



Munich Personal RePEc Archive

Price Deviations of SP 500 Index Options from the Black-Scholes Formula Follow a Simple Pattern

Li, Minqiang

Georgia Institute of Technology

2008

Online at <https://mpra.ub.uni-muenchen.de/11530/>
MPRA Paper No. 11530, posted 15 Nov 2008 04:01 UTC

Price Deviations of S&P 500 Index Options from the Black-Scholes Formula Follow a Simple Pattern

Abstract

It is known that actual option prices deviate from the Black-Scholes formula using the same volatility for different strikes. For the S&P 500 index options, we find that these deviations follow a stable pattern and are described by a simple function of at-the-money-forward total volatility. This implies that the term structure of at-the-money-forward volatilities is sufficient to determine the entire volatility surface. We also find that the implied risk-neutral density is bimodal. The patterns we find are useful in predicting future implied volatilities.

1. Introduction

It is well known (see, e.g., Rubinstein (1994)) that if one uses the Black-Scholes formula or the binomial model with the same volatility to compute option values for different strike prices but the same time to expiration, the computed option prices deviate from quoted prices. Equivalently, for a fixed time to expiration the implied volatility function obtained by inverting the Black-Scholes formula or binomial model is a non-constant function of the strike price. Depending on the shape, this relation between the option strike price and implied volatility is known as the volatility smile or skew. The smile or skew is by now viewed as a basic feature of most option markets, and its presence is apparent in even the most cursory examinations of option implied volatilities. Similarly, it is well known that option implied volatilities differ with the time remaining to expiration. Both this phenomenon and the smile or skew are captured by the implied volatility surface, which is the function mapping the strike price and time to expiration to implied volatility.

This paper presents a new approach for describing the differences between actual option prices and their Black-Scholes values (computed with the same volatility for different strikes), which we term the price deviation function, as well as the associated implied volatility function. It also presents new empirical results about the dynamic behavior of the deviations of actual S&P 500 index (SPX) option prices from their Black-Scholes values and the associated implied volatility function. While both the price deviation and implied volatility functions vary randomly over time, we find that during the period 1996–2002 SPX option prices follow a stable pattern that allows the price deviation function to be well approximated by a simple function of the at-the-money-forward “total volatility,” defined as the product of the at-the-money-forward volatility and the square root of the time to expiration. Thus changes in the slope and curvature of the implied volatility function are almost completely explained by changes in the at-the-money-forward total volatility.

A large literature addresses the volatility skew, much of it consisting of papers proposing models that produce volatility skews, smiles, or both. The simplest such models are the deterministic volatility function models in which the stock price follows a diffusion process with a diffusion coefficient or instantaneous volatility rate given by a deterministic function of the stock price and time.¹ Other approaches to capturing the smile or skew include superimposing jumps on the price process of the underlying asset, allowing for stochastic volatility by introducing an additional process to

¹The constant elasticity of variance model of Cox (1976) is the earliest model that fits this description. Derman and Kani (1994), Dupire (1994), and Rubinstein (1994) propose models in which the volatility rate is an unspecified function of the stock price and time that is chosen to match the volatility surface. A somewhat related set of papers including Melick and Thomas (1997), Ait-Sahalia and Lo (1998), and Campa, Hang, and Reider (1998) use various approaches to develop estimates of the state price density that are consistent with observed option prices. Another group of papers oriented primarily to practitioners does not attempt to explain the skew, but rather focuses on describing it. These include Derman, Kani and Zou (1996), Heynen (1993), Rosenberg (2000), and Tompkins (2001).

describe the movement of the instantaneous volatility, or both.² Recently, there is also a growing literature on using the Lévy process to model the price process, thus allowing for more flexibility in the behavior of the jumps.³ However, evidence on the empirical performance of these models is mixed. Using data on the prices of SPX options, Dumas, Fleming, and Whaley (1998) perform out-of-sample empirical tests of the deterministic volatility functions and find that they are outperformed by an *ad hoc* implementation of the Black-Scholes model that smooths the implied volatility function across strike prices and time to expiration and uses the estimated implied volatilities to calculate option prices one week later. Brandt and Wu (2002) reach similar conclusions using data on FT-SE 100 index options. Bakshi, Cao, and Chen (1997) and Bates (1998) compare the pricing and hedging errors of various models and obtain results generally consistent with the idea that a model with both discontinuous “jumps” and stochastic volatility performs better than the Black-Scholes model using the same volatility for different strike prices. Jackwerth and Rubinstein (2001) run an extensive horse race and find that, although the Black-Scholes model is outperformed by both deterministic and stochastic volatility models, the variability of the pricing errors is so large that it is not possible to rank the competing models. Moreover, it is difficult to reconcile the slope (and changes in the slope) of the implied volatility function with the parameters of the empirical distribution of returns and reasonable estimates of investors’ risk aversion (Bakshi, Cao, and Chen (1997), Bates (2000), Jackwerth (2000), and Bondarenko (2001)).

A common feature of the deterministic volatility function, stochastic volatility/jump-diffusion, and Lévy process approaches is that they attempt to find a stochastic process that generates a risk-neutral density consistent with the implied volatility function. An alternative approach to explain the skew is that option prices may differ from their no-arbitrage values due to the costs and difficulty of executing the dynamic trading strategies that appear in the derivations of options’ no-arbitrage values.⁴ Despite the lack of consensus about the cause or causes of the skew, certain stylized facts

²Examples of these approaches include the jump-diffusion models of Merton (1976) and Bates (1991), the stochastic volatility models of Hull and White (1987), Scott (1987), Wiggins (1987), Melino and Turnbull (1990, 1995), Stein and Stein (1991), and Heston (1993a), the combined stochastic volatility jump-diffusion models of Bates (1996) and Bakshi, Cao, and Chen (1997), and the general affine model analyzed in Duffie, Pan, and Singleton (2000).

³Examples of this approach include the variance-gamma model of Madan and Seneta (1990), Madan and Milne (1991), Madan, Carr, and Chang (1998), the log-gamma model of Heston (1993b), the hyperbolic model of Eberlein, Keller, and Prause (1998), the α -stable models of Janicki, Popova, Ritchken, and Woyczynski (1997), Popova and Ritchken (1998), and Hurst, Platen, and Rachev (1999), the CGMY model of Carr, Geman, Madan, and Yor (2002), and the log stable models of McCulloch (1987, 1996), and Carr and Wu (2003a). Other studies on applications of Lévy processes include Geman, Madan, and Yor (2001), Carr and Wu (2003b), and Carr and Wu (2004).

⁴Bollen and Whaley (2004) and Whaley (2002; Section 5.4) present evidence that the differences in the implied volatility functions of index and individual equity options are not explained by differences in their return distributions. Bollen and Whaley (2004) also document that movements in the implied volatilities are correlated with net buying pressure. Longstaff (1995) and Bates (1996) also find that implied volatilities are related to the level of trading activity. Jackwerth (2000) attempts to recover risk aversion functions from the prices of SPX options and concludes that the prices are “irreconcilable with a representative investor,” while Bondarenko (2001) examines the prices of out-of-the-money S&P 500 futures options and concludes that the market is inefficient. Buraschi and Jackwerth (2001; p. 523) remark that their results regarding spanning in options markets are consistent with away-from-the-money

about it are known. First, for index options the implied volatility is generally a decreasing function of the strike price, i.e. the implied volatility function is negatively sloped. Second, for individual equity options the implied volatility function is generally convex and more symmetric or “smile-shaped” than the implied volatility function for index options, and it is less negatively sloped.⁵ Third, both the level and slope of the implied volatility function for index options appear to change randomly over time.⁶

This paper presents new stylized facts about the volatility skew. We begin by expanding the deviations of actual prices from Black-Scholes values using parabolic cylinder functions. It suffices to use only two such functions, corresponding to the odd and even components of the price deviations. For a large neighborhood of the forward price of the underlying asset, these two components can be interpreted as the slope and curvature of the price deviation function, and can be mapped to the slope and curvature of the implied volatility function. Using data on SPX options for the period 1996–2002, the coefficients of both the odd and even components can be well fit as quadratic functions of the at-the-money-forward total volatility. That is, changes in the volatility skew are almost completely explained by changes in the at-the-money-forward total volatility. This relation is stable during the period 1996–2002. In particular, the term structure of at-the-money-forward volatilities determines the entire volatility surface. This finding of stability in the price deviation and implied volatility functions stands in marked contrast to the view of a large fraction of the literature. One difference is that most previous work has not focused on the slope and curvature of the price deviation function. Perhaps more importantly, the slope and curvature of both the price deviation function and the implied volatility function depend upon the total volatility $\sigma_F\sqrt{T}$, and previous authors have not focused on this quantity.

Stability in the volatility skew implies corresponding stability in at least some features of the risk-neutral density. As shown by Breeden and Litzenberger (1978), the risk-neutral density can be obtained by differentiating the value of a European call or put option with respect to the strike price (and adjusting for the discount factor). Carrying out this computation using the fitted option

index options being used by a different clientele than their at-the-money counterparts. More recently, Gârleneau, Pedersen, and Poteshman (2006) describe a model in which limits to arbitrage cause option prices to deviate from their no-arbitrage values, and present empirical evidence consistent with that model.

⁵The implied volatility function for the index options has not always been negatively sloped. Rubinstein (1994) reports that this phenomenon appears to date from the October 1987 stock market crash. Whaley (1986) and Sheikh (1991) document changes in the implied volatility function of index options during the period 1983–1985.

⁶Studies documenting those styled facts include: Rubinstein (1994), Derman (1999a, 1999b), Gemmill and Kamiyama (2000), Bakshi, Kapadia, and Madan (2001), Dennis and Mayhew (2001), and many others. Systematic studies of the skew include Bollen and Whaley (2004), who examine changes in the implied volatilities of the S&P 500 index options, and Gemmill and Kamiyama (2000), who study implied volatilities and volatility skews of the S&P 500, FT-SE 100, and Nikkei 225 index options over the period 1985–1995 and conclude that “skewness changes markedly from day to day” (p. 11). Jackwerth and Rubinstein (1996, p. 1619) take random changes in the implied volatility function for granted, writing that “. . . the shape of the smile obviously changes from day to day, . . .”. Similarly, Derman (1999a, 1999b) studies how the volatility skew in the SPX options varies with the index level, and in fact identifies seven different “regimes of volatility” during the period from September 1, 1997 to November 2, 1998.

prices, we find that the implied risk-neutral density of index returns is bimodal. This finding has implications for models with stochastic volatility or jump components, as it provides a stylized fact that proposed models must capture and rules out models that cannot produce this property. The choice of parabolic cylinder functions also has nice consequences for the implied distribution, as it, for example, integrates to one exactly and preserves put-call parity. The parametric form of the distribution should also be useful to researchers who want to introduce skewness and kurtosis to the return distribution. More generally, description of the deviations of actual option prices from Black-Scholes values in terms of parabolic cylinder functions is likely to be useful for empirical analysis because it provides a low-dimensional (i.e., small number of parameters) description of the price deviation and implied volatility functions.

In related work (Li and Pearson 2006), they find that our skew model is useful in predicting future volatilities. Specifically, they update the “horse race” among competing models originally carried out by Jackwerth and Rubinstein (2001) using more recent option data set and more recent models. In agreement with Jackwerth and Rubinstein (2001), they find that trader rules dominate mathematically more sophisticated models in predicting option prices. However, the simple pattern in the implied volatility skew can be used to further improve the performances of the trader rules. Perhaps surprisingly, the naive Black-Scholes model beats all mathematically more sophisticated models after adjusting the skew shapes using the simple pattern.

The paper is organized as follows. Section 2 describes the deviations of actual option prices from Black-Scholes values and presents the series expansion. Using this approach, Section 3 shows that during the period 1996–2002 the deviations of actual prices from Black-Scholes values are well described by a stable function of the at-the-money-forward total volatility. Section 4 connects the price deviation function to the implied volatility function, and shows that the implied volatility surface also can be modeled as a function of the at-the-money-forward volatility. Then, Section 5 works out the implications for the implied risk-neutral density. Finally, Section 6 shows that the simple function provides good conditional predictions of option prices. Section 7 briefly concludes.

2. Deviations of actual option prices from Black-Scholes values and the series expansion

We begin by modeling deviations of actual option prices (closing bid-ask midpoints) from Black-Scholes values computed using at-the-money-forward volatilities. Figure 1 illustrates the deviations observed at the close of trading on April 26, 2002 for the SPX options expiring in June 2002 as a function of their strike prices. The Black-Scholes values are computed using the at-the-money-forward volatility, defined as the implied volatility of the straddle with strike price equal to the

forward price of the index. Because none of the traded options had strike prices exactly equal to the forward price, the at-the-money-forward volatility is interpolated using cubic splines from the implied volatilities of the straddles with strike prices surrounding the forward price, which on this day is estimated to be $F = \$1,072.49$. Due to put-call parity the deviations for the puts are essentially identical to those for the calls, and the deviations are zero at the strike price equal to the forward price. As is known to be typical for the SPX options, for low strike prices the actual prices exceed the Black-Scholes values, while for high strike prices the actual option prices are less than the Black-Scholes values. One can see in Figure 1 that the deviations can be greater than \$4 (at strike $X = \$975$) and less than $-\$2$ (near $X = \$1,150$). For comparison, the bid-ask spreads of both puts and calls at both of these strikes were \$2. By simple model-independent arguments, we expect that the price deviations should approach zero as the strike prices approach infinity or zero, or equivalently as $\ln X$ approaches $\pm\infty$. Of course, options with strike prices of infinity and zero are not traded, but the curves plotted in Figure 1 are consistent with this prediction.

