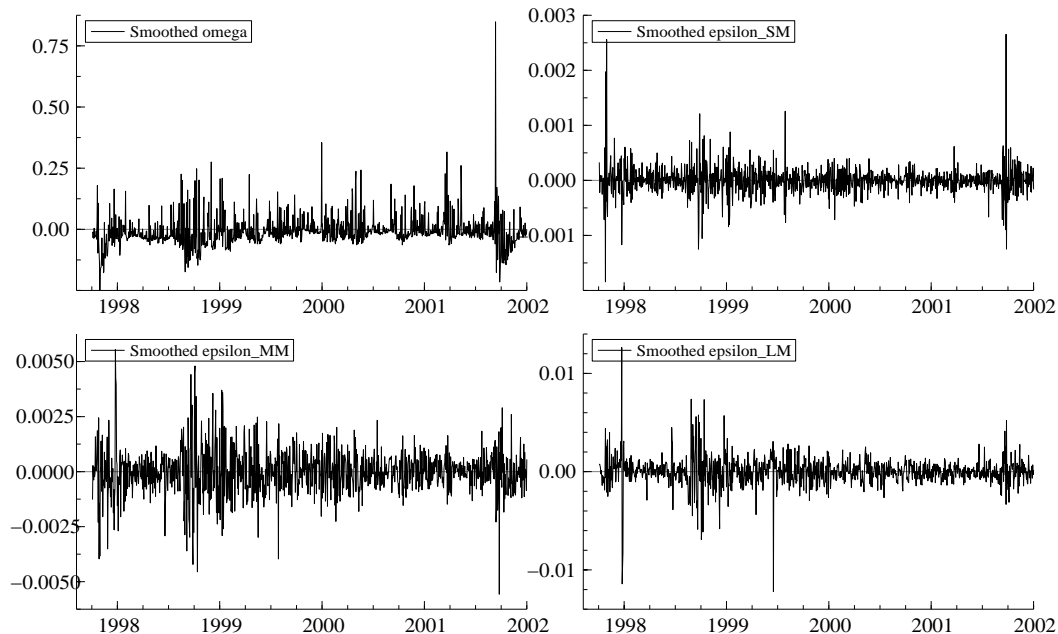Figure 6.7: Smoothed disturbances of the state space model; $\omega_t, \varepsilon_{SM,t}, \varepsilon_{MM,t}$ and $\varepsilon_{LM,t}$ (3-factor OU, Restrictions (a)).

Let us finally examine the possible presence of a leverage effect in the FTSE100-index data. Recall that the model outlined in section 1 does not account for it. If we compute the correlation between daily FTSE100-index returns and daily

smoothed stock-variance changes (obtained from the 3-factor model), we find a correlation of −0.6. The presence of the leverage effect seems evident. Nevertheless, given that our analysis concerns the maturity dimension of ATM options (and hence does not regard the cross-sectional option dimension), it seems unlikely that this misspecification invalidates most of our findings.

7.   Summary and directions for further research

In this paper we provide further empirical evidence on the existence of multiple stock volatility-driving factors. Moreover, we attach interpretations to these factors; both in terms of their contribution to the evolution of the stock volatility and their impact on option prices, and to the dynamics of the volatility term structure. Using the estimation strategy of multifactor affine SV option pricing models considered in Van der Ploeg et al. (2003), we find that the volatility dynamics present in daily FTSE100 stock-index and option data can satisfactorily be explained by three factors.

An explorative analysis of the data shows that both the level and fluctuations in the FTSE100-index volatility have substantially increased from the third quarter of 1997 onwards. To mitigate misspecification, we therefore only use data for Oct 97 - Dec 2001 in the estimations. An examination of the term structure of ATM implied volatilities shows that it can have very different shapes. Moreover, it evolves over time in complex manners. 34% of all shifts is non-parallel. A realistic model must obviously be able to reproduce these complex dynamics.

