Finance Stochast. 8, 451–477 (2004)
DOI: 10.1007/s00780-004-0126-7

Finance and Stochastics
© Springer-Verlag 2004

Maturity cycles in implied volatility

Jean-Pierre Fouque[1,*], **George Papanicolaou**[2], **Ronnie Sircar**[3,**], **Knut Solna**[4]

[1] Department of Mathematics, NC State University, Raleigh NC 27695-8205, USA (e-mail: fouque@math.ncsu.edu)
[2] Department of Mathematics, Stanford University, Stanford CA 94305, USA (e-mail: papanico@math.stanford.edu)
[3] Department of Operations Research and Financial Engineering, Princeton University, E-Quad, Princeton, NJ 08544, USA (e-mail: sircar@princeton.edu)
[4] Department of Mathematics, University of California, Irvine CA 92697, USA (e-mail: ksolna@math.uci.edu)

Abstract. The skew effect in market implied volatility can be reproduced by option pricing theory based on stochastic volatility models for the price of the underlying asset. Here we study the performance of the calibration of the S&P 500 implied volatility surface using the asymptotic pricing theory under fast mean-reverting stochastic volatility described in [8]. The time-variation of the fitted skew-slope parameter shows a periodic behaviour that depends on the option *maturity dates* in the future, which are known in advance. By extending the mathematical analysis to incorporate model parameters which are time-varying, we show this behaviour can be explained in a manner consistent with a large model class for the underlying price dynamics with time-periodic volatility coefficients.

Key words: Implied volatilities, maturity cycles, fast mean-reverting stochastic volatility, asymptotic expansions

JEL Classification: C13, G13

Mathematics Subject Classification (2000): 91B70, 60F05, 60H30

1 Introduction

The problem of calibrating market implied volatilities is often described as having three dimensions : K, T and t. The first refers to finding a model (or class of

We thank a referee for his/her comments which improved the paper.
* Work partially supported by NSF grant DMS-0071744.
** Work supported by NSF grant DMS-0090067. We are grateful to Peter Thurston for research assistance.

Manuscript received: August 2002; final version received: December 2003

models) for the underlying asset dynamics that captures the variation of implied volatilities as functions of strike (the smile, or skew in equities and indices); the second refers to the term-structure, or variation over maturities of the options; the small t refers to the variation of the surface over time: if this is not captured well by the model, it leads to instability over time of calibrated parameters. Obtaining stability is crucial for pricing path-dependent options and for hedging.

The first problem can be addressed by many types of models that have fat, and possibly asymmetric tails in returns distributions, for example stochastic volatility models or models with jumps. The other two problems are more challenging. Many approaches deal with these problems successively, trying to obtain a tight in-sample fit[1] by allowing many degrees of freedom, without regard to the underlying price dynamics. The cost is poor stability properties. In previous work [8], we have studied an approach that approximately captures the K and T problems and has fairly good stability properties. We discuss those issues in detail in this work, and introduce a simple extension which improves its performance in all three dimensions.

The calibration procedure we propose here is simply to fit the surface to a straight line in a composite variable (called the log-moneyness-to-*effective*-time-to-maturity), but despite its simplicity, it is consistent with (and derived from) popular mean-reverting models of random volatility, and their associated *no arbitrage* restrictions. When applied to S&P 500 data (across both strikes *and* maturities), the in-sample fit to the surface is fairly good for maturities on the order of a few months, as shown, for example, in Fig. 9. More importantly the variation of the estimated parameters over time is relatively small, as shown in Fig. 11, indicating that the procedure (and underlying models) have good stability properties.

The fast mean-reversion (FMR) asymptotic theory developed in [8] is designed to capture some of the important effects that fluctuations in the volatility have on derivative prices. In particular it shows how the implied volatility is affected. In the simplest case it suggests to fit the observed implied volatility surface I to an affine function of the log-moneyness-to-maturity ratio (LMMR):

$$I = a \times \text{LMMR} + b, \tag{1}$$
$$\text{LMMR} = \frac{\log(K/S)}{T-t},$$

and thereby to estimate two group market parameters a (the slope) and b (the intercept). Here K is strike price, T is expiration date and S is the asset price or index level at time t.

This approximation for the implied volatility surface arises from the following class of stochastic volatility models for the stock price (S_t) and the volatility driving process (Y_t):

$$dS_t = \mu S_t\, dt + f(Y_t) S_t\, dW_t,$$

where (W_t) is a standard Brownian motion and (Y_t) is an ergodic (or "mean-reverting") Markov Itô process with a unique invariant distribution, driven by a second Brownian motion $(\hat{Z}_t)$ that is correlated with (W_t). The function f is positive, bounded above and away from zero.

[1] By in-sample fit, we mean how well the model fits the data from which it is calibrated.

1.1 Volatility mean-reversion

To fix ideas and to allow us to refer to decorrelation speeds and rates of mean-reversion in terms of specific parameters, we shall write (Y_t) as an Ornstein-Uhlenbeck (OU) process:

$$dY_t = \alpha(m - Y_t)\,dt + \nu\sqrt{2\alpha}\,d\hat{Z}_t, \tag{2}$$

where (W_t) and $(\hat{Z}_t)$ have instantaneous correlation ρ, with $|\rho| < 1$.

The process Y has a unique invariant distribution, namely $\mathcal{N}(m, \nu^2)$, and is a simple building-block for a large class of stochastic volatility models described by the choice of $f(\cdot)$. The parameter α characterizes the *speed* of the process. It measures the exponential decorrelation rate, and ν characterizes the typical *size* of fluctuations of the volatility driving process.

The formula (1) is accurate in the regime of fast decorrelation (α large) of the process (Y_t). The precise order of this approximation is obtained in [7], at the level of call option prices, and a simple argument presented in [8, Sect. 5.3] translates this to units of implied volatility, leading to the accuracy of the approximation (1).

We stress that in our singular perturbation analysis, ν is fixed and the speed α is large. Another important asymptotic regime is characterized by α fixed and ν small. The latter small vol-of-vol regular perturbation expansions have been studied extensively by Hull and White [10] and Lewis [13]. Both these types of expansions reveal the success of stochastic volatility models in reproducing the implied volatility skew, and as functions of log-moneyness at a fixed time-to-maturity, they are similar for practical purposes. They differ, however, in predictions about the term-structure, that is the variation of implied volatility with time-to-maturity. A detailed comparison is beyond the scope of the present work.

Yet another regime is ν fixed and α small, corresponding to a slowly mean-reverting volatility process. This has been studied in [15] and [12]. At the end of the Conclusions, we comment on a work in preparation, combining fast and slow volatility factors and their associated regular and singular asymptotics.

In our approach, we are working within a class of models where the volatility function $f(y)$ and the market price of volatility risk $\Lambda(t, y)$, introduced in Sect. 3.1, are unspecified up to technical boundedness restrictions, and the resulting approximations do not depend on making a particular choice. In a number of other studies [4, 9, 13], the fully-specified Heston model is considered because closed form options prices are available, allowing comparison with expansions. This is done in the case of small vol-of-vol expansions in [4]. It would be interesting to investigate the case $f(y) = y^2$ (which leads to a square-root volatility process) and compare the fast mean-reverting asymptotics (suitably extended to relax the boundedness assumption) with the closed-form formulas, but this is beyond the topic of this paper.

1.2 Fast mean-reversion of volatility

In [6], we studied high-frequency S&P 500 data over the period of one year. The major difficulty with high-frequency data is pronounced intraday phenomena as-

sociated with microscopic trading patterns as described, for example, in [3]. In [6], it was shown how this 'periodic day effect' impacted the variogram and spectral methods used to analyze the data, and therefore how to account for it.

The result was, for the S&P 500 data examined, the presence of a *fast* volatility factor with rate of mean reversion $\alpha \sim 130 - 230$ (in annualized units). Note that the rate of mean-reversion of the unobserved volatility factor is extremely difficult to estimate *precisely*, hence the large range. This was confirmed from tests on simulated data.

Many empirical studies have looked at low-frequency (daily) data, with the data necessarily ranging over a period of years, and they have found a *low* rate of volatility mean-reversion. This does not contradict the empirical finding described above: analyzing data at lower frequencies over longer time periods would primarily pick up a slower time-scale of fluctuation and could not identify scales below the sampling frequency. At present, we will concentrate on the fast factor only and discuss extensions to include a slower factor in Sect. 6. We argue that this performs well empirically for S&P 500 options of up to the order of a few months from maturity, but that the extra factor extension is needed for longer options.

Another recent empirical study [2], this time of exchange rate dynamics, finds "the evidence points strongly toward two-factor [volatility] models with one highly persistent factor and one quickly mean-reverting factor". To pull one estimate from their Table 6, they find the rate of mean-reversion of the fast volatility factor for the US Dollar-Deutsche Mark exchange rate to be $\alpha = 237.5$ (in annualized units). For the other four exchange rates they also look at, the order of magnitude of this parameter is the same (hundreds).

1.3 Goal

The purpose of this article is to study the performance of (1) when fitted to real data and to introduce an extended theory with time-dependent periodic parameters that picks up a significant feature of the data. The criteria used here are

- goodness of the in-sample fit;
- stability of fitted parameters over time.

One common observation we would like to understand is:

For fixed $K \neq S$, the term-structure (variation with time-to-maturity $T - t$) is not like $(T - t)^{-1}$ but more like $(T - t)^{-1/2}$:

$$I \sim \frac{1}{\sqrt{T - t}}. \tag{3}$$

In Sect. 2, we look at the daily slope and intercept estimates (a and b in (1)) from S&P 500 index option implied volatilities. These shed some light on the previous observation. More strikingly, they identify a systematic almost-periodic behaviour in the slope of the skew (looked at as a function of LMMR). This is shown in Fig. 2. *We argue that just as implied volatility is a convenient transformation through which to observe deviations of options prices from the Black-Scholes theory, the LMMR*

representation is useful for picking out secondary phenomena, once the basic skew shape has been captured by a stochastic volatility theory.