Although the prices are the data directly observed from the market, many researchers and practitioners focus on implied volatilities, and the implied volatility function is a widely used representation of option prices. While working with implied volatilities is equivalent to working with prices, working with prices overcomes several problems with implied volatilities. First, for away-from-the-money options Black-Scholes option values are insensitive to changes in volatility. As a result, small changes in option prices lead to large changes in implied volatilities, suggesting that implied volatilities of away-from-the-money options may not be reliable and that any fitting procedure should place less weight on the implied volatilities of such options. Working in terms of option prices is a natural way to accomplish this “deweighting” of away-from-the-money implied volatilities. Second, dimensional analysis argues against separating the volatility σ and time to expiration τ from the combination $\sigma^2\tau$ that enters the Black-Scholes formula. It is not appropriate to treat implied volatilities for different times to expiration on equal footing. Finally, although there are some theoretical arguments (e.g., Hodges 1996, Gatheral 1999) restricting the asymptotic behavior of the implied volatility as the strike X approaches zero and infinity, such restrictions usually take the form of inequalities for the derivatives of implied volatility with respect to moneyness. Because there are usually few away-from-the-money options available, it is difficult to calculate these derivatives reliably, and these inequalities are awkward to implement. On the other hand, it is much more natural and convenient to restrict the asymptotic behavior of option prices. For these reasons we work with the deviations of actual option prices from Black-Scholes values.⁷

⁷The state-price density is yet another mechanism through which one might study the volatility structure (see, e.g., Aït-Sahalia and Duarte 2003). A difficulty with this approach is that there are shape restrictions on the option value functions, e.g., the values of both calls and puts must be convex functions of the strike, and the risk-neutral density implied from the observed option prices must integrate to one. Imposing the convexity constraint in estimation (for example, using smoothing splines or kernel regression) is not easy and complicated procedures are sometimes

Besides overcoming the shortcomings of using implied volatilities, there are several advantages of studying the price deviations directly. First, once we have fit the price deviations we can easily use Taylor series expansion to obtain an estimated implied volatility function. In addition, the implied risk-neutral density can be easily obtained by differentiating the fitted prices. Further, although our procedure does not guarantee the absence of convexity violations, it turns out that such violations are minor and occur only at strike prices that are well away from the forward price.

2.1. Series expansion of the price deviations

The Black-Scholes formulas give the values of call and put options C_{BS} and P_{BS} as

$$C_{BS} = Fe^{-r\tau} \left(N(d_1) - \frac{X}{F} N(d_2) \right), \quad \text{and} \quad P_{BS} = Fe^{-r\tau} \left(-N(-d_1) + \frac{X}{F} N(-d_2) \right), \quad (1)$$

where S is the current stock price or index value, r is the interest rate, δ is the dividend yield, σ is the volatility, τ is the time to expiration, X is the strike price, $F = S \exp((r - \delta)\tau)$ is the forward price, and d_1 and d_2 are given by

$$d_1 = \frac{\ln(F/X)}{\sigma\sqrt{\tau}} + \frac{1}{2}\sigma\sqrt{\tau}, \quad \text{and} \quad d_2 = \frac{\ln(F/X)}{\sigma\sqrt{\tau}} - \frac{1}{2}\sigma\sqrt{\tau}. \quad (2)$$

We call $\sigma\sqrt{\tau}$ the total volatility.

To compute the price deviations, we need to pick a volatility to use in the Black-Scholes formula. We choose $\sigma_F(t, \tau)$, the at-the-money-forward volatility of an option with time to expiration τ at calendar date t .⁸ We make this choice because options with strikes close to the forward price are usually actively traded with relatively small bid-ask spreads so that σ_F can be computed accurately. In addition, usually the forward price lies in the middle of the available strike prices because options with strikes close to the new index level are introduced for trading as the index level changes. Thus σ_F is obtained by interpolation (which is more reliable) rather than extrapolation.

necessary to satisfy the constraint. A further difficulty is that theoretical arguments do not restrict the tail behavior of the implied risk-neutral density. Thus, this approach provides little guidance in valuing options with strikes outside the range of strike prices for which implied volatilities are observed.

⁸Because usually none of the traded options have strike prices equal to the forward price, the at-the-money-forward volatility is interpolated from the implied volatilities of the two straddles with strike prices closest to the forward price, which is approximately equal to the average of the implied volatilities of the call and put. The computation of the forward price presents another problem. As argued in Longstaff (1995), reported index values and futures prices are unreliable estimators of the true index value, and yields on Treasury bills may not be the appropriate risk-free discount rate. Also, the observed current dividend yield might not reflect the expected dividend rate of investors. We use Shimko's (1993) procedure to overcome these problems by using the put-call parity relation. Because this procedure uses only price quotes from a single market captured at the same instant, it avoids the problems associated with asynchronous prices.

We measure the relative moneyness of an option using the quantity

$$d(t, \tau, X) = \frac{\ln(F(t, \tau)/X)}{\sigma_F(t, \tau)\sqrt{\tau}}, \quad (3)$$

which we call the total-volatility-adjusted moneyness or often simply moneyness. For example, a value of $d = 2$ roughly means that the option is two standard deviations in (for a call) or out (for a put) of the money. For brevity, we often suppress the arguments in $d(t, \tau, X)$ and simply write d , and similarly for $F(t, \tau)$ and $\sigma_F(t, \tau)$. Let $Y_A(d)$ denote the actual market price of a call or put option with moneyness d , let $Y_{BS}(d, \sigma_F)$ denote the Black-Scholes value computed using the at-the-money-forward volatility, and let

$$y = y(t, \tau, X) \equiv \frac{Y_A(d) - Y_{BS}(d, \sigma_F)}{F} \quad (4)$$

denote the difference between the actual price and the Black-Scholes value, scaled by the forward price. Scaling by the forward price is sensible because the Black-Scholes formula is homogenous of degree one in the forward price. This scaling also makes y a dimensionless quantity.

We expand the deviations from Black-Scholes prices using parabolic cylinder function series. These functions are given by

$$D_n(z) = (-1)^n e^{z^2/4} \frac{d^n}{dz^n} \left(e^{-z^2/2} \right), \quad n = 0, 1, 2, \dots \quad (5)$$

Expressions for the lowest order functions are:

$$D_0(z) = e^{-z^2/4}, \quad D_1(z) = ze^{-z^2/4}, \quad (6)$$

$$D_2(z) = (z^2 - 1)e^{-z^2/4}, \quad D_3(z) = (z^3 - 3z)e^{-z^2/4}. \quad (7)$$

A number of considerations lead us to this choice of parabolic cylinder functions. First, the functions $D_n(z)$ ($n = 0, 1, 2, \dots$) constitute a complete orthogonal function series on $(-\infty, +\infty)$. A function $f(z)$ that: (i) has continuous first and second order derivatives; and (ii) goes to zero as $|z| \rightarrow \infty$ can be expanded using $D_n(z)$ as (Wang and Guo (1989))

$$f(z) = \sum_{n=0}^{\infty} a_n D_n(z), \quad \text{where} \quad a_n = \frac{1}{n! \sqrt{2\pi}} \int_{-\infty}^{\infty} f(z) D_n(z) dz. \quad (8)$$

The parabolic cylinder functions can be used as a basis for deviations from Black-Scholes values because it is reasonable to think that the deviations from Black-Scholes values satisfy condition (i) above, and will also satisfy (ii) if the argument z is chosen to be a linear function of $\ln X$, i.e. if

$z = a + b \ln X$ for some a and b . Thus, deviations from Black-Scholes values should permit an expansion in terms of the parabolic cylinder functions. Further, the parabolic cylinder functions are a natural choice of a function series because each function $D_n(z)$ ($n = 0, 1, 2, \dots$) goes to zero as $|z| \rightarrow \infty$, suggesting that only a few terms will be needed. As shown below, this turns out to be the case. Parabolic cylinder functions also facilitate the connections between price deviations and implied volatility functions and between price deviations and the implied risk-neutral density, as will be evident later on.

Temporarily fix a calendar date t and a time to maturity τ . Expanding the scaled price deviations in equation (6) in terms of the parabolic cylinder functions, we have

$$y \equiv \frac{Y_A(z) - Y_{BS}(z, \sigma_F)}{F} \approx \sum_{n=0}^{N-1} a_n(t, \tau) D_n(z). \quad (9)$$

where N is the number of functions we use and a_n , for $n = 0, 1, \dots, N - 1$, are the expansion coefficients. This equation is for options with different strike prices for a fixed (t, τ) pair. Based on the evidence in Section 3 below, we first drop the terms of third and higher order. Next, we choose the parameter z to be $z = \sqrt{2}d$, where d is the total-volatility-adjusted moneyness defined in equation (3) above. As pointed out above, choosing z to be a linear function of $\ln X$ automatically captures the asymptotic behavior of the price deviations. The factor $\sqrt{2}$ has the effect of making the exponential factor $\exp(-z^2/4)$ in the expressions for the parabolic cylinder functions equal to $\exp(-d^2/2)$, which is similar to a term that appears in the derivatives of the Black-Scholes formula. As will be seen below, this is convenient when we relate the price deviations to implied volatilities. An advantage of using d instead of d_1 is that when the strike price equals the forward price, the quantities d , the left-hand side of (9), and $D_1(\sqrt{2}d)$ are zero. This constrains $a_0 = a_2$ on the right-hand side of (9), enabling the use of one fewer expansion function.

After making these choices, the model of price deviations from Black-Scholes values is then

$$y \equiv \frac{Y_A(d) - Y_{BS}(d, \sigma_F)}{F} = a_1 \widetilde{D}_1(\sqrt{2}d) + a_2 \widetilde{D}_2(\sqrt{2}d), \quad (10)$$

where

$$\widetilde{D}_1(z) = D_1(z) = ze^{-z^2/4}, \quad (11)$$

$$\widetilde{D}_2(z) = D_0(z) + D_2(z) = z^2 e^{-z^2/4}. \quad (12)$$

This procedure produces for each (t, τ) pair two expansion coefficients a_1 and a_2 . Empirical evidence presented in Section 3 below shows that the coefficients a_1 and a_2 themselves are well fit by quadratic

functions of the at-the-money-forward total volatility $\sigma_F(\tau)\sqrt{\tau}$:

$$a_1 = (\alpha_1\sigma_F^2(\tau)\tau + \beta_1\sigma_F(\tau)\sqrt{\tau})\exp(-r\tau), \quad (13)$$

$$a_2 = (\alpha_2\sigma_F^2(\tau)\tau + \beta_2\sigma_F(\tau)\sqrt{\tau})\exp(-r\tau), \quad (14)$$

where $r = r(t, \tau)$ is the risk-free interest rate. The details of the above fitting are explained below. The coefficients α_1 , β_1 , α_2 , and β_2 are stable over time in the sense that the same four parameters fit the volatility skew for every calendar date and option expiration date over the period January 1996–September 2002; the only parameter that varies with the calendar date and time to expiration is the at-the-money-forward volatility $\sigma_F(t, \tau)$. Loosely, the slope and curvature of the implied volatility function are determined by the at-the-money-forward total volatility $\sigma_F(t, \tau)\sqrt{\tau}$. Because on each calendar date we compute the at-the-money forward volatility for each option expiration date, equation (10) is a model of the shape of the volatility skew or smile, and not a model of the level of implied volatility. In Subsection 3.3.4 we combine equations (10) through (14) and fit the implied volatility surface in a single step.

2.2. Interpretation of the parabolic cylinder functions

The two functions \widetilde{D}_1 and \widetilde{D}_2 that appear in the expansion (10) can be interpreted as the odd and even components of the deviations of actual option prices from Black-Scholes values. Figure 2 shows schematically how the two components sum to the price deviation function. The solid line shows the deviations of option prices from Black-Scholes values computed using the at-the-money-forward volatility as a function of the volatility-adjusted moneyness d , while the lines with circles and crosses show the odd and even components, respectively. Further, around $d = 0$, that is around the forward price, the two components can be interpreted as the slope and curvature of the price deviation function.

The first of the two components is clearly the more important. When we estimate the coefficients a_1 and a_2 separately for each calendar date as described in the next section, the estimated values of a_1 are positive for the entire sample period, ranging from zero to about 0.05. The median value of a_1 is 0.0082 while the mean is 0.0094. The estimated values of a_2 are both positive and negative, mostly in the range of -0.005 to 0.005 . The median of the absolute value of a_2 is 0.00073 and the mean is 0.00085. When $\ln(F/X)/(\sigma_F\sqrt{\tau}) = 1$ (one standard deviation away from the money) we have $z = \sqrt{2}d = \sqrt{2}$ and taking the median values of a_1 and a_2 , the odd component is about $0.0082/(0.00073 \times 2) = 5.62$ times larger than the even component. For options that are close to the money (small d), the relative importance of the odd component is even greater.

3. Fitting the price deviations

In addition to the choice of expansion functions, the model in equations (10)–(14) reflects three further key choices. These are first, the decision to omit terms of third and higher order; second, the decision to set $a_0 = a_2$; and third, the choice to specify the expansion coefficients a_1 and a_2 as functions of at-the-money-forward total volatility. In this section we first describe the data we used in this paper and then present the data analysis that motivated and supports these choices. Subsection 3.4 below provides an empirical analysis of the fit of the model we used. In the next two sections, we then connect our price deviation fitting with implied volatility fitting and the implied risk-neutral density.

3.1. Data

The SPX option data used in this paper come from OptionMetrics LLC. The data include end-of-day bid and ask quotes, implied volatilities, open interest, and daily trading volume for the SPX (S&P 500 index) options traded on the Chicago Board Options Exchange from January 4, 1996 through September 13, 2002. The data also include daily index values and estimates of dividend yields, as well as term structures of zero-coupon interest rates constructed from LIBOR quotes and Eurodollar futures prices. The period from January 1996 until September 2002 includes prices for options with 88 different expiration dates, including eight expiration dates that fall after September 13, 2002.