We start off with assuming that there is only one SV factor that drives the volatility which follows an OU process. Although the in-sample fit of the 1-factor model estimated using return-SM option data seems reasonable, the out-of-sample pricing performance of longer-dated options is bad: The model severely overfits MM and LM option prices. Similar results for the Heston (1993) model are known in the literature. Estimating the 1-factor model using stock returns and data on a SM, MM and LM option series jointly, indicates that the main misspecification of the 1-factor model is that it neglects important dynamics. The 1-factor model is not capable of describing the observed rich dynamics in the joint data. The volatility term structure is fitted worst at the short end, and best at the long end.

These results naturally motivate an empirical investigation of multifactor volatility models. Although allowing for two factors yields a considerable increase in fit, dynamic misspecification still remains. The average fit of the three option series is now similar, with the LM series fitted somewhat best.

Extending to three independent OU factors seems sufficient to satisfactorily describe the dynamics observed in the four joint time series. Comparing the smoothed stock volatilities obtained from the 1-factor model (estimated using return-SM option data) with the volatilities obtained from the 3-factor model, reveals that especially in times of sudden increases in financial-markets uncertainty, the latter ones react quicker. This makes sense: decomposing the volatility in different dynamic components allows the volatility to respond quicker.

The three factors distilled from the data differ greatly in their characteristics. The first factor is extremely persistent. It has a daily persistence of 0.9999 with

shocks that have close-to permanent effects on its future evolution; they seem to last for years. The second factor is much quicker mean-reverting (persistence 0.9873) with shocks that have a half-life of about 2.5 months. The third factor is very fast mean-reverting (persistence 0.9413). Shocks to this factor lose half their impact in 11 days, and have moreover largest variance. The third factor clearly governs large volatility changes in relatively brief periods of time.

We interpret the factors in two ways. The first interpretation is that the first, second, and third factor determine the long-term, middle-long-term and short-term trends in the evolution of the stock volatility respectively. As such, they influence the prices of options in different ways. We show that the first, long-memory factor similarly impacts on all options, irrespective of maturity. The second factor affects all options as well, but gradually by smaller amounts the longer the life-time of the option. Short-maturity options are especially prone to shocks in the third factor. The impact of this factor quickly diminishes as the option maturity increases. As this factor is so quickly mean-reverting, its shocks tend to average out over a sufficiently long time span, resulting in an only marginal impact on long-dated options.

Our second interpretation of the factors concerns their impact on the shape and dynamics of the volatility term structure. We show that the first factor is mainly responsible for changes in the general level of the term structure. The second factor (but also the third) factor is largely associated with changes in the slope. The third factor is surprisingly closely related to dynamic changes in the convexity of the volatility term structure.

The fit of the volatility term structure obtained by allowing for three factors is good. The pricing errors concentrate around zero, are typically small and there does not seem any systematic pattern left in the errors over time. The SM series is fitted somewhat best. A subsequent comparison of the pricing errors from the 1-factor model with the errors of the 2-factor and 3-factor models yields a pattern that is logically understood given our previous analysis.

Specification tests of the 3-factor OU SV model reveal that the dynamics observed in the data are satisfactorily captured by the model. Nonetheless and not surprisingly, we also find evidence of level-dependent volatility-of-volatility and the leverage effect. These features of the data are not modeled by the multifactor OU SV model considered in this paper. In future research we aim at explicitly allowing for e.g. volatility feedback in the SV-driving processes. We then resort to a related but different estimation method: Extended Kalman filter QML.

Nonetheless, despite these deficiencies, our findings are clearly supportive and confirmative of the existence of multiple volatility-driving factors in practice.