We observe that the near-periodicity in the skew slope is related to the cycle of option expiration dates, which are around the third Friday of each month. An examination of the calendar shows that the "jumps" in the fitted skew slope, when the cycle returns to the beginning (it is not continuous), coincides with the closest-to-maturity options disappearing from the data. The fact that some options are expiring has a noticeable effect on the LMMR fit (1) to the dataset of a range of maturities. We model this empirical observation directly through a periodic variation in the *speed* of the process driving volatility in Sect. 4. This cycle effect may not be present in different markets with different expiration date structures, for example OTC markets where there is an almost-continuous range of maturity dates.

We present the adjusted asymptotic theory in the general case of time-dependent volatility parameters in Sect. 3, and then specialize to only the speed being time-dependent in Sect. 4. The obstacle that makes the mathematical analysis interesting is the nonsmoothness of the call (or put) option payoff. It requires a carefully-chosen regularization that allows use of the explicit Black-Scholes formula to quantify the rate of blow-up of higher order derivatives with respect to the stock variable, that are needed in the proof.

The asymptotic analysis shows that the formula (1) prevails under periodic variation in the parameters for the underlying price dynamics when we are far from maturity, but is modified when we get closer to maturity. This correction shows how the seemingly contradictory observations (1) and (3) can be bridged and explained in terms of a transition in between these two regimes.

Finally, in Sect. 5, we parametrize the near-periodic behaviour of the speed function and fit to data, leading to improved stability of the fitted parameters over time. We present the results from S&P 500 implied volatilities from January 2000 through March 2001, and remark that similar results were also obtained using the 1993 S&P 500 dataset used in [1].

2 Fit of implied volatilities to LMMR formula

2.1 Dataset

Our dataset contains implied volatilities from closing S&P 500 European option prices from January 2000 through March 2001. On each day, there are typically about 70 strikes (ranging over approximately 35% to 145% moneyness) and 7 maturities ($1, 2, 3, 6, 12, 15$ and 18 months). Throughout the empirical analysis, we will restrict ourselves to options *between* 70% *and* 102% *moneyness* so as to be on the safe side of liquidity issues. For much of the analysis we shall also restrict ourselves to options with three months or less to maturity since we are concerned with improving the fit to the shortest-dated options. We discuss an extension which could lead to improved capturing of the longer-term implied volatilities in the conclusions.

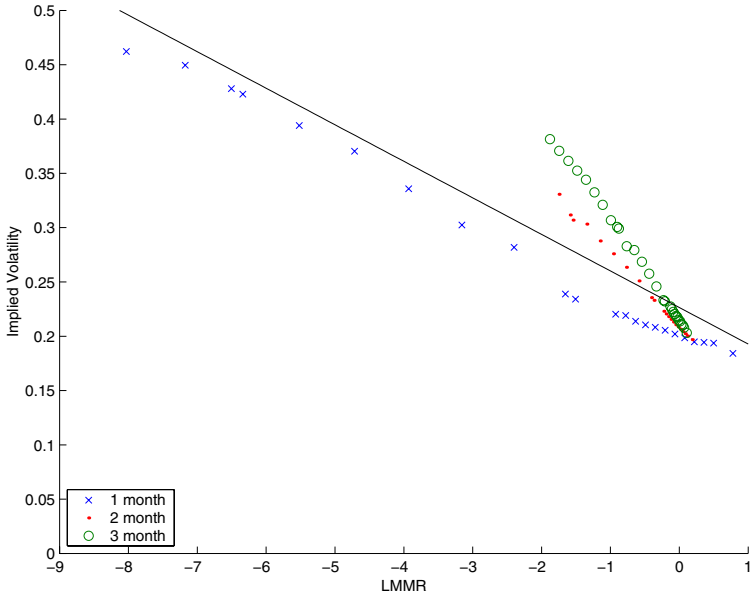

Fig. 1. Implied volatility as a function of LMMR on January 12, 2000, for options with at least two and at most 93 days to maturity. The one month options expire in January, the two month in February, and so on. Notice the distinct strands for each maturity, in particular the large gap between the one-month and two-month implied volatilities. The solid line is the least-squares straight line fit over all these maturities and strikes at once

2.2 LMMR fit across maturities up to one year

On a typical day if we plot implied volatility as a function of LMMR including options with at least two and at most 93 days (three months) to maturity, we see an array of distinct *diverging* strands corresponding to each maturity, as in Fig. 1.

This picture tells us that the sharper the (negative) slope of the implied volatility (as a function of LMMR) on this plot, the *longer* the maturity. (Typically, we think of S&P 500 implied volatilities as a function of strike becoming sharper for *shorter* maturities, so the adjustment by the time-to-maturity has over-corrected for this feature). In fact if we were to fit a straight line to each strand separately, the fits would be good, but there would be a different slope estimate for each maturity. The LMMR theory can be used satisfactorily in this way, but we would like to incorporate term-structure effects in observed implied volatilities in the way we use stochastic volatility models (that is, to capture the T dimension of the problem discussed in the introduction).

We see that if we try and fit across a range of maturities, the in-sample fit is not very good (see the solid line in Fig. 1).

More strikingly, if we look at the daily variation of the fitted parameters, as in Fig. 2, there is a pronounced near-periodic repetition in the LMMR skew slope a, the period being on the order of twenty trading days. Further, there is a clear stratification, meaning the slopes jump down once each period.

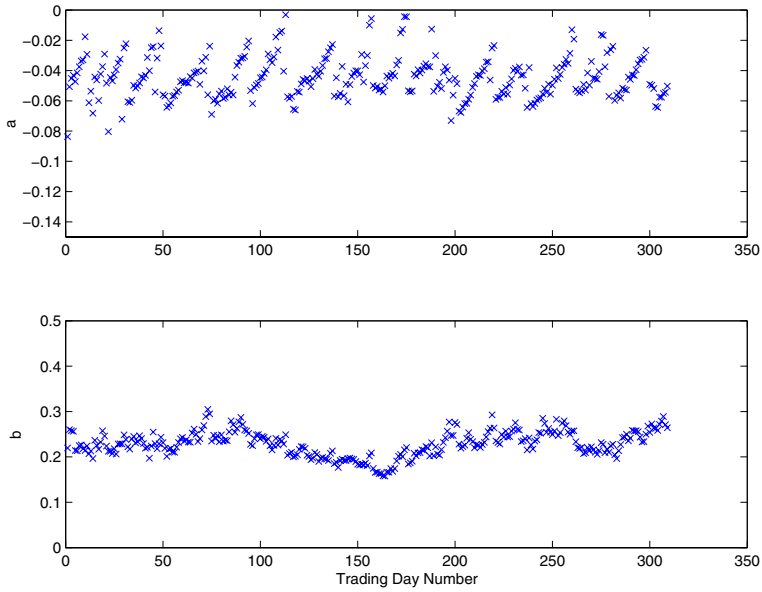

Fig. 2. Daily fitted LMMR parameters for options with at least 2 and at most 93 days to maturity, showing the periodicity corresponding to the monthly maturity cycle. Statistics: Mean $(a) = -0.0444$, Std $(a) = 0.0137$; Mean $(b) = 0.2288$, Std $(b) = 0.0278$

2.3 Examination of the breaks

To try to understand what causes the sudden but regular discontinuities in the a's, we look at in-sample fits on days before and after a jump in Fig. 2. It is clear from Fig. 2 that the breaks are roughly a month apart, and a comparison with the calendar shows they occur exactly when the shortest maturity options included in the fit disappear (in Fig. 2, two days before the third Friday of each month, for example).

Figure 3 shows fits on the last day of one stratum and the first day of the next.

The disappearance of the shortest option has sharpened the estimated LMMR slope dramatically.

We conclude from this that the option expiration dates (which are known in advance) have some bearing on the evolution of the whole implied volatility surface. In other words, the periodic pattern is not merely a commentary on the LMMR fitting procedure, but contains information about the systematic behaviour of market option prices. This phenomenon is not easy to see directly from implied volatility data, but is revealed clearly once LMMR has been used as a basis to filter out stochastic volatility effects. In Sect. 3, we present a theory based on a periodic variation of the speed of the volatility driving process that accounts for this pattern. It is based on modifying the asymptotic theory for time-varying parameters.

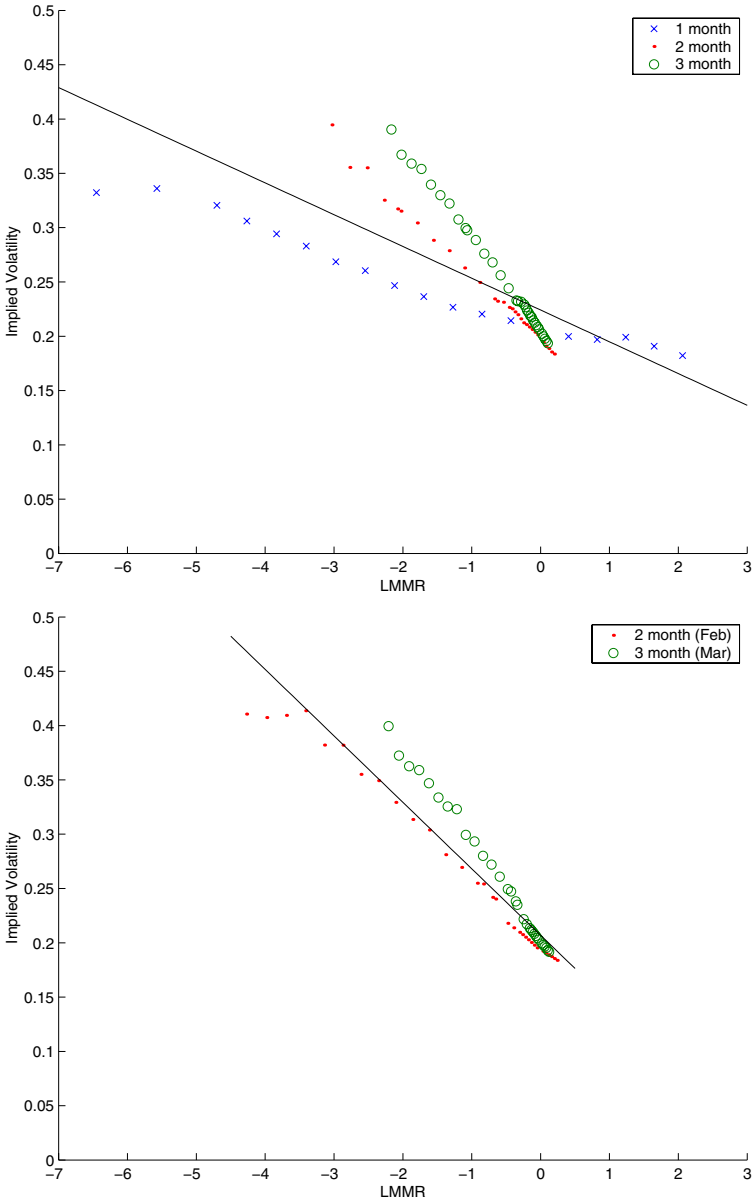

Fig. 3. In-sample LMMR fit the last day before a jump in the fitted slope (January 18, *top*) and after (January 19, *bottom*). Note that the shortest option has disappeared from data we are using in the fit

