The Finite Moment Logstable Process And Option Pricing*

Peter Carr†

Banc of America Securities

Liuren Wu‡

Graduate School of Business, Fordham University

October 5, 2000; first draft: February 21, 2000

*We welcome comments, including references to related papers we have inadvertently overlooked. The latest version of the paper can be downloaded from `http://www.bnet.fordham.edu/public/finance/lwu/index.html` or `http://www.math.nyu.edu/research/carrp/papers`.

†**Correspondence Information:** 9 West 57th Street, 40th floor, New York, NY 10019; tel: (212) 583-8529; fax: (212) 583-8569; `pcarr@bofasecurities.com`

‡**Correspondence Information:** 113 West 60th Street, New York, NY 10023; tel: (212) 636-6117; fax: (212)765-5573; `wu@fordham.edu`

The Finite Moment Logstable Process And Option Pricing

ABSTRACT

We document for the first time a striking regularity in the U.S. equity index options market. We empirically observe that when implied volatilities are graphed against a standard measure of moneyness, the implied volatility smirk does not flatten out as maturity increases. This behavior contrasts sharply with that observed in currency options and also contradicts a prediction of virtually all existing stationary option pricing models. These models all imply that the volatility smirk should flatten out as maturity increases, due to the onset of the central limit theorem. This implication holds for all models with or without jumps and stochastic volatility, so long as the model has finite skewness and kurtosis of returns and a stationary volatility process. To capture the actual behavior of the index volatility smirk across maturities, we develop a parsimonious model which has infinite skewness and kurtosis of returns, but has finite moments of all orders for stock prices themselves. We calibrate our model and demonstrate its superior empirical performance against several alternatives.

JEL Classification Codes: G12, G13, F31, C14.

Key Words: Volatility smirk; maturity effect; Lévy α-stable motion; self-similarity; option pricing.

Ever since the stock market crash of 1987, the U.S. stock index options market has been exhibiting a striking pattern documented by academics and practitioners alike. At a given maturity level, the Black and Scholes (1973) implied volatilities for out-of-the-money puts (or in-the-money calls) are much higher than those which are near or in-the-money.[1] This phenomenon is commonly referred to as the "volatility smirk."

It is less well known that implied volatilities also have a strong empirical regularity in the maturity direction. When implied volatilities are graphed against a standard measure of "moneyness", we document that the resulting implied volatility smirk does not flatten out as maturity increases. The measure of moneyness for which this observation holds is the log of the strike over the forward, normalized by the square root of maturity. In the Black Scholes model, this moneyness measure is roughly the number of standard deviations that the log of the strike is away from the log of the forward.

Despite many empirical studies of the US equity index options market, to the best of our knowledge, this invariance of the smirk to maturity has not been documented previously. There are several potential reasons why this regularity has eluded researchers' attention. First, in other option markets such as currencies, the implied volatility smirk or smile does flatten out as maturity increases.[2] Second, when the moneyness measure is not normalized by the square-root of maturity, the smirk does flatten out as maturity increases.[3] Finally, all standard option pricing models imply that the volatility smirk ultimately flattens with increasing maturity, whether or not the model allows jumps and/or stochastic volatility. This flattening phenomenon is due to the central limit theorem and arises whenever the model has finite conditional moments for stock returns and the volatility process is stationary.

It is well known that the implied volatility smirk is a direct result of conditional non-normality in stock returns. In particular, the downward slope reflects asymmetry in the risk-neutral distribution of the underlying stock return, as measured by skewness, and the

[1] See, for example, Ait-Sahalia and Lo (1998), Jackwerth and Rubinstein (1996), and Rubinstein (1994) for empirical documentation of this phenomenon in S&P 500 Index options.

[2] Backus, Foresi, and Wu (1997) document this pattern for major currency options, although the flattening occurs at a much slower rate than implied by i.i.d innovations.