Appendix

I. System matrices state space model for multifactor OU SV

The state space model associated with the $n-$factor affine OU SV derivative pricing model in which both the squared returns and the three option series are jointly used for estimation reads:

Measurement equation:

$$\boldsymbol{y}_t = \boldsymbol{a}_t + \boldsymbol{H}_t{'}\boldsymbol{\xi}_t + \boldsymbol{w}_t \;; \qquad \boldsymbol{w}_t \sim (\boldsymbol{0}, \boldsymbol{R}) \qquad t = \Delta t, \ldots, T\Delta t \qquad (\text{I.1})$$

Transition equation

$$\boldsymbol{\xi}_{t+\Delta t} = \boldsymbol{d} + \boldsymbol{F}\boldsymbol{\xi}_t + \boldsymbol{v}_{t+\Delta t}\;; \qquad \boldsymbol{v}_{t+\Delta t} \sim (\boldsymbol{0}, \boldsymbol{Q})\,. \qquad (\text{I.2})$$

Here the error series $\{\boldsymbol{w}_t\}$ and $\{\boldsymbol{v}_{t+\Delta t}\}$ are both serially and mutually uncorrelated at all points in time. The system matrices read [33]

$$(4x1)\,\boldsymbol{y}_t \equiv \begin{bmatrix} \frac{1}{\Delta t}(r_t - \hat{\mu}_{t-\Delta t}\Delta t)^2 \\ \sigma^2_{implied,1t} \\ \sigma^2_{implied,2t} \\ \sigma^2_{implied,3t} \end{bmatrix}, \qquad (4x1)\,\boldsymbol{a}_t \equiv \begin{bmatrix} 0 \\ \frac{1}{\tau_{1t}}[A(\tau_{1t}) + \mathbf{B}(\tau_{1t}){'}\boldsymbol{\theta}] \\ \frac{1}{\tau_{2t}}[A(\tau_{2t}) + \mathbf{B}(\tau_{2t}){'}\boldsymbol{\theta}] \\ \frac{1}{\tau_{3t}}[A(\tau_{3t}) + \mathbf{B}(\tau_{3t}){'}\boldsymbol{\theta}] \end{bmatrix} \qquad (\text{I.3})$$

$$(4x\,n+1)\,\boldsymbol{H}_t{'} \equiv \begin{bmatrix} 1 & \mathbf{0}{'} \\ 0 & \frac{1}{\tau_{1t}}\mathbf{B}(\tau_{1t}){'} \\ 0 & \frac{1}{\tau_{2t}}\mathbf{B}(\tau_{2t}){'} \\ 0 & \frac{1}{\tau_{3t}}\mathbf{B}(\tau_{3t}){'} \end{bmatrix}, \qquad (n+1\,x1)\,\boldsymbol{\xi}_t \equiv \begin{bmatrix} \frac{1}{\Delta t}(r_t - \hat{\mu}_{t-\Delta t}\Delta t)^2 \\ \mathbf{x}_t^* \end{bmatrix},$$

$$(4x1)\,\boldsymbol{w}_t \equiv \begin{bmatrix} 0 \\ \varepsilon_{1t} \\ \varepsilon_{2t} \\ \varepsilon_{3t} \end{bmatrix}, \qquad (4x4)\,\boldsymbol{R} \equiv \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & \sigma^2_{\varepsilon 1} & \rho_{12}\sigma_{\varepsilon 1}\sigma_{\varepsilon 2} & \rho_{13}\sigma_{\varepsilon 1}\sigma_{\varepsilon 3} \\ 0 & \rho_{12}\sigma_{\varepsilon 1}\sigma_{\varepsilon 2} & \sigma^2_{\varepsilon 2} & \rho_{23}\sigma_{\varepsilon 2}\sigma_{\varepsilon 3} \\ 0 & \rho_{13}\sigma_{\varepsilon 1}\sigma_{\varepsilon 3} & \rho_{23}\sigma_{\varepsilon 2}\sigma_{\varepsilon 3} & \sigma^2_{\varepsilon 3} \end{bmatrix},$$

$$(n+1\,x1)\,\boldsymbol{d} \equiv \begin{bmatrix} \mathbf{1}{'}\boldsymbol{\theta} \\ \mathbf{0} \end{bmatrix}, \qquad (n+1\,x\,n+1)\,\boldsymbol{F} \equiv \begin{bmatrix} 0 & \mathbf{1}{'} \\ \mathbf{0} & \exp[-\mathbf{K}_d\Delta t] \end{bmatrix},$$

$$(n+1\,x1)\,\boldsymbol{v}_{t+\Delta t} \equiv \begin{bmatrix} \omega_{t+\Delta t} \\ \mathbf{u}_{t+\Delta t} \end{bmatrix}, \qquad (n+1\,x\,n+1)\,\boldsymbol{Q} \equiv \begin{bmatrix} \sigma^2_\omega & \mathbf{0}{'} \\ \mathbf{0} & \mathbf{G}(\Delta t)\odot\boldsymbol{\Sigma\Sigma}{'} \end{bmatrix}.$$