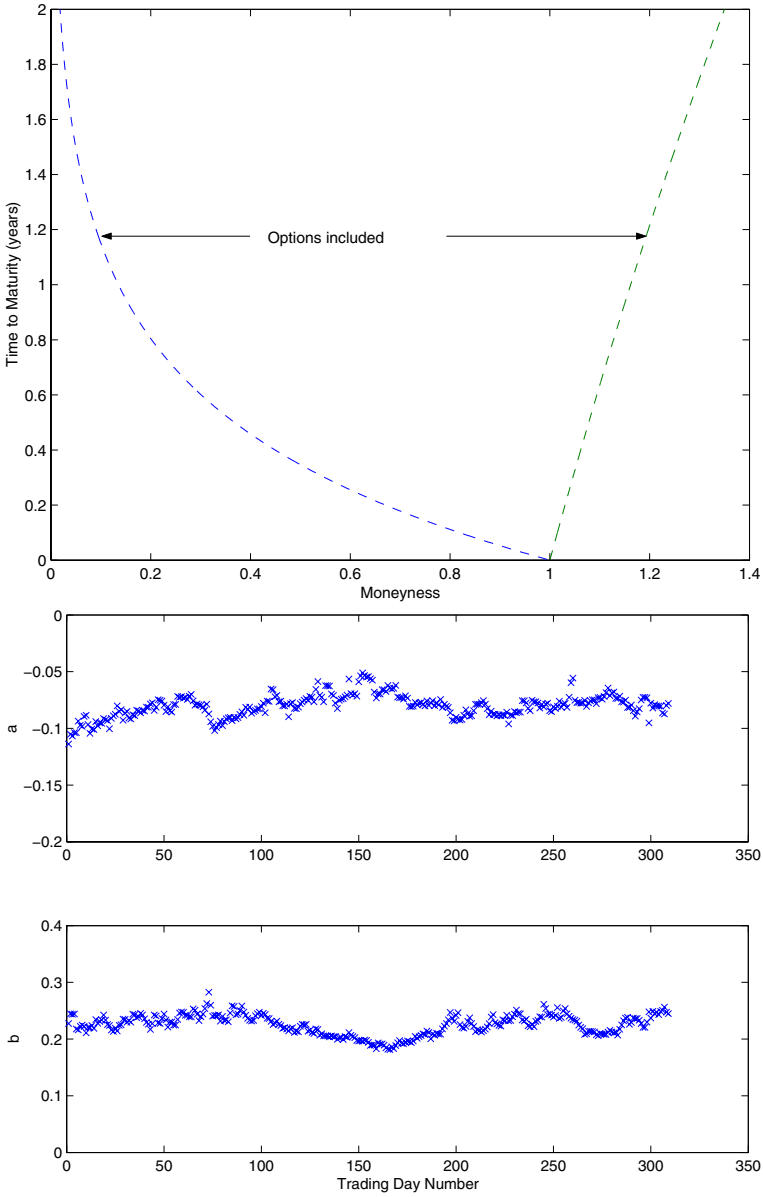

Fig. 4. Range of options included in restricted LMMR fit (*top*). The daily fitted LMMR parameters are shown in the bottom plot. Statistics: Mean $(a) = -0.0801$, Std $(a) = 0.0100$; Mean $(b) = 0.2251$, Std $(b) = 0.0180$

2.3.1 LMMR fit to restricted data

If we restrict the range of options included in the LMMR fit so that very short-dated options are mainly excluded, the stability of the fitted parameter improves considerably. We take options whose LMMR is the range

$$-2 < \text{LMMR} < 0.15.$$

This region is sketched in the top plot of Fig. 4: most long-dated options are included, but very few as the time-to-maturity gets shorter.

The bottom plot shows the improvement in parameter stability, as compared with Fig. 2. The periodic variation in the slope is no longer seen. Note however this does not solve the problem of the separate strands for option of each maturity seen in Fig. 1. In other words, the in-sample fits to the longer options is still relatively unsatisfactory. The extension we introduce in Sect. 3 will address both problems (stability and in-sample fit) without removing the short options as in Fig. 4.

2.4 Fitting LMMR maturity-by-maturity

Many pricing models give rise to predicted implied volatility skews that fit the data well at fixed maturities (for example [11,14]). This is not surprising because the implied volatility skew is quite smooth as a function of strike and any theory with reasonable properties (asymmetric return distribution, leptokurtosis) should be able to interpolate across strikes very well. The problem is going across maturities, which we try to achieve in this paper.

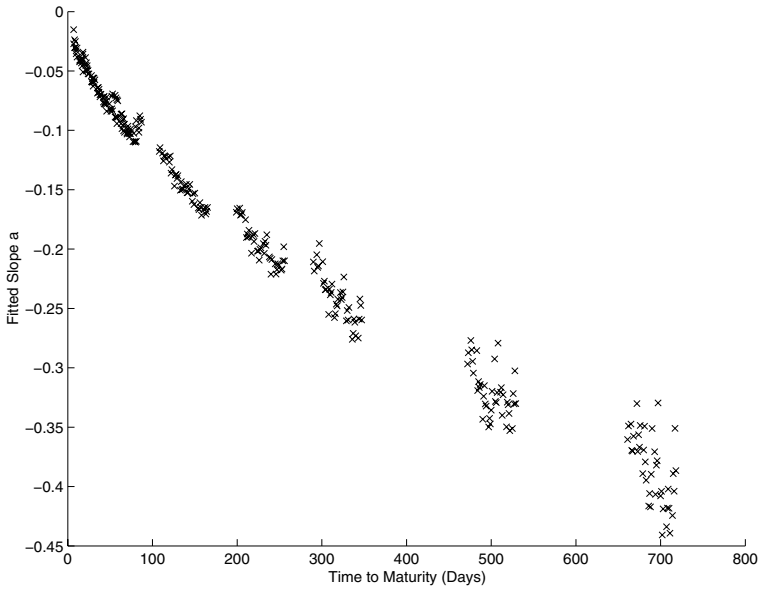

Fig. 5. LMMR slopes from fits performed daily and maturity-by-maturity as a function of time-to-expiration, using observations in January & February 2000

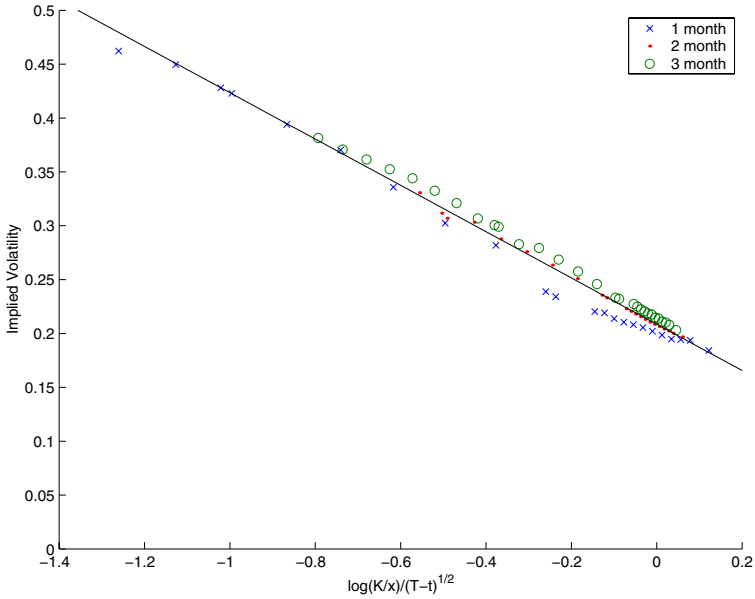

Fig. 6. Implied volatility as a function of $\log(K/S)/\sqrt{T-t}$ on January 12, 2000, for options with at least two and at most 93 days to maturity

The evidence so far indicates that the slope a in (1) should depend on t and T in any modified model that will fit the data better. Many theories are in practice calibrated maturity-by-maturity. As an empirical attempt to understand the variation in a, we fit the LMMR formula each day to obtain estimates $a_{(t,T)}$ and $b_{(t,T)}$ that depend on the day t and the maturity T. In Fig. 5, using the dataset described in Sect. 2.1, we plot $a_{(t,T)}$ against $T-t$ which shows that there is a systematic variation of a with $T-t$. From short-dated options, there is a sharp decay which seems to flatten as $T-t$ increases. Empirically then we might want to vary a like

$$a \sim -c(T-t)^q \qquad 0 < q < 1,\ c > 0.$$

As an example, we try $q = 1/2$, which captures the square-root decay of the term-structure discussed in Sect. 1. We fit implied volatilities daily to the *square-root formula*:

$$I = A\frac{\log(K/S)}{\sqrt{T-t}} + B, \tag{4}$$

and estimate the slope and intercept coefficients A and B. Figure 6 shows the fit to the new square-root formula on a particular day. The separate strands for each maturity are still there, but the slopes are more closely lined up. There is an improvement in the stability of the fitted parameters as shown in Fig. 7 over those in Fig. 2.