[3] For example, Ait-Sahalia and Lo (1998) use the ratio of strike over futures on the horizontal axis of their plots of the implied volatility smirk.

curvature of the smirk reflects the fat-tails of the risk-neutral return distribution, as measured by kurtosis. In continuous-time finance, a standard way to generate skewness and kurtosis in the distribution of stock returns is to use a Lévy process with jumps. However, when the log price is modeled as a Lévy process with jumps and finite moments, the skewness, if finite, decreases like the reciprocal of the square-root of maturity, and the kurtosis, if finite, decreases with the reciprocal of maturity (see Konikov and Madan (2000) for a proof). As a result, the conditional distribution converges to normality and the implied volatility smirk flattens as maturity increases. Incorporating stochastic volatility slows down this effect, but does not stop the convergence to normality, so long as the volatility process is stationary.[4]

To prevent the onset of the central limit theorem and the accompanying flattening of the smirk, we develop a new option pricing model for which the return distribution of the underlying stock has infinite moments for any order of two or greater. Although higher order *return* moments are infinite for our process, all moments of the stock *price* itself are finite. This delicate balancing act is accomplished by modeling returns as driven by a Lévy α-stable process with maximum negative skewness. The resulting return distribution has a very fat left tail but a thin right tail. The fat left tail engenders infinite moments for returns of orders greater than α, obviating the convergence to normality and the flattening of the volatility smirk. The thin right tail is thickened when we use the exponential function to map returns to prices, but remains thin enough so that the resulting price distribution has finite conditional moments of all orders. The finiteness of these price moments guarantees the existence of an equivalent martingale measure and, therefore, the finiteness of option prices at all maturities. Our specification also allows a simple analytical form for the characteristic function of the return. Many standard contingent claims can then be readily priced by the fast Fourier transform (FFT) method of Carr and Madan (1999).

In the Black Scholes model, the log price is modeled as an arithmetic Brownian motion, which is a standard example of a Lévy process, i.e. a process with stationary independent increments. It is well known that the Black Scholes model has a free scale parameter σ which controls the width of the risk-neutral distribution. Our Lévy α-stable process is also a Lévy

[4]See, for example, Backus, Foresi, and Wu (1997) and Das and Sundaram (1999) on the convergence to normality under stochastic volatilities.

process, but in addition to a free scale parameter σ, the process also has an additional free parameter $\alpha \in [0, 2]$ which controls the tails of the distribution. In particular, setting $\alpha = 2$ causes our log price process to degenerate to a Brownian motion, so that the Black Scholes model arises as a special case of our model. Setting $\alpha < 2$ generates a fat left tail consistent with the implied volatility smirk observed in the equity index options market.

Under our Lévy α-stable process, the slope of the implied volatility smirk is directly controlled by the tail index parameter α. The fact that the implied volatility smirk does not flatten out with increasing maturity is, in fact, a direct consequence of the *self-similarity* property of the Lévy α-stable process. Under such a self-similar process, the tail behavior of the return distribution is invariant under time aggregation.

The stochastic component of our log price process is a pure jump process with no continuous component. This implies that in our model, the risk in an option position cannot be delta-hedged away, although a riskless portfolio can be formed by dynamically trading in a continuum of options. Since this form of dynamic trading must be regarded as a physical impossibility at present, our model must in turn be interpreted as treating options as primary assets, which serve a role in helping to complete the market. Given the increasingly popular view that markets are fundamentally incomplete, there are many possible risk-neutral processes for stock and option prices which are consistent with an assumed statistical process for the stock and the absence of arbitrage. However, the selection of a particular risk-neutral process for the stock price can be interpreted as a way to indirectly select a particular risk-neutral process and a corresponding initial value for each member of the family of options written on the stock. The validity of the choice can be measured by capturing the extent to which the model option prices match the market option prices across all strikes, maturities, and dates. Thus, we regard the Lévy α-stable process which we use to specify the risk-neutral process for stock returns as a way to indirectly but exogenously impose a unique arbitrage-free stochastic process and initial value for each member of the option family.