The matrix $(nxn)\,\mathbf{G}(\Delta t)$ has $ij-$th element equal to $[\mathbf{G}(\Delta t)]_{ij} = (1 - \exp[-(k_i + k_j)\Delta t])/(k_i + k_j)$. The state space model is estimated using Kalman filter QML . For more details see Van der Ploeg et al. (2003).

---

[33] For notational convenience we index the option series with 1, 2 and 3 instead of SM, MM, and LM.

II.   Curvature measure of the volatility term structure

In this appendix we provide a rationale for our convexity or curvature measure (6.2) of the volatility term structure. Recall first that the volatility term structure plots the BS implied volatilities as a function of maturity, and that on each day we observe *three* BS implied volatilities: SM, MM and LM. The way these volatilities are located with respect to each other determines the shape and hence curvature of the volatility term structure. (Think about 3 points on the graph of a smooth function.)

Now, consider some at least twice differentiable function $f(.)$. Our aim is to obtain an approximation of the curvature of this function in a certain point y, given only the three coordinate-pairs $[y + h_1, f(y + h_1)]$, $[y, f(y)]$, and $[y - h_2, f(y - h_2)]$, with $h_1, h_2 > 0$. The second-order derivative of $f(.)$ evaluated in the point y, $f''(y)$, determines this curvature.

Our goal is then to determine the coefficients a, b and c such that

$$f''(y) \approx a f(y + h_1) + b f(y) + c f(y - h_2). \tag{II.1}$$

Second-order Taylor series approximations lead us to

$$f(y + h_1) \approx f(y) + h_1 f'(y) + \tfrac{1}{2} h_1^2 f''(y) \tag{II.2}$$

$$f(y - h_2) \approx f(y) - h_2 f'(y) + \tfrac{1}{2} h_2^2 f''(y). \tag{II.3}$$

Substitution into (II.1) followed by rewriting gives

$$f''(y) \approx (a + b + c) f(y) + (a h_1 - c h_2) f'(y) + \tfrac{1}{2}(a h_1^2 + c h_2^2) f''(y). \tag{II.4}$$

This leads us to the following restrictions

$$a + b + c \equiv 0; \; a h_1 - c h_2 \equiv 0; \; \tfrac{1}{2}(a h_1^2 + c h_2^2) \equiv 1, \tag{II.5}$$

which yield

$$a = \frac{2}{h_1(h_1 + h_2)}; \; c = \frac{2}{h_2(h_1 + h_2)}; \; b = -a - c. \tag{II.6}$$

These coefficients together with (II.1) determine the approximate curvature in the point $[y, f(y)]$. Notice in particular that if $f(y + h_1) = f(y) = f(y - h_2)$ such that the function is flat, then the curvature is zero, which is as expected. Notice moreover that in case of $h_1 = h_2 = h$, the convexity approximation reduces to the commonly-known formula

$$f''(y) \approx \frac{f(y + h) - 2f(y) + f(y - h)}{h^2}. \tag{II.7}$$

In our application the role of the function $f(.)$ is played by the BS implied volatilities seen as a function of maturity; i.e., by the volatility term structure. The maturity of the MM series $\tau_{MM,t}$ is associated with y, h_1 equals the maturity difference between the LM and MM series, and h_2 equals the maturity difference between MM and SM series. As these differences are not equal (and moreover vary over time as well) we approximate the curvature of the volatility term structure by equations (II.1) and (II.6) leading to equation (6.2) in the main text.

References

Andersen, T. G., L. Benzoni and J. Lund, 2002, An empirical investigation of continuous-time equity return models, Journal of Finance, vol. LVII, no. 3, 1239-1284.