We conclude that we would like to modify the theory so that it

– produces good fits when using options data over a range of maturities;

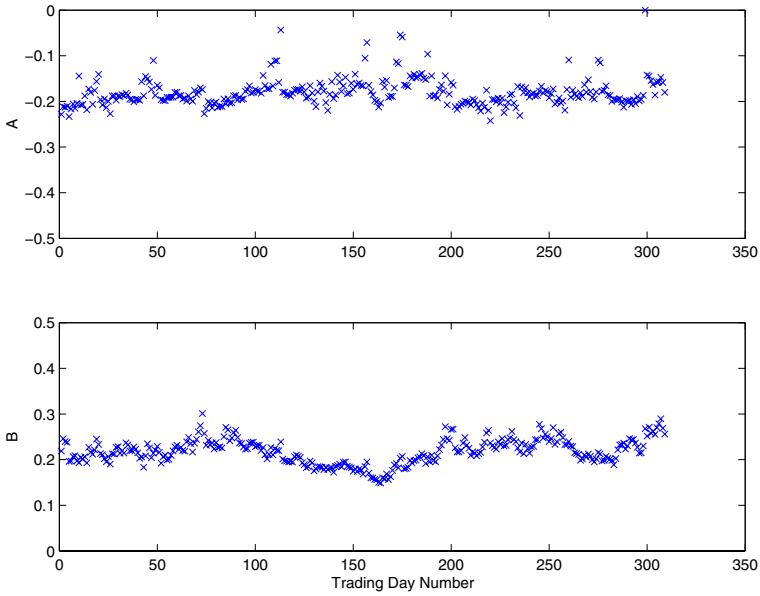

Fig. 7. Daily fitted slope A & intercepts B for options with at least 2 and at most 93 days to maturity fit to an affine function of $\log(K/S)/\sqrt{T-t}$. Statistics: Mean $(A) = -0.1806$, Std $(A) = 0.0299$; Mean $(B) = 0.2177$, Std $(B) = 0.0272$

- exhibits stability over time of estimated parameters;
- explains the periodic behaviour associated with expiration dates;
- is consistent with a stochastic volatility no-arbitrage derivation.

It is the last feature that prevents us from simply using the square-root formula. While it performs well empirically, it is not the product of a model (as far as we know), so having estimated A and B in (4), it is not clear what to do with them in terms of pricing other exotic derivatives or risk-management problems.

In the next section, we analyze a class of models that manage to display the properties listed above. We begin in Sect. 3 by establishing the asymptotic theory which takes into account time-dependent volatility parameters. The reader interested mainly in the application can skip the mathematical proofs and go directly to the implications for implied volatility approximation presented in Sect. 3.3.

3 Time-dependent parameters

We now discuss a class of mean-reverting stochastic volatility models of the underlying stock price with general (deterministic) time-dependent volatility parameters $\alpha(t), \nu(t), m(t)$ and $\rho(t)$ in (2). The main difficulty in proving the accuracy of an expansion in powers of the characteristic speed of the volatility driving process is the nondifferentiability of the option payoff at the strike. We show in Sects. 3.1-3.2.3 how to tackle this with a regularization method, and the result is given in Theorem 1 below. We outline the proof, which is an extension of that in [7] to incorporate

time-dependent parameters. We point out the main modifications and refer to [7] for certain details. The full proof for the current case is available from the authors upon request.

In Sect. 3.3, we give the important application of the mathematical results for calibration from implied volatility. We discuss in Sect. 4 the special case where *only* α is time-dependent, which is used in the data analysis of Sect. 5.

3.1 Models and derivative pricing

We assume a probability space $(\Omega, \mathcal{F}, I\!P)$ and write the class of models we study under the real-world measure $I\!P$ as

$$\begin{aligned} dS_t^\varepsilon &= \mu S_t^\varepsilon \, dt + f(Y_t^\varepsilon) S_t^\varepsilon \, dW_t \\ dY_t^\varepsilon &= \frac{1}{\varepsilon}\alpha(t)(m(t) - Y_t^\varepsilon)\, dt + \frac{1}{\sqrt{\varepsilon}}\nu(t)\sqrt{2\alpha(t)}\, d\hat{Z}_t, \end{aligned} \tag{5}$$

where W and $\hat{Z}$ are standard Brownian motions with instantaneous correlation $\rho(t) \in (-1, 1)$:

$$I\!E\{dW_t \, d\hat{Z}_t\} = \rho(t)\, dt.$$

The volatility function f, and its reciprocal $1/f$, the local mean-reversion speed α and the function ν are assumed positive and bounded. The function m is assumed bounded.

The constant $\varepsilon > 0$ is small to represent *fast* mean-reversion. Now $\alpha(t)/\varepsilon$ is a measure of the *local* speed of the process and $\nu(t)$ a measure of the local size of volatility fluctuations and we are interested in asymptotic expansions in the limit $\varepsilon \downarrow 0$.

As is standard in the practical application of such incomplete market models, we assume that the market *selects* an equivalent martingale (pricing) measure $I\!P^\star \equiv I\!P$ under which $(S^\varepsilon, Y^\varepsilon)$ evolves according to

$$\begin{aligned} dS_t^\varepsilon &= rS_t^\varepsilon \, dt + f(Y_t^\varepsilon) S_t^\varepsilon \, dW_t^\star \\ dY_t^\varepsilon &= \left[\frac{1}{\varepsilon}\alpha(t)(m(t) - Y_t^\varepsilon) - \frac{\nu(t)\sqrt{2\alpha(t)}}{\sqrt{\varepsilon}} \Lambda(t, Y_t^\varepsilon)\right] dt + \frac{1}{\sqrt{\varepsilon}}\nu(t)\sqrt{2\alpha(t)}\, d\hat{Z}_t^\star, \end{aligned} \tag{6}$$

where $W^\star$ and $\hat{Z}^\star$ are standard $I\!P^\star$-Brownian motions with instantaneous correlation $\rho(t) \in (-1, 1)$:

$$I\!E^\star\{dW_t^\star \, d\hat{Z}_t^\star\} = \rho(t)\, dt.$$

Here, $r \geq 0$ is the riskless interest rate, assumed constant, and

$$\Lambda(t, y) = \rho(t)\frac{\mu - r}{f(y)} + \gamma(t, y)\sqrt{1 - \rho(t)^2}$$

is the combined market price of volatility risk, including the volatility risk premium $\gamma(t, y)$, which characterizes the risk-neutral measure $I\!P^\star$ chosen by the market. We

assume that γ is a bounded function of (t, y) only, and therefore that Y^ε is a Markov process under $I\!P^\star$.

We consider a European call option with strike price K, for which the payoff at maturity time T is

$$H(S_T^\varepsilon) = (S_T^\varepsilon - K)^+.$$

Its price at time t is given by

$$P^\varepsilon(t, S, y) = I\!E^\star \left\{ e^{-r(T-t)} H(S_T^\varepsilon) \mid S_t^\varepsilon = S, Y_t^\varepsilon = y \right\}, \tag{7}$$

using the Markov property of the pair $(S^\varepsilon, Y^\varepsilon)$. We subsequently use the shorthand notation $I\!E^\star{}_{t,S,y}$ for such conditional expectations.

Under our assumptions, the function $P^\varepsilon(t, S, y)$ is the classical solution of the associated Feynman-Kac partial differential equation (PDE)

$$\begin{aligned} &\frac{\partial P^\varepsilon}{\partial t} + \frac{1}{2} f(y)^2 S^2 \frac{\partial^2 P^\varepsilon}{\partial S^2} + r\left(S \frac{\partial P^\varepsilon}{\partial S} - P^\varepsilon \right) \\ &+ \frac{1}{\sqrt{\varepsilon}} \nu(t) \sqrt{2\alpha(t)} \left(\rho(t) f(y) S \frac{\partial^2 P^\varepsilon}{\partial S \partial y} - \Lambda(t, y) \frac{\partial P^\varepsilon}{\partial y} \right) \\ &+ \frac{1}{\varepsilon} \alpha(t) \left(\nu^2(t) \frac{\partial^2 P^\varepsilon}{\partial y^2} + (m(t) - y) \frac{\partial P^\varepsilon}{\partial y} \right) = 0 \end{aligned} \tag{8}$$

in $S > 0, y \in I\!R, t < T$, with the terminal condition being the payoff at maturity, $P^\varepsilon(T, S, y) = H(S)$.