To show that our choice is validated by the time series and cross sectional behavior of observed option prices, we calibrate our model and test it against a number of alternatives using about one year's worth of daily S&P 500 index option data. The results indicate that

(1) despite the parsimony of our model, it has superior explanatory power over all alternatives tested, and (2) the key weakness in these alternative pure jump and jump-diffusion models is their inability to simultaneously fit volatility smirks at different maturities.

The paper is organized as follows. The next section reviews the literature on existing option pricing models and on the Lévy α-stable process. Section II documents in detail the maturity pattern of the volatility smirk for S&P 500 index options. Section III introduces our specification of the risk-neutral stock price dynamics and illustrates the procedure used to determine option prices. Section IV deals with the design and results of our empirical tests on S&P 500 index options. Section V illustrates how the model can be readily extended to (1) include a diffusion component, (2) to refine the structure of the jump, and (3) to incorporate stochastic volatility. Foe each case, we derive the closed form representation of the characteristic function for the log return. We also calibrate these extended models and discuss the merits of each extension. Section VI briefly summarizes the paper and suggests some future research.

I. Literature Review

It is well known that at a given maturity, implied volatilities of index options slope down in a slightly convex manner when graphed against various measures of moneyness. It is widely agreed that the downward slope is a direct result of asymmetry in the conditional risk-neutral distribution of returns from the underlying asset, while the curvature arises from fatter tails in both ends of this distribution. In short, the moneyness bias is a direct result of conditional non-normality in the underlying return's risk-neutral density. Many processes have been proposed to generate non-normality in asset returns. In discrete time, conditional non-normality can be directly incorporated into the innovations via (i) expansions, e.g. Abadir and Rockinger (1997), Backus, Foresi, and Wu (1997), Brenner and Eom (1997), Corrado and Su (1996), Corrado and Su (1997), Jarrow and Rudd (1982), Jondeau and Rockinger (1998), Longstaff (1995), Potters, Cont, and Bouchaud (1998), Abken, Madan, and Ramamurite (1996a), Abken, Madan, and Ramamurite (1996b), and Rubinstein (1998); (ii) direct specification of more generalized

distributions, e.g. Aparicio and Hodges (1998), Bookstaber and McDonald (1987), Posner and Milevsky (1998), Sherrick, Garcia, and Tirupattur (1995), Sherrick, Irwin, and Forster (1992), and Sherrick, Irwin, and Forster (1996); and (iii) mixtures of distributions, e.g. Melick and Thomas (1997) and Ritchey (1990). In continuous time finance, a popular way to generate conditional non-normality at moderate to long maturities is through the use of stochastic volatility models. These models also capture the well-documented volatility clustering feature of the data. Prominent examples include the stochastic volatility models by Heston (1993), Hull and White (1987) and Romano and Touzi (1997) and the GARCH option pricing models by Duan (1995), Duan, Gauthier, and Simonato (1999), Duan and Simonato (2000), and Duan and Wei (1999). A second popular way to generate non-normal innovations in continuous time is to incorporate jumps into the asset price process. Examples include processes with finite jump arrival rates such as in Merton (1976),Bates (1991) and Kou (1999), processes with infinite jump arrival rates such as the variance-gamma model of Madan and Milne (1991) and Madan, Carr, and Chang (1998), and processes that can have finite or infinite jump arrival rates such as the CGMY model by Carr, Geman, Madan, and Yor (2000). Incorporating a stochastic or GARCH-type volatility process not only