We use the bid-ask average as our measure of the option price, and use all options that satisfy the following criteria. First, we require that the option moneyness F/X satisfy $0.8 \leq F/X \leq 1.2$ and that the total-volatility-adjusted moneyness $d = \ln(F/X)/(\sigma_F\sqrt{\tau})$ satisfy $-3 \leq d \leq 3$, where F is the forward price of the index, X is an option strike price, τ is the time to expiration, and σ_F is at-the-money-forward volatility. Second, we use only options for which both bid and ask prices are strictly positive. Finally, we require that the time to expiration be greater than or equal to ten business days. However, we keep all the long-term options with times to expiration up to 2 years. We also require that for each strike price, there is a matching pair of put and call. These different criteria overlap considerably. Imposing all criteria, we have 456,280 different option prices (228,140 calls and 228,140 puts).

These criteria are motivated by several considerations. First, they eliminate the least liquid options. Second, they eliminate options for which the bid-ask average is unlikely to be a useful approximation of the option value. Specifically, many of the away-from-the-money options are bid at zero and offered at \$0.5, with a bid-ask average equal to \$0.25, independent of the time to expiration or strike price. Finally, we impose the requirement that τ be greater than or equal to

ten business days because our procedure requires prices for a range of different strikes, and for very short times to expiration often only a few options satisfy the first two criteria.

3.2. Order of the expansion

Because parabolic cylinder functions form a basis for a function space, the price deviations can be captured completely if sufficiently many functions are used. The question is then how many functions are needed to capture the price deviations with a high degree of accuracy. There is reason to think that the number of required expansion functions will be small, as Figure 2 suggests that the two functions $\widetilde{D}_1(z)$ and $\widetilde{D}_2(z)$ themselves explain a large fraction of the deviations from Black-Scholes values.

To explore this issue, let us temporarily fix a calendar day t and time to expiration τ and consider all the options with different strikes. Using k to index the options with different strikes, we have

$$y_k \equiv \frac{Y_A(z_k) - Y_{BS}(z_k, \sigma_F)}{F} \approx \sum_{n=0}^{N-1} a_n D_n(z_k), \quad (15)$$

where

$$z_k = \sqrt{2}d_k = \frac{\sqrt{2} \ln(F/X_k)}{\sigma_F \sqrt{\tau}}. \quad (16)$$

Since we are fixing t and τ for now, we have omitted the dependence of t and τ in equation (15), that is, we have written F for $F(t, \tau)$, σ_F for $\sigma_F(t, \tau)$, a_n for $a_n(t, \tau)$ and y_k for $y(t, \tau, X_k)$, etc. If we append an error term to the right-hand side, equation (15) can be thought of as an ordinary least squares regression, with dependent variable y_k and independent variables $D_n(z_k)$.

For each calendar day-expiration date pair, we estimate equation (15) using ordinary least squares to obtain the expansion coefficients a_n . We begin by looking at the squared errors. If there are K available option series on day t , define the sum of squared errors in the usual way:

$$SSE = \sum_{k=1}^K \left(y_k - \sum_{n=0}^{N-1} a_n D_n(\sqrt{2}d_k) \right)^2. \quad (17)$$

Defining the total sum of squares by $SST = \sum_{k=1}^K (y_k - \bar{y})^2$, where $\bar{y} = \sum_{k=1}^K y_k / K$, the coefficient of determination R^2 is given by $R^2 = 1 - SSE/SST$.

We illustrate the fit for different choices of N using the April 26, 2002 prices of the options expiring in June 2002. Table 1 gives the sums of squared errors (SSE) and coefficients of determination R^2 at the best fitting parameter values for $N = 1, 2, \dots, 7$. Similar patterns of sums of squared errors and R^2 's with respect to N are found on other dates. Because the coefficients are computed using least squares, computational constraints do not place any limit on the number of

functions N . However, there is no benefit to choosing N to be larger than six. First, doing this has almost no impact on the sum of squared errors and the R^2 . Second, since we are not expanding the entire price deviation function on each date (we typically have about 20 data points for each date), including more expansion functions may well result in over-fitting the data.

The discreteness of quoted option prices inherently limits the ability to obtain an exact fit with a smooth function, suggesting it may be the case that the choice of $N = 6$ already overfits the data, and that a choice as small as $N = 3$ might be reasonable. Further, the fact that the left-hand side of equation (15) is zero when $d_k = 0$, and the parabolic cylinder functions D_0, D_2, D_4, \dots , are even functions places a constraint on the coefficients of these functions. In the case of $N = 3$ this constraint takes the form $a_0 \approx a_2$, where the approximate equality comes from the fact that the expansion does not provide an exact fit. Figure 3 illustrates that this approximate equality is satisfied in the data when $N = 3$. In particular, Figure 3 shows the time series of the estimated expansion coefficients a_0 , a_1 , and a_2 for each day during the five month period prior to the expiration of the June 2002 options. The estimated coefficient a_1 on the odd function D_1 declines steadily as the time to expiration decreases from around 0.4 to zero, while the coefficients a_0 and a_2 on the even functions D_0 and D_2 are approximately constant and equal to each other. In addition, with this choice of $N = 3$ the R^2 's are generally above 0.95 for each calendar day, declining below this value only when the time to expiration becomes short. Table 1 also reports the fit for $N = 3$ but with constraint $a_0 = a_2$ (column 2') and we see that the R^2 is almost as large as that for $N = 3$ without the constraint. Similar results are obtained for the other options with the same expiration date and the options with different expiration dates.

For these reasons we chose $N = 3$ and constrain $a_0 = a_2$, and thus use the expansion

$$y_k \equiv \frac{Y_A(d_k) - Y_{BS}(d_k, \sigma_F)}{F} = a_1 \widetilde{D}_1(\sqrt{2}d_k) + a_2 \widetilde{D}_2(\sqrt{2}d_k), \quad (18)$$

where

$$\widetilde{D}_1(z) = D_1(z) = ze^{-z^2/4}, \quad (19)$$

$$\widetilde{D}_2(z) = D_0(z) + D_2(z) = z^2 e^{-z^2/4}. \quad (20)$$

Figure 4 shows the time series of the expansion coefficients a_1 and a_2 obtained by estimating equation (18). As expected, it is similar to Figure 3, with the only notable difference being that the time series of estimated expansion coefficients are smoother. Although (18) involves one less function, the fit remains excellent, with R^2 's well above 0.92 for most of the calendar days.

3.3. Exploratory data analysis: Patterns in the expansion coefficients

The estimation of equations (15) and (18) by ordinary least squares discussed above is done separately for each calendar day and expiration date, i.e. the procedure imposes no restriction on the coefficients a_1 and a_2 across different days and expiration dates. Estimating the expansion (18) separately for each calendar date-expiration date pair and imposing the constraint $a_0 = a_2$ results in a total of 12,381 estimated coefficients a_1 and a_2 . We seek stable relations that are satisfied by the coefficients a_1 and a_2 .

We begin by using dimensional analysis as a rough guide, which has proven to be very useful in physics, engineering, and even pharmacology. The left-hand side of the expansion (18) consists of price deviations divided by the forward price, and thus is dimensionless. Through the choice of $z = \sqrt{2} \ln(F/X) / (\sigma_F \sqrt{\tau})$ the quantities $\widetilde{D}_1 = ze^{-z^2/4}$ and $\widetilde{D}_2 = z^2 e^{-z^2/4}$ also are dimensionless. Thus, the coefficients a_1 and a_2 must also be dimensionless, suggesting that they can only be functions of dimensionless quantities. The available dimensionless quantities include the moneyness F/X , the discount factor $\exp(-r\tau)$ or equivalently $r\tau$, the total dividends $\delta\tau$, the total volatility $\sigma_F \sqrt{\tau}$, and d , d_1 and d_2 . Among these, we can identify the following independent quantities: d , $\sigma_F \sqrt{\tau}$, $r\tau$ and $\delta\tau$. We disregard d , because the expansion should ensure that the dependence of the price deviations on d is captured via the expansion functions, similar to the separation of variables method in solving partial differential equations. Further, it seems unlikely that the dividend yield would vary enough to drive the variation in the coefficients a_1 and a_2 . Thus, dimensional analysis suggests models of the form

$$a_n = f_n(\exp(-r\tau), \sigma_F \sqrt{\tau}), \quad n = 1, 2, \quad (21)$$

where the f_n 's are functions to be determined.

Table 2 confirms these tentative conclusions. It shows the correlations between the expansion coefficients a_1 and a_2 , the at-the-money-forward total variance $\sigma_F^2 \tau$, the time remaining to expiration τ , the volatility σ_F , and some other variables. The correlations between a_1 and both τ and σ_F are high, with values of 0.7364 and 0.5248, respectively. However, the correlation between a_1 and the total variance $\sigma_F^2 \tau$ is even higher, 0.9277. These correlations are larger than the correlations between a_1 and any of the other variables.⁹ While the magnitudes are smaller, a similar pattern is found in the correlations between a_2 and the various variables. In particular, the correlation between a_2 and $\sigma_F^2 \tau$ of -0.6521 is the largest of the correlations involving a_2 . In the next subsection, we take a closer look at the odd and even components in turn.

⁹Further, in untabulated results we find that the relatively large correlation between F and a_1 is not stable across the different years in the sample period.

3.3.1. The odd component

In the last section we preannounced that the model $a_1 = (\alpha_1 \sigma_F^2 \tau + \beta_1 \sigma_F \sqrt{\tau}) e^{-r\tau}$ (equation (13)) in which a_1 is a quadratic function of at-the-money-forward total volatility provides a good fit to the expansion coefficient a_1 . Table 3 shows the coefficient estimates and sum of squared errors (*SSE*) for a number of models for fitting a_1 . The total sum of squares *SST* of a_1 for all seven years is 0.4511. The first row of the table (model (a)) shows results for the quadratic model. The sum of squared errors for this model is 0.0562. This implies that this model explains 88 percent of the variation in a_1 . The second model (b) presents results for a version of the model that does not include the term $e^{-r\tau}$. The coefficient estimates and *SSE* are similar, but the first model that includes the term $e^{-r\tau}$ provides a slightly better fit. The next model (c) includes a constant term. Unsurprisingly, the model with a constant term provides a better fit, in the sense that the *SSE* is smaller. However, the estimated constant term is small, and not economically significant. The next set of models (d)–(h) experiment with different functional forms of $\sigma_F \sqrt{\tau}$ and $r\tau$, and find that none dominates the quadratic model (a). Finally, the last three models (i)–(k) specify the expansion coefficient a_1 as a function of τ alone rather than $\sigma_F \sqrt{\tau}$. Such models are problematic *a priori* as they are not dimensionless, and the fit of these models turns out to be poor.

We next estimate the quadratic model (13) (i.e., model (a)) separately for each year. Figure 5 presents a three-dimensional view of the fit for the calendar year 2001 when the model is estimated using the option prices from year 2001. The parameter estimates are $\alpha = 0.0437$ and $\beta = 0.0459$. In particular, Figure 5 plots the actual (dots) and fitted (stars) values of a_1 for each calendar date-expiration date pair during 2001 as a function of σ_F^2 and τ . For each actual value (each dot), the fitted value is represented by the star immediately above or below it. Figure 6 presents two-dimensional graphs obtained by projecting the data in Figure 5 onto the time to expiration, volatility, and total variance $\sigma_F^2 \tau$. Of each pair of panels, the left-hand panel shows the actual data and the right-hand panel shows the fitted values. This figure confirms that the quadratic model (13) provides an excellent fit, and also confirms that the total volatility $\sigma_F^2 \tau$ explains the bulk of the variation in the coefficient a_1 . It is clear that the quadratic model (13) explains most the variation in a_1 .

Figure 7 summarizes the fit of the quadratic model for the other years. In particular, it shows the actual and fitted values of the coefficients a_1 on the odd parabolic cylinder function \widetilde{D}_1 as a function of the total volatility $\sigma_F^2 \tau$ separately for each of the five years 1996–2000 and the period January 2001–September 2002. Each panel shows both the actual (gray dots) and fitted (black dots) values. Volatilities were relatively low in 1996, and there was no SPX option with time to expiration greater than one year in our sample. As a result, the data are concentrated toward the lower left-hand side of the 1996 panel. We see that the fit of the quadratic model is quite good for

all years except 1997 and 1998, and is reasonably good even for those years. In 1997 the quadratic model still seems like a good choice, but there is considerable dispersion around the fitted values. In 1998 there is a little less dispersion about the fitted values, and there seem to exist two patterns, with the fitted model roughly corresponding to the average of the two patterns. Also, volatilities were high in this time period, and the relation is more steeply sloped than in other periods. The fit is excellent for the remaining years, with similar shapes in 1999 and 2000 and a slightly less steeply sloped relation in 2001 and 2002.

3.3.2. The even component

Table 2 shows that the correlation between a_2 and the at-the-money-forward total variance $\sigma_F^2\tau$ is -0.6521 , which is of greater magnitude than the correlation between a_2 and any other variable. In unreported graphical analysis we plotted a_2 against the dimensionless quantities $\sigma_F\sqrt{\tau}$ and $r\tau$ as well as several other variables for the entire sample. Consistent with the argument that a_1 and a_2 should be functions of only dimensionless quantities, we found no clear relation with the forward price, volatility, or interest rate.

In the last section we preannounced that the model $a_2 = (\alpha_2\sigma_F^2\tau + \beta_2\sigma_F\sqrt{\tau})e^{-r\tau}$ (equation (14)) in which a_2 is a quadratic function of at-the-money-forward total volatility provides a good fit to the expansion coefficient a_2 . Table 4 shows the results of fitting this and alternative models for a_2 . The table displays the coefficient estimates and sum of squared errors (*SSE*) for the same set of models that are used to explain a_1 above. As expected, since a_2 is much less variable than a_1 and the time series of estimates have different properties, the coefficient estimate and *SSE*'s are much different from those in Table 3. Nonetheless, as with a_1 the quadratic model performs better than all other models, although the differences from some of the other models are small.