3.2 Main theorem

The theorem of this section describes the approximation of derivative prices we obtain from an asymptotic expansion of P^ε, the solution of (8). In order to formulate it, we define

$$\langle g \rangle_t = \frac{1}{\sqrt{2\pi\nu^2(t)}} \int_{-\infty}^{\infty} g(\xi) e^{-(\xi - m(t))^2 / 2\nu^2(t)} d\xi,$$

the average of g with respect to the distribution $\mathcal{N}(m(t), \nu^2(t))$, and we introduce the following time-dependent deterministic quantities:

$$\bar{\sigma}^2(t) = \langle f^2 \rangle_t, \tag{9}$$

$$\overline{\bar{\sigma}^2}(t; T) = \frac{1}{T-t} \int_t^T \bar{\sigma}^2(s)\, ds, \tag{10}$$

$$V_3^\varepsilon(t) = \frac{\sqrt{\varepsilon}\, \rho(t) \nu(t)}{\sqrt{2\alpha(t)}} \langle f(\cdot) \frac{\partial \phi}{\partial y}(t, \cdot) \rangle_t \tag{11}$$

$$V_2^\varepsilon(t) = 2V_3^\varepsilon(t) - \frac{\sqrt{\varepsilon}\, \nu(t)}{\sqrt{2\alpha(t)}} \langle \Lambda(t, \cdot) \frac{\partial \phi}{\partial y}(t, \cdot) \rangle_t \tag{12}$$

$$\overline{V}_2^\varepsilon(t; T) = \frac{1}{T-t} \int_t^T V_2^\varepsilon(s)\, ds \tag{13}$$

$$\overline{V}_3^\varepsilon(t; T) = \frac{1}{T-t} \int_t^T V_3^\varepsilon(s)\, ds, \tag{14}$$

where $\phi(t,y)$ satisfies the Poisson equation

$$\mathcal{L}_0\phi(t,y) = f^2(y) - \bar{\sigma}^2(t), \tag{15}$$

and $\varepsilon^{-1}\mathcal{L}_0$ is the infinitesimal generator of the process Y^ε (5) under $I\!P$. We give $\mathcal{L}_0$ explicitly in (19). By the boundedness assumption on f, we can (and do) choose ϕ to have a bounded first derivative in y. See [8, Sect. 5.2.3], for example.

Our approximation is given by the sum of the following two quantities

$$P_0(t,S) = P_{BS}\left(t,S;K,T;\sqrt{\overline{\bar{\sigma}^2}(t;T)}\right), \tag{16}$$

$$\widetilde{P_1}(t,S) = -(T-t)\left[\overline{V}^\varepsilon_2(t;T)S^2\frac{\partial^2 P_0}{\partial S^2} + \overline{V}^\varepsilon_3(t;T)S^3\frac{\partial^3 P_0}{\partial S^3}\right], \tag{17}$$

where $P_{BS}(t,S;K,T;\sigma)$ is the Black-Scholes price of the European call H with the constant volatility σ, given explicitly by the Black-Scholes formula. Note that P_0 is independent of ε, and that through the coefficients $\overline{V}^\varepsilon_2$ and $\overline{V}^\varepsilon_3$, $\widetilde{P_1}$ is proportional to $\sqrt{\varepsilon}$. In the case that the volatility parameters are constant, the approximation reduces to the one detailed in [8].

Theorem 1 *Under the boundedness assumptions on the model coefficients f,γ,m,ν,α and ρ, at any fixed time $t_0 < T$, and fixed points $S > 0, y \in I\!R$,*

$$\lim_{\varepsilon\downarrow 0}\frac{|P^\varepsilon(t_0,S,y) - (P_0(t_0,S) + \widetilde{P_1}(t_0,S))|}{\varepsilon^{1-p}} = 0,$$

for any $p > 0$.

If the payoff H was a smooth function, the proof would simply consist in writing a PDE for the approximation error and using its probabilistic representation to control it. The error expression involves S-derivatives of the Black-Scholes price $P_{BS}(t,S)$ up to order seven. However, implied volatilities are calculated from liquid call (and put) options, for which the payoff is not differentiable at the strike, and higher derivatives of the Black-Scholes price blow up close to maturity. We discuss the case of the call in this section. The case of the put can be handled trivially by put-call parity.

In the nonsmooth case of the call, the proof requires a carefully-chosen regularization of the payoff. The regularization we introduce below is consistent with the Black-Scholes PDE *with time-dependent volatility*. The proof of the theorem is a refinement of the arguments presented in [7] to deal with time-dependent volatility parameters. It uses three lemmas which we state below. The proofs of these are extensions of the corresponding lemmas in [7], and we will merely highlight the new difficulties here.

3.2.1 Operator notation

It is convenient to make the change of variable $x = \log S$. With a slight abuse of notation, we continue to denote the call option price by $P^\varepsilon(t,x,y)$, and we rewrite the PDE (8) with the log-stock variable as

$$\mathcal{L}^\varepsilon P^\varepsilon = 0, \qquad x, y \in I\!R, t < T, \tag{18}$$

with the terminal condition $P(T,x,y) = h(x) = H(e^x)$, where

$$\mathcal{L}^\varepsilon = \frac{1}{\varepsilon}\mathcal{L}_0 + \frac{1}{\sqrt{\varepsilon}}\mathcal{L}_1 + \mathcal{L}_2,$$

and the operators $\mathcal{L}_0, \mathcal{L}_1$ and $\mathcal{L}_2$ are defined by

$$\mathcal{L}_0 = \alpha(t)\left(\nu^2(t)\frac{\partial^2}{\partial y^2} + (m(t)-y)\frac{\partial}{\partial y}\right) \tag{19}$$

$$\mathcal{L}_1 = \nu(t)\sqrt{2\alpha(t)}\left(\rho(t)f(y)\frac{\partial^2}{\partial x \partial y} - \Lambda(t,y)\frac{\partial}{\partial y}\right) \tag{20}$$

$$\mathcal{L}_2 = \frac{\partial}{\partial t} + \frac{1}{2}f(y)^2\frac{\partial^2}{\partial x^2} + \left(r - \frac{1}{2}f(y)^2\right)\frac{\partial}{\partial x} - r\cdot. \tag{21}$$

Similarly, we denote the Black-Scholes call formula as a function of the log stock price x by $P_{BS}(t,x;K,T;\sigma)$ and our approximation to P^ε in terms of the log variable is

$$Q^\varepsilon(t,x) = P_0(t,x) + \sqrt{\varepsilon}P_1(t,x),$$

where

$$P_0(t,x) = P_{BS}\left(t,x;K,T;\sqrt{\bar{\sigma}^2(t;T)}\right), \tag{22}$$

$$\sqrt{\varepsilon}\,P_1 = -(T-t)\mathcal{A}^\varepsilon(t;T)P_0, \tag{23}$$

and

$$\mathcal{A}^\varepsilon(t;T) = \overline{V}_3^\varepsilon(t;T)\frac{\partial^3}{\partial x^3} + (\overline{V}_2^\varepsilon(t;T) - 3\overline{V}_3^\varepsilon(t;T))\frac{\partial^2}{\partial x^2} + (2\overline{V}_3^\varepsilon(t;T) - \overline{V}_2^\varepsilon(t;T))\frac{\partial}{\partial x}\,. \tag{24}$$

The correction $\sqrt{\varepsilon}\,P_1$ is the solution of a Black-Scholes PDE with time-dependent volatility coefficient $\bar{\sigma}(t)$ and a source term:

$$\langle\mathcal{L}_2\rangle_t(\sqrt{\varepsilon}P_1) = \left(V_3^\varepsilon(t)\frac{\partial^3}{\partial x^3} + (V_2^\varepsilon(t) - 3V_3^\varepsilon(t))\frac{\partial^2}{\partial x^2} + (2V_3^\varepsilon(t) - V_2^\varepsilon(t))\frac{\partial}{\partial x}\right)P_0,$$

where $\langle\mathcal{L}_2\rangle_t$ is given by (21) with $f(y)^2$ replaced by $\bar{\sigma}^2(t)$.

3.2.2 Regularization of the payoff

The main ingredient of the proof is an appropriate regularization of the payoff, as shown in [7]. We adapt here this regularization to the case to time-dependent volatility coefficients. For a call option, we approximate the payoff h with the Black-Scholes price of a call with fixed volatility

$$\sigma_0 := \sqrt{\overline{\bar{\sigma}^2}(t_0; T)},$$

where $\overline{\bar{\sigma}^2}$ was defined in (10), and time to maturity δ. We define

$$h^{\delta,t_0}(x) := P_{BS}(T, x; K, T + \delta; \sigma_0),$$

which can also be written $h^{\delta,t_0}(x) = P_{BS}(T - \delta, x; K, T; \sigma_0)$, because the Black-Scholes formula is a function of time-to-maturity. Note that in the unregularized case, $\delta = 0$, we have $h^{0,t_0} = h$. For $\delta > 0$, this new payoff is smooth in x. The price $P^{\varepsilon,\delta}(t, x, y)$ of the option with the regularized payoff solves

$$\begin{aligned} \mathcal{L}^{\varepsilon} P^{\varepsilon,\delta} &= 0, \qquad t < T; x, y \in I\!R \\ P^{\varepsilon,\delta}(T, x, y) &= h^{\delta,t_0}(x). \end{aligned} \tag{25}$$

Let $Q^{\varepsilon,\delta}(t, x)$ denote the following approximation to the regularized option price:

$$P^{\varepsilon,\delta} \approx Q^{\varepsilon,\delta} := P_0^{\delta} + \sqrt{\varepsilon} P_1^{\delta},$$

where $P_0^{\delta}(t, x)$ and $P_1^{\delta}(t, x)$ are defined as P_0 and P_1 in (22) and (23), with the payoff h replaced by the regularized payoff h^{δ,t_0}.

3.2.3 Lemmas and Proof of Theorem

We establish the following pathway to proving Theorem 1. The proofs of the following three lemmas are extensions of those given in [7]. The main modification in the arguments that allows us to exploit the Black-Scholes formula is to note that we can write $P_0^{\delta}(t, x)$ in terms of P_{BS}, but with a volatility parameter $\tilde{\sigma}(t; T)$, where

$$\tilde{\sigma}^2(t; T) := \frac{1}{T + \delta - t} \left(\int_t^T \bar{\sigma}^2(s)\, ds + \delta \sigma_0^2 \right),$$

that depends on t and δ. Our choice of regularization has been made so that at $t = t_0$, we have $\tilde{\sigma}^2(t_0; T) = \sigma_0^2$.

In the following, constants may depend on (t_0, T, x, y) but not on (ε, δ):

Lemma 1 *Fix the point (t_0, x, y) where $t_0 < T$. There exist constants $\delta_1 > 0$, $\varepsilon_1 > 0$ and $c_1 > 0$ such that*

$$|P^{\varepsilon}(t_0, x, y) - P^{\varepsilon,\delta}(t_0, x, y)| \leq c_1 \delta$$

for all $0 < \delta < \delta_1$ and $0 < \varepsilon < \varepsilon_1$.