While many researchers have focussed on the volatility smirk across moneyness at one fixed maturity, we argue that the maturity pattern provides us with an additional dimension which can be used to design models and test their implications. Failing to take account of the maturity dimension can lead to unreasonable and unstable parameter estimates in model calibration and to erroneous and unwarranted conclusions in model comparisons. For example, when calibrating a pure jump or jump-diffusion model with constant volatility, one obtains rather different parameter estimates depending on the maturity or range of maturities for the option data used in the calibration. The reason is simple: given the appropriate choice of parameter estimates, a constant volatility jump-diffusion model is capable of fitting the volatility smirk at any given maturity; but such a model cannot simultaneously fit smirks at different maturities. Similarly, when comparing the performance of different option pricing models, the particular maturities of the options used in the data set tends to determine the conclusions reached. In particular, when using moderate to long maturity options, one often finds that the contribution of a jump process is minimal, while the role of stochastic volatility

process is vitally important. However, the role of a jump process is much more important when calibrating models to short maturity options, since all stochastic volatility models imply that the conditional distribution of the stock return at the initial instant is normal, and are thus incapable of capturing the deep smirk observed in overnight options.

Backus, Foresi, and Wu (1997) and Das and Sundaram (1999) illustrate how stochastic volatility can interact with jumps to slow down the flattening of the volatility smile or smirk with increasing maturity. These authors argue that both jumps and stochastic volatility components are necessary to capture the variation in currency option prices across the whole spectrum of maturities and strikes. A key parameter that controls the decay of the volatility smile or smirk over maturity is the persistence of the stochastic volatility. The more persistent the volatility process is, the more slowly the smile flattens out. Backus, Foresi, and Wu (1997) give examples on how to calibrate the mean-reversion parameter for the volatility process to the cross-sectional maturity pattern of currency implied volatilities.

While a jump-diffusion model with stochastic volatility can capture the moneyness and maturity patterns of currency option implied volatilities, it cannot capture the corresponding behavior of U.S. equity index volatilities. In particular, the maturity pattern of the volatility smirk requires that the volatility process be non-stationary, while the time series data suggest the opposite. The model which we propose defuses this inherent tension successfully.

Financial asset returns are often regarded as the cumulative outcome of a large number of information arrivals occurring more or less continuously over time. According to a generalization of the standard central limit theorem, if the sum of a large number of i.i.d. random variables has a limiting distribution after appropriate shifting and scaling, then the limiting distribution must be a member of the α-*stable* class (Lévy (1925) and Feller (1971)). It is therefore natural to assume that asset returns are at least approximately governed by a stable distribution if the accumulation is additive, or by a log-stable distribution if the accumulation is multiplicative. [5]

[5] Mittnik, Rachev, and Paolella (1998), for example, illustrate how different compounding schemes lead to different stable processes in the limit.

The normal distribution is the most familiar and computationally tractable stable distribution with a tail index $\alpha = 2$. It has routinely been postulated to govern asset returns. However, since the pioneering work of Fama (1965), it has been known that returns are often much more leptokurtotic than implied by normality. This naturally leads to the consideration of non-Gaussian stable distributions with $\alpha < 2$, as proposed by Mandelbrot (1963b), Mandelbrot (1963a), and Fama (1965).

Because stable distributions are infinitely divisible, they are particularly attractive for continuous time modeling, as emphasized by Samuelson (1965) and McCulloch (1978). The stable generalization of the familiar standard Brownian motion is often called the *Lévy α-stable motion* and is the subject of two recent monographs by Samorodnitsky and Taqqu (1994) and Janicki and Weron (1994). The relevance of stable motions for option pricing has been recognized previously. For example, Janicki, Popova, Ritchken, and Woyczynski (1997) and Popova and Ritchken (1998) derive option pricing bounds in an α-stable security market.

However, in the nondegenerate case of $\alpha < 2$, any finite interval almost surely contains an infinite number of discontinuities. This feature implies that the risk in an option position cannot be delta-hedged away and that the variance and higher moments of asset returns are not defined. In particular, the infinite variance may imply that the risk premium is not finite. McCulloch (1987) and McCulloch (1996) circumvent this problem in pricing options by an innovative specification of the marginal utility of the asset.