Andersen, T.G. and J. Lund, 1997, Estimating continuous-time stochastic volatility models of the short-term interest rate, Journal of Econometrics, 77, 343-377.

Alizadeh, S., M.W. Brandt, and F.X. Diebold, 2002, Range-based estimation of stochastic volatility models, Journal of Finance, vol. LVII, no.3, 1047-1091.

Bakshi, G. and N. Kapadia, 2003, Delta-hedged gains and the negative market volatility risk premium, Review of Financial Studies, vol.16, no.2, 527-566.

Barndorff-Nielsen, O.E. and N. Shephard, 2001, Non-Gaussian Ornstein-Uhlenbeck–based models and some of their uses in financial economics, Journal of the Royal Statistical Society, Series B, 63, 167-241.

Barndorff-Nielsen, O.E. and N. Shephard, 2002, Econometric analysis of realized volatility and its use in estimating stochastic volatility models, Journal of the Royal Statistical Society, Series B, 64, 253-280.

Bates, D.S., 1996b, Jumps and stochastic volatility: exchange rate processes implicit in Deutsche Mark options, Review of Financial Studies, vol.9, no.1.

Bates, D. S., 2000, Post-'87 crash fears in S&P500 futures options, Journal of Econometrics 94 (1-2), 181 – 238.

Black, F., 1976, Studies of stock price volatility changes, Proceedings of the 1976 Meetings of the Business and Economic Statistics Section, American Statistical Association, 177-181

Buraschi, A. and J. Jackwerth, 2001, The price of a smile: hedging and spanning in option markets, Review of Financial Studies, vol.14, no. 2, 495-527.

Campbell, J.Y., A.W. Lo and A.C. MacKinlay, 1997, The econometrics of financial markets (Princeton University Press, Princeton NJ).

Carr, P. and L. Wu, 2004, Variance risk premia, Working paper, Bloomberg L.P. And Courant Institute, Zicklin School of Business, New York.

Chernov, M., A.R. Gallant, E. Ghysels and G. Tauchen, 2003, Alternative models for stock price dynamics, Journal of Econometrics vol. 116, issues 1-2, 225-257.

Chernov, M. and E. Ghysels, 2000, A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation, Journal of Financial Economics 56, 407 – 458.

Cont, R. and J. Fonseca (2002), Dynamics of implied volatility surfaces, Quantitative Finance, vol.2, 45-60.

Coval, J.D. and T. Shumway, 2001, Expected option returns, Journal of Finance, vol. LVI, no.3, 983-1009.

Dai, Q. and K.J. Singleton, 2000, Specification analysis of affine term structure models, The Journal of Finance, Vol. LV, no. 5, 1943 – 1978.

Driessen, J. and P. Maenhout, 2003, A portfolio perspective of option pricing anomalies, Working paper, University of Amsterdam and INSEAD (France).

Duffee, G. R. and R. H. Stanton, 2001, Estimation of dynamic term structure models, Working paper, Haas School of Business, U.C. Berkeley, California, USA.

Duffie, D. and R. Kan, 1996, A yield-factor model of interest rates, Mathematical Finance 6, 379 – 406.

Duffie, D., J. Pan and K. Singleton, 2000, Transform analysis and asset pricing for affine jump diffusions, Econometrica 68, 1343-1376.

Durbin, J. and S.J. Koopman, 2001, Time series analysis by state space methods, (Oxford University Press, Oxford, UK).

Engle, R.F. and G.G.J. Lee, 1999, A permanent and transitory component model of stock return volatility, in R.F. Engle and H. White (eds.), Cointegration, Causality and Forecasting: A Festschrift in Honour of Clive W.J. Granger, 475-497, Oxford University Press, Oxford UK.

Eraker, B., M. Johannes and N. Polson, 2003, The impact of jumps in volatility and returns, Journal of Finance, vol. LVIII, no.3, 1269-1299.