This establishes that the solutions to the regularized and unregularized problems are close.

Lemma 2 *Fix the point (t_0, x, y) where $t_0 < T$. There exist constants $\delta_2 > 0$, $\varepsilon_2 > 0$ and $c_2 > 0$ such that*

$$|Q^{\varepsilon}(t_0, x) - Q^{\varepsilon,\delta}(t_0, x)| \leq c_2\delta$$

for all $0 < \delta < \delta_2$ and $0 < \varepsilon < \varepsilon_2$.

This establishes that the approximations to the regularized and unregularized prices are close.

Lemma 3 *Fix the point (t_0, x, y) where $t_0 < T$. There exist constants $\delta_3 > 0$, $\varepsilon_3 > 0$ and $c_3 > 0$ such that*

$$|P^{\varepsilon,\delta}(t_0, x, y) - Q^{\varepsilon,\delta}(t_0, x)| \leq c_3 \left(\varepsilon |\log \delta| + \varepsilon \sqrt{\frac{\varepsilon}{\delta}} + \varepsilon \right),$$

for all $0 < \delta < \delta_3$ and $0 < \varepsilon < \varepsilon_3$.

This establishes that for fixed $\delta > 0$, when the payoff is smooth, the approximation is $\mathcal{O}(\varepsilon)$.

Theorem 1 follows easily from these lemmas:

Proof of Theorem 1. Take $\bar{\delta} = \min(\delta_1, \delta_2, \delta_3)$ and $\bar{\varepsilon} = \min(\varepsilon_1, \varepsilon_2, \varepsilon_3)$. Then, using Lemmas 1, 2 and 3, we obtain

$$\begin{aligned} |P^{\varepsilon} - Q^{\varepsilon}| &\leq |P^{\varepsilon} - P^{\varepsilon,\delta}| + |Q^{\varepsilon,\delta} - Q^{\varepsilon}| + |P^{\varepsilon,\delta} - Q^{\varepsilon,\delta}| \\ &\leq 2\max(c_1, c_2)\delta + c_3 \left(\varepsilon |\log \delta| + \varepsilon \sqrt{\frac{\varepsilon}{\delta}} + \varepsilon \right), \end{aligned}$$

for $0 < \delta < \bar{\delta}$ and $0 < \varepsilon < \bar{\varepsilon}$, where the functions are evaluated at the fixed (t_0, x, y). Taking $\delta = \varepsilon$, we have

$$|P^{\varepsilon} - Q^{\varepsilon}| \leq c_4(\varepsilon + \varepsilon |\log \varepsilon|),$$

for some fixed $c_4 > 0$, and Theorem 1 follows.

3.3 Implications for implied volatility approximation

To summarize, the asymptotic approximation at a generic time $t < T$ for the call option price P^{ε} (in the original stock variable S), defined in (7), is given by

$$P^{\varepsilon}(t, S, y) = P_{BS} - (T-t) \left[\overline{V}_2^{\varepsilon}(t; T) S^2 \frac{\partial^2 P_{BS}}{\partial S^2} + \overline{V}_3^{\varepsilon}(t; T) S^3 \frac{\partial^3 P_{BS}}{\partial S^3} \right] + \mathcal{O}(\varepsilon^p),$$

for any $p < 1$, where the argument of the Black-Scholes formula P_{BS} is $\left(t, S; K, T; \sqrt{\overline{\sigma^2}(t; T)} \right)$, $\overline{\sigma^2}$ is defined in (10) and $\overline{V}_2^{\varepsilon}(t; T)$ and $\overline{V}_3^{\varepsilon}(t; T)$ in (13)

and (14) respectively. This is analogous to the constant parameter case given in [8] with the replacements

$$\bar{\sigma}^2 \mapsto \overline{\bar{\sigma}^2}(t;T); \qquad V_2 \mapsto \overline{V}^{\varepsilon}_2(t;T); \qquad V_3 \mapsto \overline{V}^{\varepsilon}_3(t;T).$$

The corresponding approximation for the implied volatility, can be obtained from its defining relationship

$$P_{BS}(t,S;K,T;I) = P^{\varepsilon}(t,S,y).$$

Formally inserting an expansion for implied volatility

$$I = I_0 + \sqrt{\varepsilon} I_1 + \varepsilon I_2 + \cdots,$$

a simple calculation, given for example in [8, Sect. 5.3], yields

$$\begin{aligned} I_0 &= \sqrt{\overline{\bar{\sigma}^2}(t;T)}, \\ \sqrt{\varepsilon} I_1 &= -(T-t)\frac{\left[\overline{V}^{\varepsilon}_2(t;T) S^2 \frac{\partial^2 P_{BS}}{\partial S^2} + \overline{V}^{\varepsilon}_3(t;T) S^3 \frac{\partial^3 P_{BS}}{\partial S^3}\right]}{\frac{\partial P_{BS}}{\partial \sigma}}, \end{aligned}$$

where the volatility parameter used in P_{BS} is I_0.

Using the explicit expressions, we obtain the approximation

$$I \approx a(t;T)\frac{\log(K/S)}{T-t} + b(t;T), \tag{26}$$

where

$$\begin{aligned} a(t;T) &= -\frac{1}{(\overline{\bar{\sigma}^2}(t;T))^{3/2}} \overline{V}^{\varepsilon}_3(t;T) \qquad (27) \\ b(t;T) &= \sqrt{\overline{\bar{\sigma}^2}(t;T)} - a(t;T)\left(r + \frac{3}{2}\overline{\bar{\sigma}^2}(t;T)\right) - \frac{\overline{V}^{\varepsilon}_2(t;T)}{\sqrt{\overline{\bar{\sigma}^2}(t;T)}}, \end{aligned}$$

and we do not show explicitly the ε dependence of a and b.

Finally, we note that there is an extension to the asymptotic approximations for exotic options (for example American puts, barriers and Asians), similar to those presented in [8], to incorporate time-dependent volatility parameters. In this case, the approximations depend on $V^{\varepsilon}_2(t)$ and $V^{\varepsilon}_3(t)$.

4 Calibration

Typically, time-dependent parameters are calibrated by first interpolating the observed term-structure of implied volatilities and then differentiating with respect to maturity. This is not the path we follow. Instead we will take advantage of the systematic and near-periodic time-variation of a quantity derived from S&P 500 implied volatilities, as seen in Fig. 2.

We will show that taking the speed $\alpha(t)$ of the volatility driving process to be periodically time-varying, with the typical size of the fluctuations ν, the mean-level m and the correlation ρ constant, is sufficient to capture the periodicity in the fitted skew parameter observed in Fig. 2. In this case, the averaging $\langle\cdot\rangle$ does *not* depend on t, but we still have time-dependent $V_2^\varepsilon(t), V_3^\varepsilon(t), a(t;T), b(t;T)$. Furthermore, $\overline{\bar{\sigma}^2}(t;T) = \bar{\sigma}^2$ is constant and can be estimated from historical daily returns as the annualized sample standard deviation over a period of several months, which is long compared to the mean-reversion time of volatility fast scale. This choice of estimator is justified *a posteriori* by checking that the calibrated $b - \bar{\sigma}$ are small. This is addressed in Sect. 5 for our dataset.

Since now only α is time-dependent, it follows from (13), (14) and (27) that

$$a(t;T) = \frac{k_a}{T-t}\int_t^T \frac{1}{\sqrt{\alpha(s)}}\,ds \tag{28}$$
$$b(t;T) = \bar{\sigma} + \frac{k_b}{T-t}\int_t^T \frac{1}{\sqrt{\alpha(s)}}\,ds,$$

for some small constants k_a and k_b (of order $\sqrt{\varepsilon}$). (We remark that an alternative way to obtain a similar calibration procedure is to assume only the correlation ρ to be time-dependent, but as this seems less economically plausible, we do not pursue it here).

4.1 Calendar function

It remains to choose a model for the time-dependence. Based on our previous analysis of the data in Sect. 2, we propose the following form:

$$\frac{1}{\sqrt{\alpha(t)}} = c\,(T_{n(t)} - t)^{1/2}, \tag{29}$$

for some constant c. Here $\{T_k\}$ are the option maturity dates (the third Friday of each month), and

$$n(t) = \inf\{n : T_n \geq t\},$$

the next expiration after time t. In other words, $\alpha(t)$ is a near-periodic function described by the time till the next maturity date. Clearly, there are many ways to generalize this idea, the simplest being to vary the power in (29). We will show in Sect. 5 that this choice captures well the maturity-dependent behaviour of the implied volatilities that we observed in Sect. 2.

Note that $\alpha(t)$ as chosen in (29) is unbounded. For convenience, we perform the calculations below with this α, but a large cutoff can be included to be consistent with the boundedness assumption made in Sect. 3, with a negligible effect on the results.

For simplicity, we assume the expiration dates are evenly spaced:

$$T_{k+1} - T_k =: \Delta T,$$

where in applications ΔT will be a month. From (28), we have that

$$a(t,T) = \frac{k_a}{T-t}\int_t^T (T_{n(s)} - s)^{1/2} ds.$$

Given a time t and an option expiration date T, if we decompose the time-to-maturity as

$$T - t = \Delta T(m_0 + \eta)$$

with $0 \leq \eta < 1$, we have

$$\begin{aligned}
&\frac{1}{T-t}\int_t^T (T_{n(s)} - s)^{1/2} ds \\
&= \frac{1}{T-t}\left(\int_{(1-\eta)\Delta T}^{\Delta T} (\Delta T - s)^{1/2} ds + m_0 \int_0^{\Delta T} (\Delta T - s)^{1/2} ds\right) \\
&= \frac{2(\Delta T)^{1/2}}{3}\frac{m_0 + \eta^{3/2}}{m_0 + \eta}.
\end{aligned}$$

Using (26), we get the following approximation for the implied volatility

$$I = A\frac{\log(K/S)}{\mathcal{T}(t,T)} + B\frac{(T-t)}{\mathcal{T}(t,T)} + \bar{\sigma}, \tag{30}$$

where we define the *effective time to maturity (ETM)* as

$$\mathcal{T}(t,T) = \frac{(T-t)}{\left(\frac{m_0+\eta^{3/2}}{m_0+\eta}\right)}. \tag{31}$$

Empirically, we show below that this choice of time variation performs well, but of course other choices may fit the data even better; we have not investigated this here.