Our approach to the modeling of stock index returns is unique. In light of the extreme negative skewness evidenced by equity index option prices, we incorporate maximum negative skewness into the Lévy α-stable motion. This incorporation allows us to not only capture the asymmetry of the risk-neutral distribution exhibited in stock option prices, but also to obtain finite moments of all orders for the asset price. We are thereby assured of the finiteness of the risk premium, and the corresponding finiteness of option prices.

II. Maturity Variation in the Volatility Smirk

A. Data and estimation issues

To document the maturity variation in the volatility smirk, we have obtained daily closing bid and ask implied volatility quotes on out-of-the-money S&P 500 index options across all strikes and maturities from April 6th, 1999 to May 31st, 2000 (290 business days). We also have daily closing futures prices corresponding to each option maturity. Our exclusive use of out-of-the-money options is an industry convention, arising from their greater liquidity and model sensitivity than their in-the-money counterparts.

We apply the following filters to the data: (1) the time to maturity is greater than five business days; (2) the bid implied volatility quote is positive; (3) the ask is no less than the bid. After applying these filters, we also plot the mid implied volatility for each day and maturity against strike prices to visually check for obvious outliers. After removing these outliers, we have 62,950 option quotes left over a period of 290 business days.

B. The smoothed implied volatility surface

In Figure 1, we plot the average implied volatility surface across maturity and moneyness, where the latter is defined as:

$$d = \frac{\ln(K/F)}{\sigma\sqrt{\tau}}, \tag{1}$$

for σ as some measure of the average volatility of the index. K and F are respectively the strike price and the futures price corresponding to the implied volatility quote. The use of the constant σ in the denominator of (1) is an industry convention designed to allow comparison across stocks, and to allow a simple interpretation of this moneyness measure as roughly the number of standard deviations that the log strike is away from the log forward price in the Black Scholes model[6] . In Figure 1, we use $\sigma = 27.4\%$, which is the average of all the implied volatility quotes.

[6]Note that d is also the average of $-d_1$ and $-d_2$ appearing in the standard Black Scholes put formula.

We use the mid point of the bid and ask for the implied volatility quotes. The implied volatility (IV) curve as a function of moneyness and maturity is obtained via nonparametric kernel regression. This smoothing method is especially suitable to cases where the functional structure is unknown, but extensive data are available. The kernel regression estimator can be expressed as:

$$\widehat{IV}(\mathbf{Z}) = \frac{\sum_{i=1}^{N} K\left(\frac{|\mathbf{Z}-\mathbf{Z_i}|}{H}\right) IV_i}{\sum_{i=1}^{N} K\left(\frac{|\mathbf{Z}-\mathbf{Z_i}|}{H}\right)}, \tag{2}$$

where $K(\cdot)$ is the kernel function, $\mathbf{Z} = (d, \tau)$ is the information set, N is the number of observations, and H is called the *bandwidth* matrix. Equation (2) is essentially a weighted average of the implied volatility quotes (IV_i). The kernel $K(\cdot)$ assigns a weight to each implied volatility quote IV_i based on the distance between the corresponding information $\mathbf{Z}_i = (d_i, \tau_i)$ and the conditional information $\mathbf{Z} = (d, \tau)$: $|\mathbf{Z} - \mathbf{Z_i}|$. The kernel is generally decreasing with the distance. The bandwidth matrix H adjusts the sensitivity of the kernel (weight) to the distance.

The choice of the kernel function typically has little influence on the end result. We use the most popular one, the independent Gaussian kernel:

$$K(\cdot) \quad = \quad \Pi_{j=1}^{2} k(x_j)\,, \quad k(x_j) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x_j^2}{2}\right),$$

with $x_j = (z_j - z_{ij})/h_j$ where z_j is an element of $\mathbf{Z}$ and z_{ij} is an element of $\mathbf{Z_i}$. With the independent kernel, we choose a diagonal bandwidth matrix with the j-th diagonal element given by h_j.