Figure 8 shows the actual and fitted values of the coefficient a_2 as a function of time to expiration τ , at-the-money-forward variance σ_F^2 , and at-the-money-forward total variance $\sigma_F^2\tau$ for calendar year 2001. Although both a_1 and a_2 are fit well by quadratic functions of the total volatility $\sigma_F\sqrt{\tau}$, their patterns are quite different. While a_1 seems to be monotone increasing in $\sigma_F\sqrt{\tau}$, a_2 first increases and then decreases against $\sigma_F\sqrt{\tau}$ and is nonnegative only for roughly the region $\sigma_F\sqrt{\tau} < 0.06$. As we will see later in Section 5, this has interesting implications for the second central moment of the implied risk-neutral density. Of each pair of panels in Figure 8, the left-hand panel shows the actual data and the right-hand panel shows the fitted values. This figure confirms that the quadratic model (14) provides an excellent fit, and also confirms that the total volatility $\sigma_F^2\tau$ explains the bulk of the variation in the coefficient a_2 . The fact that the fit for a_2 is not as good as that of a_1 is not surprising, as the even component explains a much smaller part of the price deviations, and any errors will be larger relative to the explained part.

Figure 9 summarizes the fit of the quadratic model for a_2 for the other years. In particular, it shows the actual and fitted values of the coefficients a_2 on the even parabolic cylinder function \widetilde{D}_2 as a function of the total volatility $\sigma_F^2\tau$ separately for each of the five years 1996-2000 and the period January 2001– September 2002. Each panel shows both the actual (gray dots) and fitted (black dots) values. As with Figure 7, for 1996 the data are all toward the left-hand side because there was no SPX option with time to expiration greater one year. We see that the fitted quadratic model is generally similar in different years, though there is a fair amount of dispersion about the fitted values in 1997 and 1998.

3.4. One-step estimation

The above exploratory data analysis used a two-step process. We first estimated the expansion coefficients a_1 and a_2 separately for each calendar date-expiration date pair, and then we searched for models that captured the changes in a_1 and a_2 . Having identified such a model, it is natural to do the analysis in a single step. The one-step model is

$$y = \alpha_1 X_1 + \beta_1 X_2 + \alpha_2 X_3 + \beta_2 X_4, \quad (22)$$

where

$$y = \frac{Y_A(d) - Y_{BS}(d, \sigma_F)}{F}, \quad (23)$$

and the right-hand side variables are

$$X_1 = \sqrt{2}\sigma_F^2\tau d \exp(-d^2/2 - r\tau), \quad (24)$$

$$X_2 = \sqrt{2}\sigma_F\sqrt{\tau}d \exp(-d^2/2 - r\tau), \quad (25)$$

$$X_3 = 2\sigma_F^2\tau d^2 \exp(-d^2/2 - r\tau), \quad (26)$$

$$X_4 = 2\sigma_F\sqrt{\tau}d^2 \exp(-d^2/2 - r\tau). \quad (27)$$

Notice that in equation (22) we are fitting all 456,280 call and put options simultaneously (228,140 (t, τ, X) triplets for both calls and puts). Estimating this model using the entire sample, we obtain coefficient estimates $\alpha_1 = 0.1003$, $\beta_1 = 0.0437$, $\alpha_2 = -0.0746$, and $\beta_2 = 0.0166$, and $R^2 = 0.9528$, as reported in Table 5. These estimates are similar to estimates for the quadratic models in Tables 3 and 4, where we have $\alpha_1 = 0.1037$, $\beta_1 = 0.0417$, $\alpha_2 = -0.0661$ and $\beta_2 = 0.0150$. These results imply that the simple model with four parameters that are constant over time explains more than 95 percent of the deviations from Black-Scholes prices.

The above fit is excellent, in light of the fact that the discreteness of quoted prices inherently limits the ability to obtain a perfect fit. Further, our model enforces put-call parity, which is not

exactly satisfied in the data. Notice that in our sample, we have both periods of rising index levels and falling index levels. However, the changes in index levels do not seem to affect the reported pattern although they affect the implied volatility levels. Also, we have options with time to expiration ranging from 10 days to 2 years. The fact that long-term and short-term options both follow the same pattern is somewhat striking. In particular, this implies that the relative expensiveness of out-of-the-money long-term options with respect to at-the-money long-term options is not very different from the relative expensiveness for short-term options. Our results also imply that SPX option traders have been pricing these options in a consistent manner for the period 1996-2002.

4. Implied volatilities

The deviations of actual option prices from Black-Scholes values computed with the same volatility for different strikes can be mapped into the implied volatility function through the inverse function of the Black-Scholes formula. This implies that the model (10)–(14) of deviations from Black-Scholes values has an equivalent representation in terms of the implied volatility function. In particular, the finding that the odd and even components of the price deviations are well fit by simple functions of the at-the-money forward total volatility suggests that implied volatilities of different strike prices can also be expressed in terms of the at-the-money forward total volatility. In this section we use a Taylor expansion to find a simple approximation of this relation, and thus connect our model of price deviations to a simple model of the implied volatility function. Our data set is the same one we used for the price deviation fitting. For each of the 12,381 calendar date-expiration date (t, τ) pairs there is an implied volatility function computed from the straddles for the different strike prices available. All together, there are 228,140 observations on the implied volatility σ . Subtracting each of these implied volatilities from their corresponding at-the-money-forward volatility for the same calendar and expiration dates, there are 228,140 observations of implied volatility deviations $\Delta\sigma$, defined as $\Delta\sigma = \sigma - \sigma_F$.

4.1. A simple model of the implied volatility function

The model of price deviations (10)–(14) or equation (22) can be written as

$$\frac{Y_A(d) - Y_{BS}(d, \sigma_F)}{F \exp(-r\tau) \exp(-d^2/2)} = \sqrt{2}(\alpha_1 \sigma_F^2 \tau + \beta_1 \sigma_F \sqrt{\tau})d + 2(\alpha_2 \sigma_F^2 \tau + \beta_2 \sigma_F \sqrt{\tau})d^2. \quad (28)$$

Letting $\sigma(t, \tau, X)$ denote the implied volatility at calendar date t for time to expiration τ and strike X , define the deviation from at-the-money-forward volatility to be $\Delta\sigma(t, \tau, X) \equiv \sigma(t, \tau, X) -$

$\sigma_F(t, \tau)$, where as above $\sigma_F(t, \tau)$ is the at-the-money-forward volatility. Recognizing that $Y_A(d) = Y_{BS}(\sigma) = Y_{BS}(\sigma_F + \Delta\sigma)$, we can use a first-order Taylor expansion to obtain the approximation

$$\frac{Y_A - Y_{BS}(d, \sigma_F)}{F \exp(-r\tau)} \approx \mathcal{V} \Delta\sigma \sqrt{\tau}, \quad (29)$$

where \mathcal{V} is the modified dimensionless vega defined by

$$\mathcal{V} \equiv \left. \frac{\partial(Y_{BS}(d, \sigma)/F \exp(-r\tau))}{\partial(\sigma \sqrt{\tau})} \right|_{\sigma=\sigma_F}. \quad (30)$$

Using the Black-Scholes formula, we see $\mathcal{V} = N'(d_1) \approx N'(d) \exp(-d\sigma_F \sqrt{\tau}/2)$. Combining (28) and (29) yields

$$\frac{\mathcal{V} \Delta\sigma \sqrt{\tau}}{\exp(-d^2/2)} = \sqrt{2}(\alpha_1 \sigma_F^2 \tau + \beta_1 \sigma_F \sqrt{\tau})d + 2(\alpha_2 \sigma_F^2 \tau + \beta_2 \sigma_F \sqrt{\tau})d^2. \quad (31)$$

In our data, the mean value of d is 0.24 and mean value of σ_F is 0.21. Thus we can approximate $\exp(d\sigma_F \sqrt{\tau}/2)$ by $1 + d\sigma_F \sqrt{\tau}/2$. Using this, we get the following regression model

$$\Delta\sigma \sqrt{\tau} = \alpha d \sigma_F \sqrt{\tau} + \beta d^2 \sigma_F \sqrt{\tau} + \gamma d \sigma_F^2 \tau + \delta d^2 \sigma_F^2 \tau. \quad (32)$$

Table 6 presents the results of estimating the model (32) and some other models over the entire sample period. The first set of results is for the model (32), which explains 95.56 percent of the variation in $\Delta\sigma$. The second set of results for the slightly simpler model

$$\Delta\sigma \sqrt{\tau} = \alpha d \sigma_F \sqrt{\tau} + \beta d^2 \sigma_F \sqrt{\tau} + \gamma d \sigma_F^2 \tau \quad (33)$$

that omits the term $d^2 \sigma_F^2 \tau$ explains 95.50 percent of the variation in $\Delta\sigma$. Due to the almost identical fits and the small magnitude of the term $d^2 \sigma_F^2 \tau$ in equation (32), we use the simpler model (33). The third and fourth sets of results are for two models obtained by restricting some parameters in the full model to be zero, while the fifth and sixth sets are for two other models that are not based on our expansion. The model $\alpha(X/1000) + \beta(X/1000)^2$ (the fifth set of results) is the quadratic-fit model of Shimko (1993). Although it is known that the Shimko model fits the volatility skew well when estimated separately on each day, when estimated with the same parameters across calendar days the Shimko model can only explain 22% of the volatility deviations. The last model $\alpha d + \beta d^2$ is a quadratic fit in the volatility-adjusted relative moneyness d . While it performs much better than the Shimko model its performance is still dominated by those of the first four models.

Figure 10 shows a three-dimensional plot of actual (dots) and fitted (stars) values from the

model in equation (33) for each calendar date-expiration date pair as a function of d and $d\sigma_F\sqrt{\tau}$ (only one out of 200 points are shown). The fit is perfect on the planes defined by $d = 0$ and $d\sigma_F\sqrt{\tau} = 0$ because on these planes the options are at-the-money forward, and hence their implied volatilities do not deviate from the implied volatility of the at-the-money forward option from which deviations are measured. Comparing the locations of the fitted values to the actual values, it is clear that the quadratic model (33) explains almost all of the variation in the values of $\Delta\sigma$.

Figure 10, while indicating a good fit, actually understates the quality of the fit because the data density is not homogeneous on the graph. For example, there are many more options with $\tau = 0.3$ than options with $\tau = 2.0$, and the bulk of the options with $\tau = 0.3$ have fitting errors that are very close to zero. However, these very small errors are hidden in the center of the figure where the data density is high. The following numbers give us some sense of the quality of fit. Out of all observations, 359 options (0.16%) have a fitting error larger than 0.05, 1,239 options (0.54%) have a fitting error larger than 0.04, 3,615 options (1.58%) larger than 0.03, 10,548 options (4.62%) larger than 0.02, and 39,161 options (17.17%) larger than 0.01. In the other tail of the distribution, 14,345 options (62.88%) have a fitting error smaller than 0.005, and 36.79% of the options have an error less than 0.002. Also, the large errors concentrate in the years 1997 and 1998 (especially in the months of August in both years), which is not surprising in light of Figure 7.

4.2. Implications of the implied volatility model

Our model captures the main features of the implied volatility surface and can be used to analyze the slope and curvature of the skew. Fix a (t, τ) pair and hold the at-the-money-forward implied volatility σ_F constant. Letting σ_X denote the implied volatility at strike price X for that (t, τ) pair, we have

$$\left. \frac{\partial \sigma_X}{\partial \ln X} \right|_{X=F} = -\frac{\alpha}{\sqrt{\tau}} - \gamma\sigma_F. \quad (34)$$

This equation says that there are two components to the slope of the skew. The first component, $-\alpha/\sqrt{\tau}$, is a time-varying component that goes to zero as time to expiration increases to infinity. The second component, $-\gamma\sigma_F$, is a volatility-level driven component. This component is present even for $\tau = \infty$ and is proportional to the stochastic at-the-money-forward volatility level. Taken together, the two components imply that the skew is negatively sloped, and becomes less so as the time to expiration τ increases. Also, the skew tends to be more negatively sloped when volatility is high. Note that α and γ come from the odd component of the price deviations only, implying that the odd component of the price deviation is directly responsible for the slope in the implied volatility skew.

It is also interesting to examine the curvature of the skew at $X = F$. From equation (33),

$$\left. \frac{\partial^2 \sigma_X}{\partial (\ln X)^2} \right|_{X=F} = \frac{2\beta}{\sigma_F \tau}. \quad (35)$$

Since β comes from the even component of the price deviations, this shows that the even component in the price deviations is directly responsible for the curvature of the skew. The slope and curvature of the skew can work together to produce a smile pattern. We see that the curvature is inversely proportional to at-the-money-forward volatility level and time to expiration. As time to expiration increases, the curvature decreases to zero.

The model (33) also provides an answer to the question of when we can expect to see a “smile” in the implied volatility function. The implied volatility curve reaches its minimum at

$$d = -\frac{\alpha + \gamma \sigma_F \sqrt{\tau}}{2\beta}, \quad (36)$$

which is about $-3.5 - 10\sigma_F \sqrt{\tau}$ for the parameter estimates we have. From the definition of d , we have $X = F \exp(-d\sigma_F \sqrt{\tau})$. If this turning point is low, then we are more likely to observe a smile pattern in the traded options, which usually are roughly centered around the current index level. Thus our model implies that we are likely to see a volatility smile more often when $\sigma_F \sqrt{\tau}$ is small (for example, in the relatively short-term options in year 1996) and when r is small or the dividend rate δ is large. We see that the volatility curve begin to increase again at $X = 1.34F$, when $\sigma_F \sqrt{\tau}$ is about 0.07. This agrees with most stochastic volatility option pricing models, in which it is easier to produce a smile when the instantaneous volatility and time to expiration are both small.