4.2 Effective time-to-maturity

We plot $\mathcal{T}(t,T)$ for the model described above in Fig. 8.

Note that the effective time runs faster close to the monthly maturity dates. Moreover, observe that if $T - t < \Delta T$ then $m_0 = 0$ and $\eta = (T-t)/(\Delta T)$ so that

$$\mathcal{T}(t,T) \sim \sqrt{T-t},$$

whereas if $T - t >> \Delta T$ then $\eta << m_0$, and therefore

$$\mathcal{T}(t,T) \sim (T-t).$$

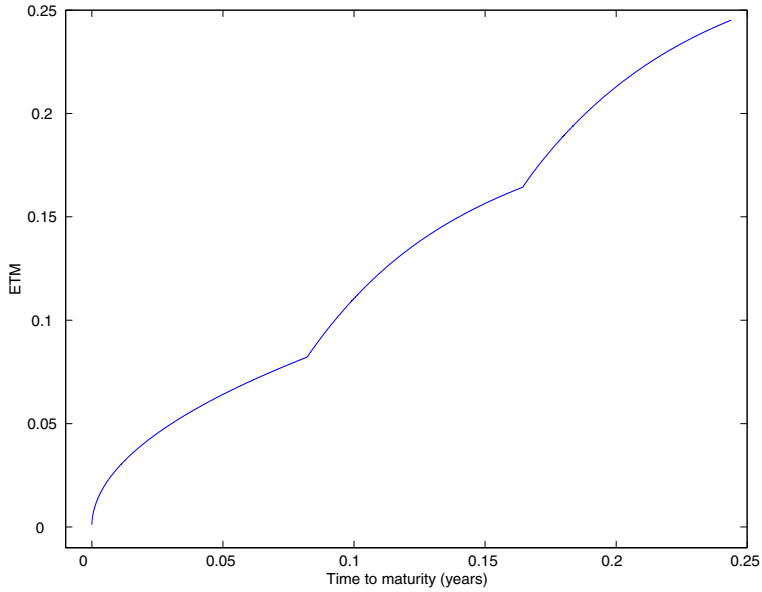

Fig. 8. Effective time to maturity (ETM) as a function of time-to-maturity for a three month option

This gives for the implied volatility

$$I \approx A\frac{\log(K/x)}{\sqrt{T-t}} + B\sqrt{T-t} + \bar{\sigma} \qquad \text{for short maturities}$$

$$I \approx A\frac{\log(K/x)}{(T-t)} + B + \bar{\sigma} \qquad \text{for long maturities.}$$

Thus, for long dated options the implied volatility is approximately affine in LMMR as in (1). For short dated options the implied volatility is approximately affine in $\log(K/x)/\sqrt{T-t}$ as suggested by the observation (3) referred to in Sect. 1.

The kinks in Fig. 8 correspond to different maturity dates T_k, and the interval in between these are here taken to be 30 days. Note that the effect of the maturity-date-dependent rate of mean reversion averages out and become negligible when the option is several maturity intervals from the actual maturity-date. However, in the last maturity period (when the option is close to maturity), it becomes a very pronounced effect.

5 Data revisited using time-dependent theory

The modified theory derived in the last section suggests to fit the implied volatility surface to a function that is affine in the two composite variables

$$\frac{\log(K/x)}{\mathcal{T}(t,T)}; \qquad \frac{(T-t)}{\mathcal{T}(t,T)}.$$

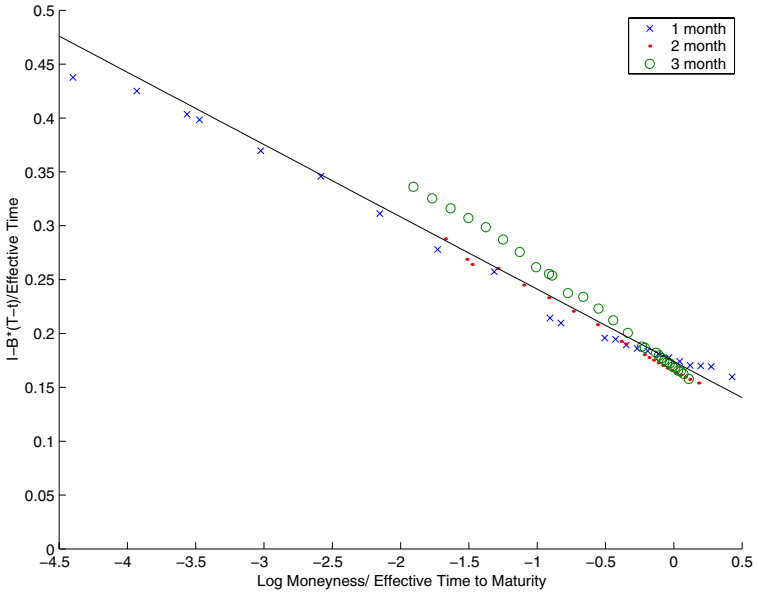

Fig. 9. Maturity-adjusted implied volatility as a function of log-moneyness to ETM on January 12, 2000, for options with at least two and at most 93 days to maturity. Notice the strand for the shortest maturity has been squeezed closer to the strand for the next maturity. Compare Fig. 1

Further, the regression constant should be close to $\bar{\sigma}$ which we estimate from historical S&P 500 returns in 2000 as $\bar{\sigma} = 0.17$.

This is a multiple linear regression, but there are far more points in the direction of the first variable than the second because there are typically dozens of liquid strikes, but only three different maturities (in the data we are restricting ourselves to). Therefore in our estimation procedure we do not place the same emphasis on both variables: we first estimate B maturity-by-maturity on each day by regressing $I - \bar{\sigma}$ on $\log(K/x)/\mathcal{T}(t,T)$ to obtain estimates $B_{(t,T)}$. Since these are analogous to the intercept parameter b in the LMMR fit, they are quite stable (as were the b's in Fig. 2) because they describe the basic level of the surface. Our estimate of B is fixed as the average of these estimates $B = 0.0447$.

This is approximately the average extra level of at-the-money implied volatilities over the historical volatility $\bar{\sigma}$.

Then we fit the *maturity adjusted implied volatility*

$$I - B\frac{(T-t)}{\mathcal{T}(t,T)}$$

as an affine function of $\log(K/x)/\mathcal{T}(t,T)$ by least-squares to estimate A. The residual (which includes the error in the fit and the constant $\bar{\sigma}$ is denoted R).

Figure 9 shows the improvement in the goodness-of-fit and how the effective time-to-maturity has squeezed the previously diverging strands in Fig. 1 together. Next we look at the days before and after the shortest dated options disappear in a particular month. As Fig. 10 shows, there is no longer a dramatic jump in the

estimated slope A. This is a persistent improvement that is also observed around the other options expiration dates in the dataset.

Finally, Fig. 11 shows the daily estimated slope A and the residual R. There is a slight improvement in stability of the slope, as measured by standard deviation, over Fig. 2. Most strikingly, the periodic effect of the expiration dates has been removed.

As mentioned in the introduction similar goodness-of-fit and stability results were obtained with the 1993 S&P 500 dataset constructed in [1].

6 Conclusions

In this article, we have derived and tested a calibration procedure for implied volatility surfaces based on a class of stochastic volatility models. It is able to fit the S&P 500 data well across a few maturities at once (Fig. 9), and also results in parameter estimates that are stable over time (Fig. 11).

The asymptotic analysis of stochastic volatility models with constant parameters leads to the LMMR calibration procedure. When used to analyze S&P 500 implied volatilities, it identifies a significant near-periodic trend, closely associated with the option expiration dates the third Friday of each month.

This observation leads us to modify the class of volatility models to incorporate a time-dependent *periodic* rate of mean-reversion, which is a function of the time to the next expiration date. The asymptotic analysis produces the improved calibration method based on the effective time-to-maturity.

We do not speculate here on a market mechanism that might produce a periodic characteristic speed of mean-reversion in the volatility driving process. A behavioural model in which market makers react to impending "witching dates" and increased market volatility associated with them is a topic of future investigation.

6.1 Longer maturities and slow volatility factor

Finally we comment on the situation with options of longer maturity. In this work, we restricted our study to options of at most three months to maturity. However longer options also trade quite actively in the S&P 500 and other markets. The time-dependent theory introduced here is not designed to help with these options. As shown in Fig. 12, these are not captured well by the one-factor stochastic volatility asymptotic theory.

We conjecture that inclusion of a second slower (mean-reverting) factor in the volatility dynamics (as discussed in Sect. 1.2 and [6]) will help with these longer maturities. Figure 12 suggests that the implied volatility skews of longer maturity options are less sensitive to maturity than those of the shorter options (compare the relative closeness of the two-year skew to the one-year skew against the gap between the three- and six-month skews). We can also see this indirectly in Fig. 5, where the fitted LMMR slope $a_{(t,T)}$ wants to dull the $(T-t)^{-1}$ predicted behaviour of the term-structure at small maturities. In [5], we study the empirical performance of a two-factor stochastic volatility model with a fast and a slow factor.