The degree of smoothing depends to a large extent on the choice of the bandwidth matrix H. The larger is the bandwidth, the greater is the degree of smoothing. The optimal choice is a tradeoff between bias and variance, which is best captured by minimizing the mean integrated squared error(MISE):

$$MISE = \mathbb{E}\left[\int_{\mathbf{Z}} \left(\widehat{IV}(\mathbf{Z}) - IV(\mathbf{Z})\right)^2 d\mathbf{Z}\right].$$

In practice, MISE is expanded and attention is restricted to the contribution from only the leading terms of the expansion, which is called the asymptotic MISE (AMISE). Assuming that the explanatory variables ($\mathbf{Z_i}$) are multivariate normally distributed and applying a Gaussian kernel, minimizing the AMISE gives the optimal bandwidth matrix,

$$H = \left(\frac{4}{\delta + 2}\right)^{1/(\delta+4)} \Sigma^{1/2} N^{-1/(\delta+4)},$$

where Σ is the variance-covariance matrix of the the information set $\mathbf{Z_i}$ and δ is the dimension of $\mathbf{Z}$. In the case of independent Gaussian kernels with $\delta = 2$, we have:

$$h_j = \sigma_j N^{-1/6}, \tag{3}$$

where σ_j is the standard deviation of z_{ij}. When the information set is not multivariate normal, or when the time series are autocorrelated, the bandwidth choice changes accordingly. Sometimes, cross-validation is necessary to empirically determine an optimal bandwidth.[7] For our purpose, (3) suffices as a benchmark choice for the bandwidth. We have experimented with different bandwidths by applying a different multiplier to the bandwidth in (3). We have also performed nonparametric smoothing with a local linear Gaussian kernel. In every case, the basic features of the implied volatility surface remain the same, regardless of changes in the bandwidth or the kernel.

The nonparametrically smoothed implied volatility surface is plotted in Figure 1. As expected, at each maturity, implied volatilities decline as moneyness increases: in fact, this decrease is almost linear[8]. A striking feature of Figure 1 is the lack of variation in the slope of this line as the maturity increases. This lack of variation stands in sharp contrast to the flattening out effect observed in the currency options market. Figure 2 slices Figure 1 into several two-dimensional graphs and overlays the volatility smirks at different maturities. As maturity increases, the overall level of the volatility smirk rises slightly, but does not flatten out at all. If anything, the slope for long maturity smirks is slightly steeper.

[7] See Simonoff (1996) for a detailed discussion on the choice of kernels and bandwidth.
[8] We recognize that this linear relation cannot be maintained asymptotically.

Ait-Sahalia and Lo (1998) report a slightly flattening in their nonparametrically estimated implied volatility smirk (see their Figure 4). However, the horizontal axis of their plot uses moneyness defined as K/F. If the horizontal axis of our Figure 1 were changed to this moneyness measure, then our volatility skew would flatten out due to the stretching of our original horizontal axis as maturity increases. Thus, it is possible that their results would be consistent with ours if they were plotted on the same axes. Consistent with this conjecture, their nonparametrically estimated skewness becomes more negative as maturity increases, while their kurtosis estimates increase with maturity (see their Figure 7).

C. A formal test of the maturity pattern

While Figure 1 is illustrative, it lacks statistically rigor. In this section, we perform a statistical test to determine whether the slopes of the implied volatility smirk is flattening or steepening with increasing maturity. Suppose that for each day, options have m different maturities indexed by $j = [1, 2, \cdots, m]$. At each maturity, we regress implied volatility on the moneyness, d,

$$IV_j = a + b_j d_j + e_j,$$

where b_j captures the slope of the smirk. As shown in Figure 1, the volatility smirk is approximately linear within the observed moneyness range. However, as shown in Figure 2, the slope in short term implied volatility does appear to change as the moneyness goes from negative to positive. Including the implied volatilities generated from out-of-the-money calls would bias the estimate of the slope, whose variation with maturity is our prime concern. Thus, for the regression, we restrict the moneyness to be in the range $d \in [-2, 0]$. The lower bound on moneyness is imposed because puts struck further out-of-the-money are not as liquid as at short maturities, and quotes at these strikes are not available at long maturities. For the regression to be robust, we further require at least five data points at each maturity.