Some features of our results, however, are in contrast with one of the most common classes of stochastic volatility models. Models in this class have no jumps, and specify an auxiliary instantaneous volatility process that is not explicitly dependent on the stock price levels, though the Brownian motions driving the stock process and volatility process can be correlated. Specifically, a typical model is specified as follows:

$$dS = rS dt + \sqrt{V}S dW_S, \quad (37)$$

$$dV = b(V) dt + a(V) dW_V, \quad (38)$$

where the two Brownian motions are correlated with coefficient $\rho(V)$ and the instantaneous variance process V is assumed to be mean-reverting. This class of models includes the now famous Heston (1993a) model. Under some general assumptions, Lewis (2000) (see also Lee 2002) works out the

asymptotic behavior of the implied volatilities as $\tau \rightarrow 0$ (p. 143) and $\tau \rightarrow \infty$ (p. 187):

$$\lim_{\tau \rightarrow 0} \frac{\partial V_X(\tau)}{\partial \ln X} \Big|_{X=F} = \frac{\rho(V_F(0))a(V_F(0))}{2\sqrt{V_F(0)}}, \quad (39)$$

where $V_X = \sigma_X^2$ is the square of implied volatility at strike price X , and

$$V_X(\tau) \approx V_F^\infty + \frac{A \ln(F/X)}{\tau} + \frac{B \ln^2(F/X)}{\tau^2} + O(\tau^{-3}) \quad \text{as } \tau \rightarrow \infty, \quad (40)$$

where A , B and V_F^∞ are some constants depending on the drift and diffusion functions of the stochastic volatility process. That is to say, for very short-maturity options, the slope of the skew tends to approach a constant while for long-maturity options, the skew tends to flatten to a common asymptotic value regardless of the moneyness $\ln(F/X)$.

In our model, for short-maturity options,

$$\frac{\partial V_X}{\partial \ln X} \Big|_{X=F} = -\frac{2\alpha\sqrt{V_F}}{\sqrt{\tau}} - 2\gamma V_F. \quad (41)$$

In particular, the slope in our model at $X = F$ approaches infinity at the speed of $\sqrt{\tau}$. In comparing equation (41) with (39), one must be cautious because we have excluded options with maturities less than 10 days and used a certain parametric form. However, in light of the fact that our model explains over 95% of the variation in the implied volatilities it is hard to reconcile our empirical results with the generic behavior of stochastic volatility models for short-term options.

Our results for long-maturity options are also at odds with the generic behavior of stochastic volatility models. In our model,

$$V_X(\tau) = V_F + \frac{2\alpha\sqrt{V_F\tau} + 2\gamma V_F\tau}{\tau} \ln(F/X) + \frac{(\alpha + \gamma\sqrt{V_F\tau})^2 + 2\beta}{\tau} \ln^2(F/X) + O(\ln^3(F/X)). \quad (42)$$

At a first glance, equations (40) and (42) seem similar because both express V_X as a quadratic function of $\ln(F/X)$. However, in our model V_F does not approach a constant level but rather is stochastic. Thus, there can be considerable variation in at-the-money-forward implied volatilities even for long-maturity options. A second key difference is that in our model the slope of the skew does not approach zero as $\tau \rightarrow \infty$, but rather approaches a stochastic slope $-2\gamma V_F$. A third difference is the speed of convergence. In our model, if we hold V_F constant, the speed of convergence is $1/\sqrt{\tau}$. Interestingly, although it violates the long-maturity behavior of $O(1/\tau)$ predicted by stochastic volatility models, this convergence agrees with the rule of thumb widely used by practitioners that approximates the decay rate of the skew slope as $\tau^{-1/2}$ (see Lee 2002).¹⁰

¹⁰While our empirical results differ from predictions of stochastic volatility models with no jumps, it remains

5. Implied risk-neutral density

If the risk-free interest rate is constant, the value of a European call option can be written as

$$C = e^{-r\tau} \int_0^\infty (S_T - X)^+ f(S_T) dS_T, \quad (43)$$

where $f(S_T)$ is the marginal risk-neutral density of S_T , the stock price at the expiration date T . Prices of European put options can be characterized similarly. A standard result dating from Breeden and Litzenberger (1978) is that the risk-neutral density is given by

$$f(S_T) = e^{r\tau} \left. \frac{\partial^2 Y}{\partial X^2} \right|_{X=S_T}, \quad (44)$$

where Y is the value of either a call or a put option. Since volatility is itself stochastic, this implied density should be interpreted as the marginal density of some multi-dimensional risk-neutral distribution.

The price fitting model (10) can be written as

$$Y = Y_{BS} + \Delta Y \equiv Y_{BS} + F(\sqrt{2}a_1 d + 2a_2 d^2)e^{-d^2/2}, \quad (45)$$

where $a_1 = (\alpha_1 \sigma_F^2 \tau + \beta_1 \sigma_F \sqrt{\tau})e^{-r\tau}$, $a_2 = (\alpha_2 \sigma_F^2 \tau + \beta_2 \sigma_F \sqrt{\tau})e^{-r\tau}$, $F = S_t \exp((r - \delta)\tau)$, $d = \ln(F/X)/(\sigma_F \sqrt{\tau})$, and we have written r for $r(t, \tau)$, σ_F for $\sigma_F(t, \tau)$ etc. Differentiating, the implied risk-neutral density f^{imp} for S_T is given by

$$f^{\text{imp}} = f_{BS} + \Delta f, \quad (46)$$

where f_{BS} is the lognormal density function

$$f_{BS} = \frac{1}{\sqrt{2\pi\sigma_F^2\tau}S_T} \exp\left(-\frac{(\log S_T - \log S_t - (r - \delta - \sigma_F^2/2)\tau)^2}{2\sigma_F^2\tau}\right), \quad (47)$$

and

$$\Delta f = \frac{Fe^{-\bar{d}^2/2}e^{r\tau}}{S_T^2\sigma_F\sqrt{\tau}}(\lambda_0 + \lambda_1\bar{d} + \lambda_2\bar{d}^2 + \lambda_3\bar{d}^3 + \lambda_4\bar{d}^4), \quad (48)$$

unclear whether stochastic volatility models with jumps will better conform to out empirical findings. We leave it to future research.

where

$$\bar{d} = \frac{\ln F - \log S_T}{\sigma_F \sqrt{\tau}}, \quad \lambda_0 = \sqrt{2}a_1 + \frac{4a_2}{\sigma_F \sqrt{\tau}}, \quad \lambda_1 = 4a_2 - \frac{3\sqrt{2}a_1}{\sigma_F \sqrt{\tau}}, \quad (49)$$

$$\lambda_2 = -\sqrt{2}a_1 - \frac{10a_2}{\sigma_F \sqrt{\tau}}, \quad \lambda_3 = \frac{\sqrt{2}a_1}{\sigma_F \sqrt{\tau}} - 2a_2, \quad \lambda_4 = \frac{2a_2}{\sigma_F \sqrt{\tau}}. \quad (50)$$

Apart from the factor $e^{-r\tau}$, a_1 and a_2 are quadratic functions of $\sigma_F \sqrt{\tau}$, implying that the λ 's are as well. Figure 11 plots the implied risk-neutral density computed using (46)–(50) above and the lognormal density on the same graph. The parameters used are $\sigma = 0.15$, $\tau = 0.2$, $S = 1,000$, $r = 0.01$, $\alpha_1 = 0.1003$, $\beta_1 = 0.0437$, $\alpha_2 = -0.0746$, and $\beta_2 = 0.0166$. These α 's and β 's are the estimates from the one-step regression reported in Table 5.

A nice feature of this implied density given by (46)–(50) is that it is in closed form. Further, other than the at-the-money-forward volatilities, it requires only four parameters, the α 's and β 's, that describe the density for all strikes and times to expiration. Since one of the main results of this paper is that these four parameters are very stable over time, one can use the parameters estimated in this paper and treat them as constants. Thus, to use our parametric implied distribution, the only input is the term structure of observed implied volatilities $\sigma_F(t, \tau)$.

The implied risk-neutral density has a number of other desirable properties. The first nice property is that our implied density is self-consistent, that is, if we let the strike price equal the forward price, then the option price equals the Black-Scholes price evaluated at the at-the-money-forward implied volatility:

$$\int_0^\infty e^{-r\tau} (S_T - F(t, \tau))^+ f^{\text{imp}}(S_T) dS_T = C_{BS}(\sigma_F(t, \tau)). \quad (51)$$

Similar results hold for the put options. Second, Δf integrates to zero exactly, implying

$$\int_0^\infty f^{\text{imp}}(S_T) dS_T = 1, \quad (52)$$

a necessary condition for f^{imp} to be a sensible density. Third, the implied density gives the same expected future stock price as the lognormal density, i.e.,

$$\int_0^\infty S_T f^{\text{imp}}(S_T) dS_T = F. \quad (53)$$

The above equation says that our implied density satisfies the necessary condition to be a risk-neutral density. It also implies that our implied risk-neutral density preserves the forward price

and put-call parity since it implies

$$\int_0^\infty e^{-r\tau} (S_T - K)^+ f^{\text{imp}}(S_T) dS_T - \int_0^\infty e^{-r\tau} (K - S_T)^+ f^{\text{imp}}(S_T) dS_T = (F - K)e^{-r\tau}. \quad (54)$$

The above four equations can be easily derived, either by straightforward calculations or by repeated use of integration by parts. The implied density behaves so nicely because we have carefully chosen σ_F as the base volatility and our model is, by construction, consistent with the prices of options; and also because the parabolic cylinder functions used in our price model have good tail behavior, that is, they go to zero exponentially as the moneyness d goes to $\pm\infty$.

It is interesting to examine the higher moments of our implied distribution. Using integration by parts, it is easy to show that

$$\int_0^\infty (S_T - F)^2 f^{\text{imp}} dS_T = F^2(e^{\sigma_F^2\tau} - 1) + F^2\sqrt{8\pi\sigma_F^2\tau} \left(2a_2(1 + \sigma_F^2\tau) - \sqrt{2}a_1\sigma_F\sqrt{\tau}\right) e^{(r+\sigma_F^2/2)\tau}. \quad (55)$$

The first term on the right-hand side comes from the lognormal density, while the second term incorporates the skew. With a bit of calculation it can be shown that the second central moment from implied distribution will be equal to that from the lognormal distribution when either $\sigma_F\sqrt{\tau} = 0$ or

$$\sigma_F\sqrt{\tau} \doteq \frac{\sqrt{2}\alpha_2 - \beta_1 + \sqrt{(\sqrt{2}\alpha_2 - \beta_1)^2 + 4\sqrt{2}\beta_2\alpha_1}}{2\alpha_1}, \quad (56)$$

which with our fitted parameters equals 0.142. If $\sigma_F = 0.2$, then this occurs with $\tau = 0.5009$. For a wide range of values of τ , the second central moments of the two distributions are close. When τ is less than 0.5009 the second central moment from the implied distribution is larger than the second central moment of the lognormal distribution. When τ is greater than 0.5, the second central moment from the implied distribution is smaller than that from the lognormal distribution, and the relative difference of the two goes to zero exponentially fast when $\sigma_F\sqrt{\tau}$ goes to ∞ . This behavior comes from the fact that a_2 is positive for small $\sigma_F\sqrt{\tau}$ and becomes increasingly negative as $\sigma_F\sqrt{\tau}$ increases.

Our implied risk-neutral density can also be viewed as a correction to the normality assumption for $\ln S_T$ in the Black-Scholes formula. Our density approximation bears some similarity with the often-used Gram-Charlier series expansion (see, e.g. Corrado and Su 1996, Backus, Foresi, and Wu 2004) to correct for skewness and kurtosis. Both densities have a correction term involving a fourth-order polynomial. However, a key difference is that our approach specifies the correction terms as simple functions of at-the-money-forward total volatility. In the Gram-Charlier approxi-

mation, one has to estimate or otherwise specify the skewness and kurtosis parameters μ_3 and μ_4 separately for each (t, τ) pair.

Finally, with suitable parameters the implied distribution can easily produce a volatility “smile.” In particular, the even component is responsible for the fat right tail in our risk-neutral density. It is also responsible for the upward curvature that appears toward the right-hand side of the implied volatility function shown in the lower panel, i.e. the fat right tail produces the volatility smile. Although we do not pursue the issue here, the ability of the model to produce a volatility smile suggests that it may be able to work well with individual equity options for which the smile is more pronounced than it is in index options (although the American feature of the individual equity options complicates matters).

The most straightforward use of the implied density would be to value European-style options that are not actively traded, i.e. options with strike prices that are not actively traded, or digital options. A more interesting aspect of the implied density is that we observe a second mode or peak in the density. A small and relatively “flat” second mode can be seen toward the left-hand side of Figure 11. This second mode is produced by the even component, and its size depends both upon the parameters α_2, β_2 of the even component and the level of the total base volatility $\sigma_F \sqrt{\tau}$.

Figure 12 compares the lognormal and implied densities for a range of times to expiration, namely 7, 31, 61, 91, 122, 182, 365 days and 2 years. The graphs in the various panels are constructed using $\sigma = 0.25$, $S = 1,000$, and $r = 0.05$, with the α 's and β 's estimated from options in our sample with times to expiration within a small window of the times to expiration indicated on the panels. Except for the shortest maturity graph (the one labeled 7/365), the parameters are estimated from options with times to expiration falling within an eight day window centered on the time to expiration indicated on the panel. The shortest maturity graph is based on parameter estimates obtained from the options with times to expiration of between 10 and 15 days. We use an eight day window to ensure that we have enough observations to calculate the local estimates for α 's and β 's reliably. It is interesting to see that on all the panels the implied densities are bimodal, with the second mode gradually developing as time to expiration increases.