Gallant, A.R., C. Hsu and G. Tauchen, 1999, Using daily range data to calibrate volatility diffusions and extract the forward integrated variance, Review of Economics and Statistics, 81, vol.4, 617-631.

Gallant, A.R. and G. Tauchen, 1996, Which moments to match?, Econometric Theory 12, 657 – 681.

Hamilton, J. D., 1994, Time Series Analysis (Princeton University Press, Princeton, New Jersey).

Harvey, A.C., 1989, Forecasting, structural time series models and the Kalman filter, Cambridge University Press, Cambridge, UK.

Harvey, A.C., E. Ruiz and N. Shephard, 1994, Multivariate stochastic variance models, Review of Economic Studies, 61, 247 – 264.

Heston, S.L., 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency options, The Review of Financial Studies 6, 327-43.

Ho, T. and S. Lee, 1986, Term structure movements and pricing interest rate contingent claims, Journal of Finance 41, 1011-1029.

Hull, J.C., 2003, Options, Futures and Other Derivatives, Prentice-Hall International.

Hull, J.C. and A. White, 1987, The pricing of options on assets with stochastic volatilities, Journal of Finance 42, 2, 281-300.

Jones, C.S., 2003, The dynamics of stochastic volatility: evidence from underlying and options markets, Journal of Econometrics 116, 181-224.

de Jong, F., 2000, Time series and cross-section information in affine term-structure models, Journal of Business and Economic Statistics, Vol. 18, No. 3, 300 – 314.

Kim, S., Shephard, N. and S. Chib, 1998, Stochastic volatility: likelihood inference and comparison with ARCH models, Review of Economic Studies 65, 361-93.

Langetieg, T., 1980, A multivariate model of the term structure of interest rates, Journal of Finance, 35, 71-97.

Lund, J., 1997, Econometric analysis of continuous-time arbitrage-free models of the term-structure of interest rates, Working paper, Aarhus School of Business (Aahrus, Denmark).

Meddahi, N., 2002, Moments of continuous time stochastic volatility models, Working paper, CIRANO

Nicolato, E. and E. Vernardos (2002), Option pricing in stochastic volatility models of the Ornstein-Uhlenbeck type, working paper, Centre for Analytical Finance, University of Aarhus, Denmark.

Pan, J., 2002, The jump-risk premia implicit in options: evidence from an integrated time-series study, Journal of Financial Economics 63, 3-50.

Renault, E., 1997, Econometric models of option pricing errors, in Advances in Economics and Econometrics: Theory and Applications, Vol. 3, Kreps, D. and Wallis, K. eds., Cambridge University Press.

Sandmann, G. and S.J. Koopman, 1998, Estimation of stochastic volatility models via Monte Carlo maximum likelihood, Journal of Econometrics 87, 271-301.

Tauchen, G., 2004, Remarks on recent developments in stochastic volatility: Statistical modelling and general equilibrium analysis, Paper prepared for the joint 59[th] European Meeting of the Econometric Society and 19[th] annual congress of the European Economics Association, Universidad Carlos III de Madrid.

Van der Ploeg, A.P.C., H.P. Boswijk and F. de Jong, 2003, A state space approach to the estimation of multifactor affine stochastic volatility option pricing models, Working paper, UvA Econometrics, University of Amsterdam, Holland. (See also the proceedings of the conference *New Frontiers in Financial Volatility Modelling*, May 25 –27, 2003, Florence, Italy.

Van der Ploeg, A.P.C., 2003, A detailed investigation of the statistical properties of the multifactor affine stochastic volatility process, PhD thesis (in progress), UvA Econometrics, University of Amsterdam.

Van der Ploeg, A.P.C., 2004a, A Monte-Carlo study towards the 1-factor OU SV option pricing model, PhD thesis (in progress), UvA Econometrics, University Of Amsterdam.

Van der Ploeg, A.P.C., 2004b, Empirical results for multifactor OU SV: Examining UK financial markets, PhD thesis (in progress), UvA Econometrics, University of Amsterdam.

Vasicek, O.A., 1977, An equilibrium characterization of the term structure, Journal of Financial Economics 5, 177-88.