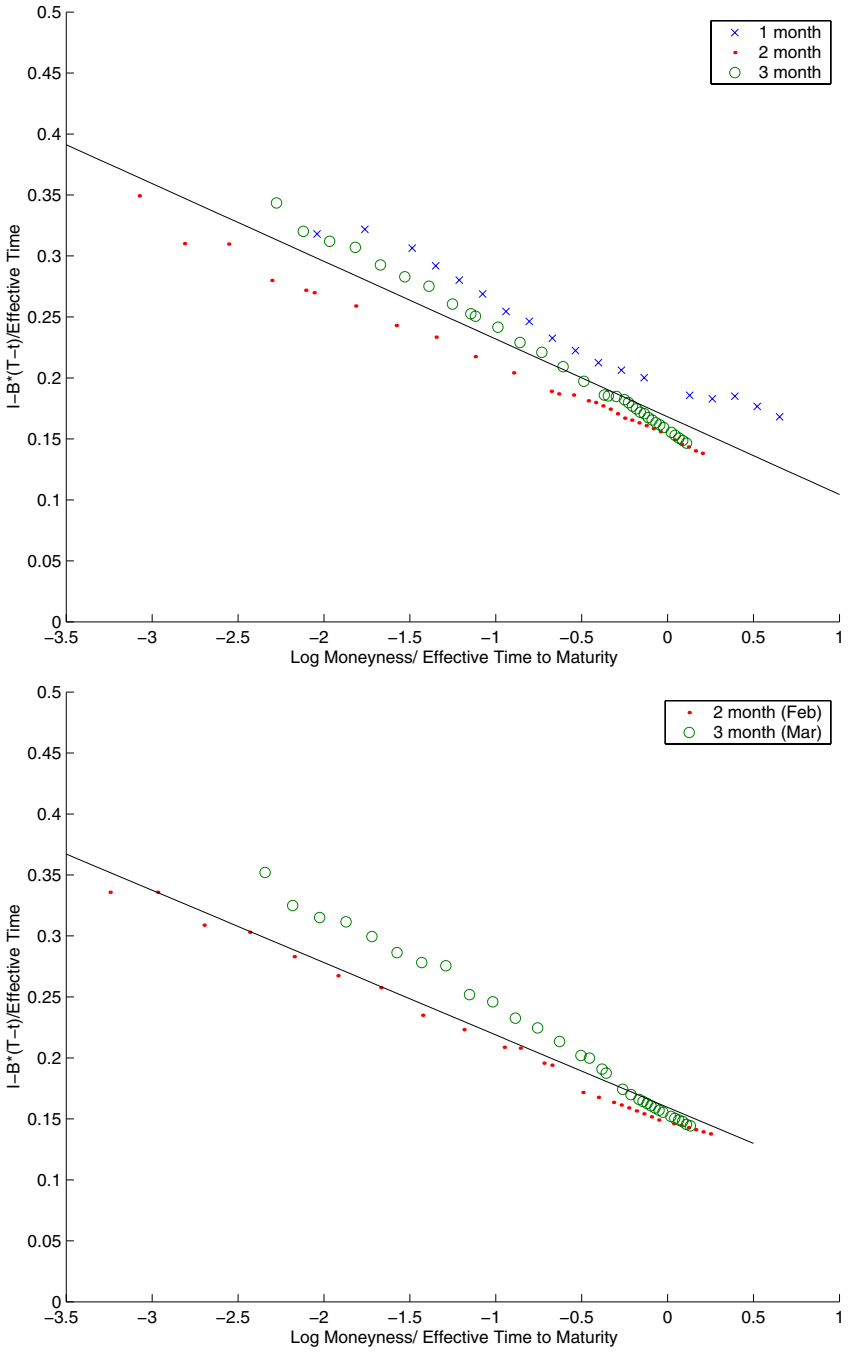

Fig. 10. Maturity-adjusted implied volatility as a function of log-moneyness to ETM on the last day before the shortest dated options disappear (January 18, 2000, top) and after (January 19, 2000, bottom). Notice the strands for each maturity are tighter than with the regular LMMR fit in Fig. 3

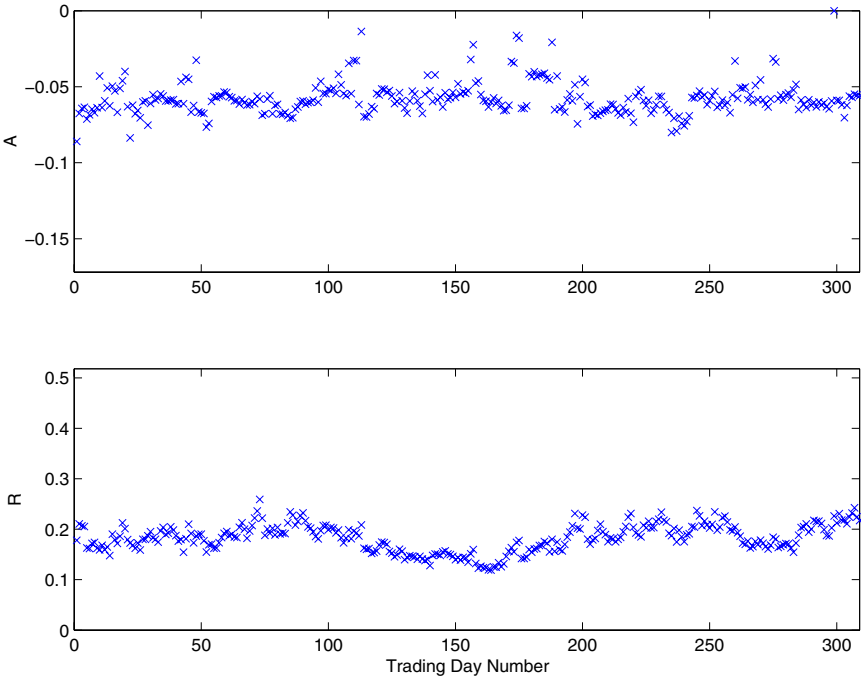

Fig. 11. Daily fitted slope A & residual R for options with at least 2 and at most 93 days to maturity fit to ETM formula. Statistics: Mean $(A) = -0.0577$, Std $(A) = 0.0107$; Mean $(R) = 0.1827$, Std $(R) = 0.0271$

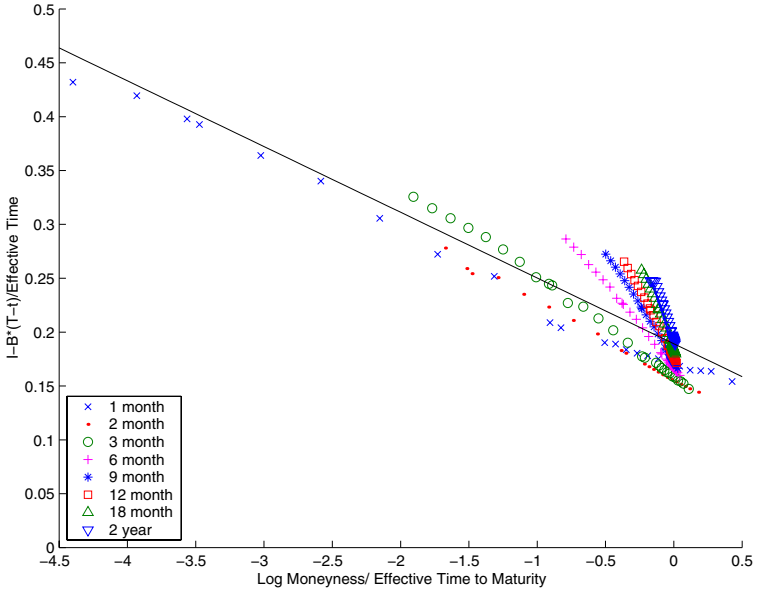

Fig. 12. ETM fit to all maturities (including 6 month, 9 month, 1, 1.5 and 2 year options) on January 14, 2000. Notice the diverging strands for the longer maturities

References

1. Aït-Sahalia, Y., Lo, A.: Nonparametric esimation of state-price densities implicit in financial asset prices. J. Finance **53**(2) April (1998)
2. Alizadeh, S., Brandt, M., Diebold, F.: Range-based estimation of stochastic volatility models. J. Finance **57**(3), 1047–1091 (2002)
3. Andersen, T., Bollerslev, T.: Intraday periodicity and volatility persistence in financial markets. J. Empirical Finance **4**, 115–158 (1997)
4. Ball, C., Roma, A.: Stochastic volatility option pricing. J. Financial Quant. Anal. **29**(4) December, 589–607 (1994)
5. Fouque, J.-P., Papanicolaou, G., Sircar, K.R., Solna, K.: Multiscale stochastic volatility asymptotics. SIAM J. Multiscale Modeling Simulation **2**(1), 22–42 (2003)
6. Fouque, J.-P., Papanicolaou, G., Sircar, K.R., Solna, K.: Short time-scale in S&P 500 volatility. J. Comput. Finance **6**(4), 1–23 (2003)
7. Fouque, J.-P., Papanicolaou, G., Sircar, K.R., Solna, K.: Singular perturbations in option pricing. SIAM J. Appl. Math. **63**(5), 1648–1665 (2003)
8. Fouque, J.-P., Papanicolaou, G., Sircar, K.R.: Derivatives in financial markets with stochastic volatility. Cambridge: Cambridge University Press 2000
9. Heston, S.: A closed-form solution for options with Stochastic Volatility with applications to bond and currency options. Rev. Financial Stud. **6**(2), 327–343 (1993)
10. Hull, J., White, A.: An analysis of the bias in option pricing caused by a stochastic volatility. Adv. Futures Options Res. **3**, 29–61 (1988)
11. Kou, S.: A jump diffusion model for option pricing. Manag. Sci. **48**, 1086–1101 (2002)
12. Lee, R.: Local volatilities under stochastic volatility. Int. J. Theoret. Appl. Finance **4**(1), 45–89 (2001)
13. Lewis, A.: Option valuation under stochastic volatility. Newport Beach, CA: Finance Press 2000
14. Madan, D., Carr, P., Chang, E.: The variance gamma process and option pricing. Eur. Finance Rev. **2**(1), 79–105 (1999)
15. Sircar, K.R., Papanicolaou, G.C.: Stochastic volatility, smile and asymptotics. Appl. Math. Finance **6**(2) June, 107–145 (1999)