This finding of a bimodal risk-neutral density has implications for models with stochastic volatility and/or jump-diffusion processes, as it provides a stylized fact that proposed models must capture and thus rules out models that cannot produce this property. A possible concern is that the bimodality is due to some artifact of the methodology of expanding price deviations in terms of parabolic cylinder functions, and is not a robust feature of the data. One response is that the second peak is due to the even component, and Figure 2 illustrating how the odd and even components sum to the price deviations makes clear that the even component is important in capturing the price deviations. A second response is that other authors have also found some evidence of

bimodality in the risk-neutral density using other approaches. Third, mathematically it is perfectly possible for the risk-neutral distribution to be bimodal while the physical distribution is unimodal. Bakshi, Kapadia and Madan (2003) shows how risk-neutral density of the return can have nonzero skewness while physical distribution is symmetric through exponential tilting which comes from the Radon-Nikodym derivative. The same mechanism can also introduce bimodality in the risk-neutral distribution.¹¹

6. Conditional out-of-sample forecasting of the skew

The previous analyses show that the models of the skew provide better in-sample fits than a wide range of alternatives. In this section we conduct a conditional out-of-sample forecasting exercise to verify that the proposed models perform better in this sense as well. This exercise is conditional in that we do not predict the entire price or implied volatility function, but rather predict option prices (i.e., the skew) conditional on the index value and the term structure of at-the-money-forward implied volatilities. Specifically, for each calendar date t we start by estimating the parameters in the selected models from a rolling window of past data. We then step ahead to a forecast date $t + \Delta t$, and, given the index value and at-the-money-forward volatility on $t + \Delta t$, we use the estimated parameters and the model to predict the price (or equivalently, volatility) deviations, that is to predict the skew. Thus, this exercise only tests whether our models can correctly value the away-from-the-money options, given the valuation of the at-the-money-forward options.

Interestingly, this exercise is similar to the way the models might be used in risk measurement, for example value-at-risk calculations. In some methods of computing value-at-risk one simulates realizations of the underlying market factors that drive price changes, and then revalues the portfolio of financial instruments for each realization of the market factors. If one takes the market factors to be the underlying index value and the at-the-money-forward volatilities, then the exercise in this section corresponds to this potential use of the models.

We compare a total of seven models. Models I and II are the models of price and volatility deviations from Sections 3 and 4, respectively. Model III is a special case of Model II with one fewer parameter. Then Models IV through VII are various plausible alternative models. For each calendar date, we estimate the model parameters using data from the previous two months. We

¹¹Consider the following artificial example. Suppose the physical distribution of the future stock return R is

$$f(R) = \frac{A \exp(-\lambda R^2)}{1 + BR^2}, \quad \text{where } \lambda > 0, B > 0. \quad (57)$$

This physical distribution is unimodal. Now with exponential tilting factor $\exp(-kR)$, if the parameters λ , k and B are such that $k + 2(B + \lambda)x + Bkx^2 + 2B\lambda x^3 = 0$ has three distinct real roots, then the risk-neutral density can be bimodal.

then look some fixed period ahead to a “forecast” date, and use the parameter estimates and the at-the-money-forward volatility for the forecast date to compute the option prices. If that forecast date is not a business day, we use the next business day. We repeat this for every date.

The various models are as follows:

Model I: This is the price deviation model from Section 3,

$$\hat{Y} = Y_{BS} + F \left(a_1 \widetilde{D}_1(\sqrt{2}d) + a_2 \widetilde{D}_2(\sqrt{2}d) \right), \quad (58)$$

where for each date t the parameters a_1 and a_2 are estimated using data from the two months prior to t and \hat{Y} is the predicted price of a call or a put option on $t + \Delta t$. We implement the model in a one-step approach. That is, during the two-month estimation period we estimate the four parameters $\gamma_1, \dots, \gamma_4$ from the regression:

$$\frac{Y_A - Y_{BS}}{F e^{-r\tau}} = \gamma_1 \sigma_F^2 \tau d e^{-d^2/2} + \gamma_2 \sigma_F \sqrt{\tau} d e^{-e^2/2} + \gamma_3 \sigma_F^2 \tau d^2 e^{-d^2/2} + \gamma_4 \sigma_F \sqrt{\tau} d^2 e^{-d^2/2}. \quad (59)$$

Model II: This is model for volatility deviations from Section 4. In this model, $\hat{Y} = Y_{BS}(\hat{\sigma})$, where

$$\hat{\sigma} \sqrt{\tau} = \sigma_F \sqrt{\tau} + \gamma_1 d \sigma_F \sqrt{\tau} + \gamma_2 d^2 \sigma_F \sqrt{\tau} + \gamma_3 d \sigma_F^2 \tau, \quad (60)$$

where $\hat{\sigma}$ is the predicted value of σ . From the Taylor expansion in Section 4, this model is approximately equivalent to Model I, except that we dropped the term in $d^2 \sigma_F^2 \tau$. As a result this model has three rather than four parameters. However, there is no guarantee that a model with more parameters will perform better. On the contrary, Dumas, Fleming, and Whaley (1998) suggest that a more parsimonious model will perform better in out-of-sample tests.

Model III: We also estimate a simpler version of Model II. In this model, $\hat{Y} = Y_{BS}(\hat{\sigma})$, where

$$\hat{\sigma} \sqrt{\tau} = \sigma_F \sqrt{\tau} + \gamma_1 d \sigma_F \sqrt{\tau} + \gamma_2 d \sigma_F^2 \tau. \quad (61)$$

This model might perform better or worse than Model II. If Model II is the true model, then we would expect this model to perform worse than Model II; however, if Model II is overfit, then Model III, being more parsimonious, might outperform Model II.

Model IV: In this crudest benchmark model we make no adjustments for the skew, that is for each time to expiration we use the Black-Scholes values computed from the at-the-money-forward volatility. Thus,

$$\hat{Y} = Y_{BS}(\sigma_F). \quad (62)$$

However, if the parameters in the quadratic models of the price deviation components change over time, it is not clear that our model will beat this naive model. The sum of squared errors of this model can be viewed as the total sum of squares for all other models.

Model V: In this model the implied volatility function is approximated by a quadratic function of the moneyness, defined in terms of $X - F$. Thus, $\hat{Y} = Y_{BS}(\hat{\sigma})$, where

$$\hat{\sigma} = \sigma_F + \gamma_1(X - F) + \gamma_2(X - F)^2. \quad (63)$$

Again, the coefficients γ_1 and γ_2 are estimated using data from the two months prior to date t . This is a natural benchmark.

Model VI: Model VI is similar to V, except that moneyness is defined in terms of $\ln(X/F)$. Thus, $\hat{Y} = Y_{BS}(\hat{\sigma})$, where

$$\hat{\sigma} = \sigma_F + \gamma_1 \ln(X/F) + \gamma_2 \ln^2(X/F), \quad (64)$$

and the coefficients γ_1 and γ_2 are estimated using data from the two months prior to date t . This is another natural benchmark. One might expect it to perform better than Model V because it measures moneyness in relative terms.

Model VII: Model VII is similar to V and VI, except that it uses time-adjusted moneyness $\ln(X/F)/\sqrt{\tau}$:

$$\hat{Y} = Y_{BS}(\hat{\sigma}), \quad (65)$$

where

$$\hat{\sigma} = \sigma_F + \gamma_1 \frac{\ln(X/F)}{\sqrt{\tau}} + \gamma_2 \left(\frac{\ln(X/F)}{\sqrt{\tau}} \right)^2, \quad (66)$$

and again the coefficients γ_1 and γ_2 are estimated using data from the two months prior to date t . This benchmark tries to capture the dependence of the volatility skew on the time to expiration. This model is very similar to those used by Natenberg (1994) and Tompkin (1999), who argue that including the time to expiration helps in fitting the skew. We expect this model to perform well because we know that to first order the total volatility deviations are linear in $\ln(X/F)$. However, in this model the slope and curvature at $X = F$ do not depend on σ_F .

We compare the performances of Models I through III to the performances of the Black-Scholes benchmark Model IV and the alternative Models V, VI, and VII. Following Huang and Wu (2004), we compute the mean-squared errors (MSE) in each model for each forecasted day. Table 7 presents the means, medians and standard deviations of the daily forecast MSE's for various horizons Δt ranging from one to 150 calendar days, or approximately 5 months. Table 8 presents t -statistics for pairwise comparisons of the daily out-of-sample MSE's. In this table, Models V, VI and VII

are compared with the performance of Model III (the model we propose) and the Black-Scholes benchmark (Model IV). Models I, II and III are all variants of each other, and although Model II performs slightly better than Model III we choose Model III as the reference model because this model is the most parsimonious among all three and has the same number of parameters as Models V, VI and VII. We calculate the t -statistic comparing two models A and B as follows:

$$t\text{-statistic} = \frac{\text{mean}(\text{MSE}^A - \text{MSE}^B)}{\text{stderr}(\text{MSE}^A - \text{MSE}^B)}, \quad (67)$$

where the standard errors are auto-correlation adjusted based on Newey and West (1987), with the number of lags chosen based on Andrews (1991).

Table 7 reveals that for each model the mean MSE increases as the forecasting horizon increases, as do the medians and the standard deviations. A striking feature of the results is that Models I, II and III, in which the slope and curvature of the price deviation or implied volatility functions are determined by functions of at-the-money-forward total volatility, clearly dominate the other models. All measures of performance are markedly superior for this group of models, and the t -statistics in Table 8 reveal that the differences in the mean MSE's between Models V, VI, and VII and Model III are highly significant. At the one day horizon the mean forecast errors are between \$0.60 and \$0.68, increasing only to between \$0.95 and \$1.01 for the 31-day horizon and to between \$2.46 and \$2.58 for the 150 day horizon. The median forecast errors at the one day horizon are between \$0.34 and \$0.46, increasing only to between \$0.94 and \$1.14 at the 150-day horizon. Within the group of three models Model II seems to perform slightly better than Models I and III, but the differences are small. The results for these three models are consistent with the idea that Model II reflects the best tradeoff between flexibility and parsimony, but this may well be over-interpreting the small differences in performance.

Of the other models, as one would expect the crude benchmark Black-Scholes Model IV that does not attempt to capture the skew has the largest mean MSE's for all horizons. In Model V, introducing the skew with moneyness measured in terms of $X - F$ decreases the mean MSE's significantly—the mean MSE decreases from 48.59 to 19.74 for the one-day horizon, from 49.23 to 22.73 for the 31-day horizon, and from 51.59 to 35.50 for the 150-day horizon. The t -statistics in Table 8 reveal that the differences in the mean MSE's are highly significant. Model VI, in which moneyness is defined as $\ln(X/F)$, offers additional improvement over Model V. Finally, Model VII that measures the skew in terms of time-adjusted moneyness $\ln(X/F)/\sqrt{\tau}$ provides a very large improvement over Models V and VI. Still, Model VII is clearly inferior to Models I–III.

7. Conclusion

This paper presents a new approach for describing the differences between the actual prices of S&P 500 index options and their Black-Scholes values (computed with the same volatility for different strikes) by expanding the deviations of actual prices from Black-Scholes values using parabolic cylinder functions. It suffices to use only two functions, corresponding to the odd and even components of the price deviations, respectively. While the price deviations from the Black-Scholes vary randomly over time, we find that the price deviations are systematic and can be fit quite well as quadratic functions of the total volatility. This relation is very stable.

Price deviations from the Black-Scholes formula have an equivalent representation in terms of the implied volatility surface and also in terms of the (marginal) risk-neutral distribution of the index level. We develop the relation between the price deviations and the implied volatilities, and find that the differences between the total volatility at the various strike prices and the at-the-money-forward total volatility satisfy a simple function of total-volatility-adjusted moneyness and at-the-money-forward total volatility. Loosely, the slope and curvature of the implied volatility function are determined by the at-the-money-forward total volatility. These findings of stability in the price deviation and implied volatility functions stand in marked contrast to the view of a large fraction of the existing literature. Turning to the risk-neutral distribution, we find that the implied (marginal) risk-neutral density of index returns is bimodal.

Since the first appearance of Jackwerth and Rubinstein (2001), a challenge for option pricing models has been to improve on the predictive performance of the simple trader rules known as the absolute and relative smile models. Li and Pearson (2006) update the Jackwerth-Rubinstein (2001) horse race among competing models using more recent data set and more recent models. They find that the simple trader rules still dominate all other models and their performance can be further improved by adjusting the predicted skew shapes using the volatility skew model proposed in this paper. Perhaps surprisingly, after shape adjustment even the Black-Scholes model generally performs better than mathematically more sophisticated models, even if the latter models are also shape adjusted. With conditioning on actual at-the-money implied volatilities, the performance of every model improves dramatically after shape adjustment.

Our findings of stability in the price deviations and the implied volatility skew, and our finding of bimodality in the implied risk-neutral density have implications for models with stochastic volatility and/or jump-diffusion processes, as they provide stylized facts that proposed models must capture and thus rule out models that cannot produce these properties. More generally, the description of the deviations of option prices from Black-Scholes values in terms of parabolic cylinder functions is likely to be useful for empirical analysis because it provides a low-dimensional (i.e., small number of parameters) description of the price deviation and implied volatility functions.

References

- Aït-Sahalia, Yacine, and Jefferson Duarte, 2003, Nonparametric option pricing under shape restrictions, *Journal of Econometrics* 116, 9–47.
- Aït-Sahalia, Yacine, and Andrew W. Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance* LIII, No. 2, 499–547.
- Andrews, Donald, 1991, heteroskedasticity and autocorrelation consistent covariance matrix estimation, *Econometrica* 59, 817–858.
- Backus D., S. Foresi, L. Wu, Accounting for biases in Black-Scholes, (2004), working paper.
- Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical Performance of Alternative Option Pricing Models, *Journal of Finance* 52 (December), 2003–2049.
- Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan, 2001, Stock return characteristics, skewness laws, and the differential pricing of individual equity options, *Review of Financial Studies* 16, No. 1, 101–143.
- Bates, David S., 1991, The Crash of 1987: Was it expected? The evidence from options markets, *Journal of Finance* XLVI, No. 3 (July), 1009–1044.
- Bates, David S., 1996, Jumps and stochastic volatility: exchange rate processes implicit in Deutsche Mark options, *Review of Financial Studies* 9, No. 1 (Winter) 69–107.
- Bates, David S., 2000, Post-'87 crash fears in the S&P 500 futures option market, *Journal of Econometrics* 94, 181–238.
- Black, Fischer, 1975, Fact and fantasy in the use of options, *Financial Analysts Journal* 31 (July/August), 36–41 and 61–72.
- Bollen, Nicholas P.B. and Robert E. Whaley, 2004, Does net buying pressure affect the shape of implied volatility functions? *Journal of Finance* LIX, No. 2 (April), 711–753.
- Bondarenko, Oleg, 2001, On market efficiency and joint hypothesis, Working paper, University of Illinois at Chicago.
- Bondarenko, Oleg, 2003, Estimation of risk-neutral densities using positive convolution approximation, *Journal of Econometrics* 116, 85–112.
- Brandt, Michael W., and Tao Wu, 2002, Cross-sectional tests of deterministic volatility functions, *Journal of Empirical Finance* 9, 525–550.
- Breeden, Douglas T. and Robert H. Litzenberger, 1978, Prices of State Contingent Claims Implicit in Option Prices, *Journal of Business* 51, 621–652.
- Buraschi, Andrea and Jens Jackwerth, 2001, The price of a smile: Hedging and spanning in option markets, *Review of Financial Studies* 14, No. 2 (Summer), 495–527.
- Campa, J., K. Chang, and R. Reider, 1998, Implied exchange rate distributions: Evidence from the OTC options market, *Journal of International Money and Finance* 17, 117–160.