For each day in our sample period, we obtained m regression slopes, one for each maturity. The top panel of Figure 3 depicts the term structure of the smirk slope for each of the 290 business days tested. Visual inspection indicates that the term structure is in most cases (1)

not linear in maturity, and (2) downward sloping. Bearing in mind that smirk slopes are negative, the downward sloping term structure implies that smirk slopes are becoming more negative as maturity increases. To determine the statistical significance of this effect, on each day we regress the smirk slope on the time to maturity to capture the overall slope of the term structure:

$$b_j = \alpha + \beta\tau_j + \epsilon.$$

The slope estimate β is negative for most days, i.e. the smirk usually steepens with increasing maturity, as suggested by the top panel in Figure 3. For 195 out of the 290 days tested, the t-statistic for β is less than -1.96, while for no days is the t-statistic greater than 1.96. Figure 4 depicts the nonparametrically smoothed distribution of the t-statistics on the estimates of β over the 290 days tested. The t-statistic is overwhelmingly negative and skewed to the left. In particular, the probability that the t-statistic is greater than 1.96 is less than 1%.

As a further test, we repeat the regression using the pooled term structure data for all days. The slope estimate from the pooled data (β) is -0.6083, with a t-statistic of -19.32, which is significantly negative for any reasonable confidence interval. The regression fit is depicted in the bottom panel of Figure 3.

Both the smoothed implied volatility surface and the formal regression tests point to a robust feature of the S&P 500 index options market. When graphed against the appropriate measure of moneyness, the implied volatility smirk does not flatten out as maturity increases. As a result, the risk-neutral density of the index return is *NOT* converging to normality as maturity increases. The central limit theorem, which is used by Backus, Foresi, and Wu (1997) to explain the long term behavior of almost all option pricing models with jumps and stationary stochastic volatility, does not apply to the market for S&P 500 index options. Clearly, a new class of models is required in order to account for the behavior of US equity index option prices.

III. The FMLS Process and Option Pricing

The key implication of the empirical results in the previous section is that stock returns exhibit self-similar properties under the risk-neutral measure. The tail behavior of the risk-neutral distribution is approximately invariant under time aggregation. Modeling the stock return by a Brownian motion, as in the Black Scholes model, captures such a self-similar property: stock returns remain normal under time aggregation. However, it is well-known that the implied density of stock returns is highly non-normal, and in particular, is heavily skewed to the left. The only process that possesses both non-normality and self-similarity is the Lévy α-stable process with the tail index parameter $\alpha < 2$. When $\alpha = 2$, the Lévy α-stable process degenerates to a Brownian motion.

Unfortunately, a *symmetric* Lévy α-stable motion for log prices with $\alpha < 2$ has infinite expected arithmetic return. This led Merton (1976) to conjecture that such a specification would make a 5-minute call option worth 100 percent of its underlying stock. Merton (1976) further conjectures that this specification for the market portfolio might also require the equilibrium interest rate to be infinite.

To circumvent these difficulties, we model the stock return by a Lévy α-stable process with *maximum negative skewness*. Our use of a Lévy α-stable process retains the key advantages of a non-normal return distribution and self-similarity. Our imposition of maximum negative skewness is required to deliver *finite* conditional moments of all orders for the stock price, relieving Merton's concerns. As a result, we christen this process the **Finite Moment Log Stable** (FMLS) process. Although our intent in maximizing negative skewness is to produce finite option prices, our specification has the bonus of capturing the highly skewed feature of the implied density for stock returns, a feature that cannot be captured by either a Brownian motion or a symmetric Lévy α-stable process.