- Carr, Peter and Dilip Madan, 1999, Option pricing and the fast fourier transform, *Journal of Computational Finance* 2, No. 4 (Summer), 61–73.
- Carr, Peter, Hlyette Geman, Dilip Madan, and Marc Yor, 2002, The fine structure of asset returns: An empirical investigation, *Journal of Business* 75, No. 2 (April), 305–332.
- Carr, Peter and Liuren Wu, 2003a, The finite moment log stable process and option pricing, *Journal of Finance* 58, No. 2, 753–777.
- Carr, Peter, and Liuren Wu, 2003b, What type of process underlies options? A Simple Robust Test, *Journal of Finance* 58, No. 6 (December), 2581–2610.
- Carr, Peter, and Liuren Wu, 2004, Time-changed Levy processes and option pricing, *Journal of Financial Economics* 71, No. 1 (January), 113–141.
- Corrado, Charles J. and Tie Su, 1996, Skewness and kurtosis in S&P 500 index returns implied by option prices, *Journal of Financial Research* 19 (Summer), 175–192.
- Cox, John C. and Stephen A. Ross, 1976, The valuation of options for alternative stochastic processes, *Journal of Financial Economics* 3, 145–166.
- Dennis, Patrick, and Stewart Mayhew, 2002, Risk-neutral skewness: Evidence from stock options, *Journal of Quantitative Analysis* 37, No. 3 (September), 471–493.
- Derman, Emanuel, 1999a, Regimes of volatility: Some observations on the variation of S&P 500 implied volatilities. Goldman Sachs: Quantitative Strategies Research Notes (January).
- Derman, Emanuel, 1999b, Regimes of volatility, *Risk* 12, 55–59.
- Derman, Emanuel and Iraj Kani, 1994, Riding on the smile, *Risk* 7, 32–39.
- Derman, Emanuel, Iraj Kani, and J. Zou, 1996, The local volatility surface: Unlocking the information in index option prices, *Financial Analysts Journal* 52, 25–36.
- Duffie, Darrel, Jun Pan, and Kenneth Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343–1376.
- Dumas, Bernard, Jeff Fleming, and Robert E. Whaley, 1998, Implied volatility functions: empirical tests, *Journal of Finance* LIII, No. 6 (December), 2059–2106.
- Dupire, Bruno, 1994, Pricing with a smile, *Risk* 7, 18–20.
- Eberlein, Ernst, Ulrich Keller, and Karsten Prause, 1998, New insights into smile, mispricing and value at risk: the hyperbolic model, *Journal of Business* 71, No. 3 (July), 371–405.
- Gârleneau, Nicolae, Lasse Heje Pedersen, and Allen M. Poteshman, 2006, Demand-based option pricing, working paper, University of Illinois at Urbana-Champaign.
- Gatheral Jim, 1999. The volatility skew: Arbitrage constraints and asymptotic behavior. Merrill Lynch.
- Geman, Hélyette, Dilip B. Madan and Marc Yor, 2001, Time changes for Lévy processes, *Mathematical Finance* 11, 79–96.

- Gemmill, Gordon and Naoki Kamiyama, 2000, International transmission of option volatility and skewness: When you're smiling, does the whole world smile? Working paper, City University Business School.
- Hentschel, Ludger, 2003, Errors in implied volatility estimation, *Journal of Financial and Quantitative Analysis*, Vol. 38, Issue 4, 779–810.
- Heston, Stephen L., 1993a, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* 6, 327–343.
- Heston, Stephen L., 1993b, Invisible parameters in option prices, *Journal of Finance* 48, 933–947.
- Heynen, R., 1993, An empirical investigation of observed smile patterns, *Review of Futures Markets* 13, 317–353.
- Hodges, Hardy M., 1996. Arbitrage bounds on the implied volatility strike and term structure of European-style options. *Journal of Derivatives*, 23-35.
- Huang, Jing-zhi and Liuren Wu, 2004, Specification analysis of option pricing models based on time-changed Lvy processes, *Journal of Finance* 59, No. 3 (June), 1405–1439.
- Hull, John and Alan White, 1987, The pricing of options on assets with stochastic volatilities, *Journal of Finance* 42, 281–300.
- Hurst, S. R., E. Platen, and Svetlozar T. Rachev, 1999, Option pricing for a logstable asset price model, *Mathematical and Computer Modeling* 29, No. 10, 105–119.
- Jackwerth, Jens Carsten, 2000, Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* 13, No. 2 (Summer), 433–451.
- Jackwerth, Jens Carsten and Mark Rubinstein, 1996, Recovering Probability Distributions from Option Prices, *Journal of Finance* LI, No. 5 (December), 1611–1631.
- Jackwerth, Jens Carsten and Mark Rubinstein, 2001, Recovering stochastic process from option prices, working paper, University of Wisconsin at Madison and Haas School of Business, University of California at Berkeley.
- Janicki, A. W., Ivilina Popova, Peter H. Ritchken, and Wojbor A. Woyczynski, 1997, Option pricing bounds in an α -stable security market. *Communications in Statistics: Stochastic Models* 13, No. 4, 817–839.
- Lee, Roger W., 2002. Implied volatility: statics, dynamics and probabilistic interpretation, in *Recent Advances in Applied Probability*, 2005, Springer.
- Lewis, Alan L., 2000. *Option Valuation under Stochastic Volatility*. Finance Press, Newport Beach, California.
- Li, Minqiang, and Neil D. Pearson, 2006. A “Horse race” among competing option pricing models using S&P 500 index options,” working paper, Georgia Institute of Technology and University of Illinois at Urbana-Champaign.
- Longstaff, Francis A., 1995, Option pricing and the martingale restriction, *Review of Financial Studies* 8 (Winter), 1091–1124.

- McCulloch, J. Huston, 1987, Foreign exchange option pricing with log-stable uncertainty, in Sarkis J. Khoury and Alo Ghosh, eds.: *Recent Developments in International Banking and Finance* (Lexington Books, Lexington, MA).
- McCulloch, J. Huston, 1996, Financial applications of stable distributions, in G. S. Maddala and Calyampudi R. Rao, eds.: *Statistical Methods in Finance* (North Holland, Amsterdam).
- Madan, Dilip B., Peter P. Carr, and Eric C. Chang, 1998, The variance gamma process and option pricing, *European Finance Review* 2, No. 1, 79–105.
- Madan, Dilip B., and Franke Milne, 1991, Option pricing with VG martingale components, *Mathematical Finance* 1, No. 4, 39–56.
- Madan, Dilip B., and E. Seneta, 1990, The VG model for Share Market Returns, *Journal of Business* 63, No. 4, 511–524.
- Melick, William R. and Charles P. Thomas, 1997, Recovering an asset's implied probability density function from option prices: An application to oil prices during the Gulf crisis, *Journal of Financial and Quantitative Analysis* 32 (March), 91–115.
- Melino, Angelo and Stuart M. Turnbull, 1990, Pricing foreign currency options with stochastic volatility, *Journal of Econometrics* 45, 239–265.
- Melino, Angelo and Stuart M. Turnbull, 1995, Misspecification and the pricing and hedging of long-term foreign currency options, *Journal of International Money and Finance* 14, 373–393.
- Merton, Robert C., 1976, Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 125–143.
- Natenberg, S., 1994, *Option Volatility & Pricing: Advance Trading Strategies and Techniques*, McGraw-Hill Publishing.
- Newey, Whitney K. and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, No. 3, 703–708.
- Popova, Ivilina and Peter H. Ritchken, 1998, On bounding option prices in Paretian stable markets, *Journal of Derivatives* 5 No. 4, 32–43.
- Rosenberg, J.V., 2000, Implied volatility functions: A reprise, *Journal of Derivatives* 7, 51–64.
- Rubinstein, Mark, 1994, Implied binomial trees, *Journal of Finance* 49, 771–818.
- Scott, Louis O., 1987, Option pricing when the variance changes randomly: Theory, estimation, and an application, *Journal of Financial and Quantitative Analysis* 22, 419–438.
- Sheikh, Amir M., 1991, Transaction data tests of S&P 100 call option pricing, *Journal of Financial and Quantitative Analysis* 26 (December), 459–475.
- Shimko, David, 1993, Bounds of probability, *Risk* 6, No. 4, 33–37.
- Stein, Elias M. and Jeremy C. Stein, 1991, Stock price distributions with the stochastic volatility: An analytic approach, *Review of Financial Studies* 4, 727–752.

Tompkins, R.G., 2001, Implied volatility surfaces: Uncovering regularities for options on financial futures, *European Journal of Finance*, Volume 7, No. 3 (September), 198 – 230

Wang and Guo, *Special Functions*, World Scientific, 1989 (first published by Science Press in Chinese in 1978).

Whaley, Robert E., 1986, Valuation of American futures options: Theory and empirical tests, *Journal of Finance* 41 (March), 127–150.

Whaley, Robert E., 2002, Derivatives, *Handbook of the Economics of Finance*, ed. by George Constantinides, Milton Harris, and Rene Stulz.

Wiggins, James, 1987, Option values under stochastic volatility: Theory and empirical estimates, *Journal of Financial Economics* 19, 351–372.

Table 1
Sums of squared errors SSE and R^2 's for different choices of N

Using the April 26, 2002 prices of the SPX options expiring in June 2002, we estimate the model

$$y_k \equiv \frac{Y_A(z_k) - Y_{BS}(z_k, \sigma_F)}{F} \approx \sum_{n=0}^{N-1} a_n D_n(z_k),$$

where k indexes the different options, $d_k = \ln(F/X_k)/(\sigma_F\sqrt{\tau})$, D_n is the n -th parabolic cylinder function given by

$$D_n(z) = (-1)^n e^{z^2/4} \frac{d^n}{dz^n} \left(e^{-z^2/2} \right), \quad n = 0, 1, 2, \dots$$

and

$$z_k = \sqrt{2}d_k = \frac{\sqrt{2} \ln(F/X_k)}{\sigma_F\sqrt{\tau}}.$$

The table reports the sum of squared errors SSE and R^2 's for different choices of N . The column $2'$ is for $N = 3$ but with the constraint $a_0 = a_2$. The total sum of squares is 241.7974.

N	1	2	$2'$	3	4	5	6	7
SSE	230.08	22.197	9.9890	6.0089	0.2543	0.2540	0.0840	0.0839
R^2	0.0491	0.9082	0.9587	0.9751	0.9989	0.9989	0.9997	0.9997

Table 2
Correlations between the expansion coefficients and other variables

This table reports the correlations between the expansion coefficients a_1 and a_2 for the odd and even components of the price deviations, the dimensionless total volatility $\sigma_F^2\tau$, and other variables. The expansion coefficients are estimated separately for each calendar day-expiration date pair by estimating the model of price deviations (18) using ordinary least squares.

	a_1	a_2	τ	σ_F	$\sigma_F^2\tau$	F	r	$S \exp(-\delta\tau)$
a_1	1.0000							
a_2	-0.6245	1.0000						
τ	0.7364	-0.5164	1.0000					
σ_F	0.5248	-0.2439	0.0297	1.0000				
$\sigma_F^2\tau$	0.9277	-0.6521	0.8643	0.4033	1.0000			
F	0.3333	-0.1763	0.0846	0.4754	0.2826	1.0000		
r	0.0165	-0.0491	0.0218	-0.1013	-0.0024	0.0423	1.0000	
$S \exp(-\delta\tau)$	0.2340	-0.1046	-0.0444	0.4756	0.1679	0.9897	0.0053	1.0000

Table 3
Models to explain the coefficient a_1 on the odd expansion function \widetilde{D}_1

The table shows the parameter estimates and sums of squared errors (*SSE*) for a variety of models to explain the expansion coefficient a_1 for the odd component. The number of observations is 12,381. The total sum of squares is 0.451127. The standard errors for the parameter estimates are given in the parentheses.

Model	f_1	γ	α	β	<i>SSE</i>
(a)	$(\alpha(\sigma_F^2\tau) + \beta(\sigma_F\sqrt{\tau})) \exp(-r\tau)$		0.1037 (0.0015)	0.0417 (0.00036)	0.0562
(b)	$\alpha(\sigma_F^2\tau) + \beta(\sigma_F\sqrt{\tau})$		0.0795 (0.0015)	0.0432 (0.00036)	0.0597
(c)	$\gamma + \alpha(\sigma_F^2\tau) + \beta(\sigma_F\sqrt{\tau})$	0.0011 (0.000081)	0.1120 (0.0027)	0.0300 (0.0010)	0.0587
(d)	$\gamma + \alpha(\sigma_F^2\tau) + \beta(\sigma_F^2\tau)^2$	0.0030 (0.000037)	0.2170 (0.0018)	-0.2496 (0.015)	0.0616
(e)	$\gamma + \alpha(\sigma_F^2\tau)^2 + \beta(\sigma_F\sqrt{\tau})$	-0.0001 (0.000055)	0.4484 (0.0099)	0.0547 (0.00041)	0.0572
(f)	$\alpha(\sigma_F^2\tau)^\beta$		0.0903 (0.00017)	0.6270 (0.0017)	0.0619
(g)	$\gamma + \alpha(\sigma_F^2\tau)^\beta$	0.0019 (0.000033)	0.1213 (0.00042)	0.7857 (0.0017)	0.0594
(h)	$\alpha(\sigma_F^2\tau)^\beta \exp(-r\tau)$		0.1062 (0.00017)	0.6598 (0.0018)	0.0588
(i)	$\alpha\tau + \beta\sqrt{\tau}$		0.0007 (0.00015)	0.0119 (0.00016)	0.1951
(j)	$\alpha\tau^2 + \beta\sqrt{\tau}$		0.0001 (0.000052)	0.0124 (0.000081)	0.1954
(k)	$\alpha\tau^2 + \beta\tau$		-0.0061 (0.000094)	0.0199 (0.00014)	0.2095

Table 4
Models to explain the coefficient a_2 on the even expansion function \widetilde{D}_2

The table shows the parameter estimates and SSE 's for a variety of models to explain the expansion coefficient a_2 for the even component. The number of observations is 12,381. The total sum of squares for a_2 is 0.0162. The standard errors for the parameter estimates are given in the parentheses.

Model	f_2	γ	α	β	SSE
(a)	$(\alpha\sigma_F^2\tau + \beta\sigma_F\sqrt{\tau}) \exp(-r\tau)$		-0.0661 (0.00060)	0.0150 (0.00014)	0.0088
(b)	$\alpha(\sigma_F^2\tau) + \beta(\sigma_F\sqrt{\tau})$		-0.0622 (0.00057)	0.0143 (0.00014)	0.0089
(c)	$\gamma + \alpha(\sigma_F^2\tau) + \beta(\sigma_F\sqrt{\tau})$	0.00037 (0.000032)	-0.0516 (0.0011)	0.0099 (0.00039)	0.0088
(d)	$\gamma + \alpha(\sigma_F^2\tau) + \beta(\sigma_F^2\tau)^2$	0.00093 (0.000014)	-0.0111 (0.00067)	-0.1367 (0.0058)	0.0089
(e)	$\gamma + \alpha(\sigma_F^2\tau)^2 + \beta(\sigma_F\sqrt{\tau})$	0.00103 (0.000022)	-0.1820 (0.0039)	-0.00233 (0.00016)	0.0090
(f)	$\alpha(\sigma_F^2\tau)^\beta$		-0.8265 (0.0145)	2.8295 (0.0076)	0.0138
(g)	$\gamma + \alpha(\sigma_F^2\tau)^\beta$	0.00088 (0.000009)	-0.0863 (0.00086)	1.5434 (0.0048)	0.0089
(h)	$\alpha(\sigma_F^2\tau)^\beta \exp(-r\tau)$		-0.9232 (0.0161)	2.8403 (0.0078)	0.0138
(i)	$\alpha\tau + \beta\sqrt{\tau}$		-0.0028 (0.000038)	0.0030 (0.000040)	0.0119
(j)	$\alpha\tau^2 + \beta\sqrt{\tau}$		0.0009 (0.000014)	0.0012 (0.000021)	0.0131
(k)	$\alpha\tau^2 + \beta\tau$		-0.0012 (0.000025)	0.0016 (0.000036)	0.0143

Table 5

One-step analysis for the deviations of actual option prices from Black-Scholes values

The one-step model is

$$y = \alpha_1 X_1 + \beta_1 X_2 + \alpha_2 X_3 + \beta_2 X_4,$$

where y is the normalized price deviation given by

$$y = \frac{Y_A(d) - Y_{BS}(d, \sigma_F)}{F},$$

and

$$\begin{aligned} X_1 &= \sqrt{2}\sigma_F^2\tau d \exp(-d^2/2 - r\tau), \\ X_2 &= \sqrt{2}\sigma_F\sqrt{\tau}d \exp(-d^2/2 - r\tau), \\ X_3 &= 2\sigma_F^2\tau d^2 \exp(-d^2/2 - r\tau), \\ X_4 &= 2\sigma_F\sqrt{\tau}d^2 \exp(-d^2/2 - r\tau). \end{aligned}$$

The parameters α_1 and β_1 are related to the odd component while α_2 and β_2 are related to the even component. Total number of observations is 456,280. Estimated standard errors are in parentheses below the coefficient estimates.

α_1	β_1	α_2	β_2	SST	SSE	R^2
0.1003	0.0437	-0.0746	0.0166	14.0628	0.6635	0.9528
(0.00029)	(0.000061)	(0.00030)	(0.000052)			

Table 6
Models to explain the volatility deviations $\Delta\sigma\sqrt{\tau}$

The table shows the parameter estimates and R^2 's for several of models to explain the total volatility deviations $\Delta\sigma\sqrt{\tau}$. The number of observations is 228,140. The total sum of squares for $\Delta\sigma\sqrt{\tau}$ is 107.28. Estimated standard errors are in parentheses below the coefficient estimates.

Model	α	β	γ	δ	R^2
$\alpha d\sigma_F\sqrt{\tau} + \beta d^2\sigma_F\sqrt{\tau} + \gamma d\sigma_F^2\tau + \delta d^2\sigma_F^2\tau$	0.1352 (0.00027)	0.0308 (0.00019)	0.4366 (0.0015)	-0.0899 (0.046)	0.9556
$\alpha d\sigma_F\sqrt{\tau} + \beta d^2\sigma_F\sqrt{\tau} + \gamma d\sigma_F^2\tau$	0.1410 (0.00025)	0.0207 (0.000087)	0.3995 (0.0013)		0.9550
$\alpha d\sigma_F\sqrt{\tau} + \beta d^2\sigma_F\sqrt{\tau}$	0.2098 (0.00013)	0.0124 (0.00010)			0.9372
$\alpha d\sigma_F\sqrt{\tau} + \gamma d\sigma_F^2\tau$	0.1695 (0.00025)		0.3020 (0.0014)		0.9442
$\alpha(X/1000) + \beta(X/1000)^2$	0.0500 (0.00017)	-0.0370 (0.00013)			0.2163
$\alpha d + \beta d^2$	0.0212 (0.00004)	-0.0011 (0.00002)			0.6414

Table 7
Out-of-sample forecasting performance of Models I–VII in Section 6

For each calendar date t , we estimate parameters in these models using 2 months of past data. We then use these parameters to predict the option prices for date $t + \Delta t$. N is the number of forecasted trading days. We report the mean, median and standard deviations of mean-squared errors on each day for each of the models.

Horizon Δt (days)		N	Model (Number of parameters)						
			I (4)	II (3)	III (2)	IV (BS) (0)	V (2)	VI (2)	VII (2)
1	1630	mean	0.65	0.60	0.68	48.59	19.47	16.24	1.96
		median	0.41	0.34	0.46	36.84	14.71	13.33	1.08
		std. dev.	0.95	0.95	0.90	37.93	16.46	11.87	2.10
3	1629	mean	0.66	0.62	0.69	48.68	19.65	16.30	1.98
		median	0.41	0.36	0.47	36.75	14.55	13.27	1.11
		std. dev.	0.91	0.92	0.86	37.96	16.90	11.98	2.10
7	1626	mean	0.72	0.68	0.75	48.99	20.30	16.61	2.01
		median	0.43	0.37	0.49	36.97	14.65	13.35	1.15
		std. dev.	1.09	1.10	1.03	38.14	18.03	12.40	2.20
31	1609	mean	1.01	0.95	1.01	49.23	22.73	17.48	2.17
		median	0.54	0.47	0.58	37.25	14.63	13.52	1.26
		std. dev.	1.54	1.55	1.46	37.89	22.57	13.95	2.53
61	1588	mean	1.47	1.38	1.44	49.96	26.47	18.93	2.45
		median	0.73	0.65	0.75	37.93	14.95	13.63	1.32
		std. dev.	2.23	2.23	2.14	38.03	30.00	16.39	3.07
91	1567	mean	1.92	1.82	1.87	50.68	30.07	20.23	2.75
		median	0.85	0.77	0.91	38.29	15.69	13.56	1.44
		std. dev.	2.75	2.76	2.63	37.95	38.18	19.09	3.53
150	1527	mean	2.58	2.46	2.51	51.59	35.50	21.74	3.29
		median	1.07	0.94	1.14	39.05	16.90	14.25	1.85
		std. dev.	3.18	3.18	3.09	37.50	51.59	22.10	4.12

Table 8**Pairwise t -statistics for Model Comparisons in conditional out-of-sample forecasting**

Entries report the t -statistics defined in equation (67). Every one of the models V, VI and VII are compared with the performance of model III (this paper) and model IV (Black-Scholes benchmark). We compare the differences in performance for seven different forecasting horizons. N is the number of forecasted trading days. All corresponding p -values are smaller than 10^{-5} .

Horizon Δt	N	V-III	VI-III	VII-III	V-IV	VI-IV	VII-IV
1	1630	24.06	27.80	16.68	-20.49	-21.58	-25.99
3	1629	23.60	27.55	16.53	-20.32	-21.53	-25.99
7	1626	22.72	26.97	16.27	-20.13	-21.54	-26.01
31	1609	19.87	24.53	14.06	-17.79	-20.81	-26.17
61	1588	17.01	21.88	11.44	-14.01	-19.72	-26.29
91	1567	14.95	19.58	9.01	-10.34	-18.49	-26.56
150	1527	12.78	17.49	7.11	-6.24	-17.25	-26.90




Figure 1. Differences between the actual prices and Black-Scholes values for both SPX call and put options expiring in June 2002 observed on April 26, 2002. The actual prices are the bid-ask midpoints from the close of trading, and the Black-Scholes values are computed using the at-the-money-forward implied volatility for the straddles.




Figure 2. Odd and even components of the price deviation function. The horizontal axis is the total-volatility-adjusted moneyness d . The solid line is the price deviation function, the line with circles is the odd expansion function \widetilde{D}_1 , and the line with crosses (“+”) is the even expansion function \widetilde{D}_2 . The three lines illustrate the approximate decomposition of the price deviation function into the odd and even components.




Figure 3. Expansion coefficients a_0 , a_1 , and a_2 from equation (9) for the five months prior to the expiration of the June 2002 option.




Figure 4. Expansion coefficients a_1 and a_2 from equation (18) for the five months prior to the expiration of the June 2002 option using two expansion functions.




Figure 5. Actual and fitted values of the coefficients a_1 on the odd parabolic cylinder function \widetilde{D}_1 as a function of the implied variance σ_F^2 and the remaining time to expiration τ for each calendar date-expiration date pair during 2001. The actual values are shown by dots \cdot , and the fitted values by stars $*$. The coefficients a_1 are estimated separately for each calendar date-expiration date pair by estimating equation (13) using ordinary least squares.




Figure 6. Actual and fitted values of the coefficients a_1 on the odd parabolic cylinder function \widetilde{D}_1 as a function of the time to expiration τ , the implied volatility σ_F , and the total volatility $\sigma_F^2 \tau$ for each calendar date-expiration date pair during 2001. The actual values are shown in the left-hand panels, and the fitted values in the right-hand panels. The coefficients a_1 are estimated separately for each calendar date-expiration date pair by estimating the quadratic model (13) using ordinary least squares.




Figure 7. Actual and fitted values of the coefficients a_1 on the odd parabolic cylinder function \widetilde{D}_1 as functions of the total volatility $\sigma_F^2 \tau$ separately for each of the five years 1996-2000 and the period January 2001 - September 2002. Each panel shows both the actual (gray dots) and fitted (black dots) values. The coefficients a_1 are estimated separately for each calendar date-expiration date pair by estimating the quadratic model (13) using ordinary least squares.




Figure 8. Actual and fitted values of the coefficients a_2 on the even parabolic cylinder function \widetilde{D}_2 as a function of the time to expiration τ , the implied volatility σ_F , and the total volatility $\sigma_F^2 \tau$ for each calendar date-expiration date pair during 2001. The actual values are shown in the left-hand panels, and the fitted values in the right-hand panels. The coefficients a_2 are estimated separately for each calendar date-expiration date pair by fitting (14) using ordinary least squares.




Figure 9. Actual and fitted values of the coefficients a_2 on the odd parabolic cylinder function \widetilde{D}_2 as a function of the total volatility $\sigma_F^2 \tau$ separately for each of the five years 1996–2000 and the period January 2001 - September 2002. Each panel shows both the actual (gray dots) and fitted (black dots) values. The coefficients a_1 are estimated separately for each calendar date-expiration date pair by fitting (14) using ordinary least squares.




Figure 10. Actual and fitted values of the implied volatility deviations $\Delta\sigma$ as a function of d and $d\sigma_F\sqrt{\tau}$ for one out of every 200 calendar date-expiration date pairs during the entire sample period. The actual values are shown by dots \cdot , and the fitted values by stars $*$.




Figure 11. Implied risk-neutral density computed by differentiating the price deviation model (10) compared to the lognormal density. The parameters used are: $\sigma = 0.15$, $\tau = 0.2$, $S = 1,000$, $r = 0.01$. The parameters α_i and β_i used are those from the one-step regression, that is, $\alpha_1 = 0.1003$, $\beta_1 = 0.0437$ and $\alpha_2 = -0.0661$, $\beta_2 = 0.0150$.




Figure 12. Implied and lognormal risk-neutral densities for various horizons, computed using parameters estimated from the entire sample. The graphs in the various panels are constructed using $\sigma = 0.25$, $S = 1,000$, and $r = 0.05$, with the α 's and β 's estimated from all options over the entire sample period with times to expiration within a small window of the times to expiration indicated on the panels. Except for the shortest maturity graph (the one labeled 7/365), the parameters are estimated from options with times to expiration falling within an eight day window centered on the time to expiration indicated on the panel. The shortest maturity graph is based on parameter estimates estimated from the options with times to expiration of between 10 and 15 days.