Formally, let S_t denote the stock price at time t. We assume that the risk-neutral measure Q is such that under it, the stock price follows the log-stable process:

$$dS_t/S_t = (r - q)\, dt + \sigma dL_t^{\alpha,-1}, \quad t \in [0, T], \tag{4}$$

where $L_t^{\alpha,\beta}$ denotes a Lévy α-stable process. The increment $dL_t^{\alpha,\beta}$ is distributed α-stable with zero drift, dispersion of $dt^{1/\alpha}$, and a skew parameter β: $L_\alpha\left(0, dt^{1/\alpha}, \beta\right.$. In (4), we set $\beta = -1$ to achieve finite moments for stock prices and negative skewness in the return density. We further restrict $\alpha \in (1,2)$ in conformity with much empirical evidence.

In general, a stable random variable $x \sim L_\alpha(\theta, \sigma, \beta)$ with $\alpha \in (0,2], \theta \in \Re, \sigma \geq 0$, and $\beta \in [-1,1]$ can be represented by its characteristic function:

$$\phi_x(u) \equiv \mathbb{E}e^{iux} = \begin{cases} \exp\left[iu\theta - |u|^\alpha \sigma^\alpha \left(1 - i\beta\left(\text{ sgn } u\right)\tan\frac{\pi\alpha}{2}\right)\right] & \text{if} \quad \alpha \neq 1, \\ \exp\left[iu\theta - |u|\sigma\left(1 + i\beta\frac{2}{\pi}\left(\text{ sgn } u\right)\ln|u|\right.\right] & \text{if} \quad \alpha = 1. \end{cases} \tag{5}$$

When $\alpha < 2$, the tails are "fat" and the tail probabilities behave like $\lambda^{-\alpha}$:

$$\lim_{\lambda\to\infty} \lambda^\alpha P(x > \lambda) \quad = \quad C_\alpha \frac{1+\beta}{2}\sigma^\alpha, \quad \lim_{\lambda\to\infty} \lambda^\alpha P(x < -\lambda) = C_\alpha \frac{1-\beta}{2}\sigma^\alpha,$$

where:

$$C_\alpha \equiv \left(\int_0^\infty x^\alpha \sin x dx\right)^{-1} = \begin{cases} \frac{1-\alpha}{\Gamma(2-\alpha)\cos(\pi\alpha/2)} & \text{if} \quad \alpha \neq 1, \\ 2/\pi & \text{if} \quad \alpha = 1. \end{cases}$$

The only exception is if $|\beta| = 1$. When $\beta = -1$, the right tail is a "thin" decaying faster than $\lambda^{-\alpha}$. As $\lambda \to \infty$,

$$P(x > \lambda) \sim \begin{cases} \frac{1}{\sqrt{2\pi\alpha(\alpha-1)}}\left(\frac{\lambda}{\alpha\hat{\sigma}_\alpha}\right.^{-\alpha/(2(\alpha-1))} \exp\left(-(\alpha-1)\left(\frac{\lambda}{\alpha\hat{\sigma}_\alpha}\right.^{\alpha/(\alpha-1)}\right) & \text{if} \quad \alpha > 1, \\ \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{(\pi/2\sigma)\lambda - 1}{2} - e^{\pi/2\sigma)\lambda - 1}\right. & \text{if} \quad \alpha = 1, \\ 0 & \text{if} \quad \alpha < 1, \end{cases}$$

where:

$$\hat{\sigma}_\alpha = \sigma\left(\cos\frac{\pi}{2}(2-\alpha)\right)^{-1/\alpha}.$$

The same formula applies to the left tail when $\beta = 1$ because $P(x < -\lambda) = P(-x > \lambda)$ and $x \sim L_\alpha(0, \sigma, 1)$ implies $-x \sim L_\alpha(0, \sigma, -1)$. See Zolotarev (1986), Theorem 2.5.3.

As a result of this tail behavior, the Laplace transform of an α-stable variable $x \sim L_\alpha(0, \sigma, \beta)$, is not finite unless $\beta = 1$. When $\beta = 1$, the Laplace transform is given by:

$$\mathcal{L}_x(\lambda) \equiv \mathbb{E}\left[e^{-\lambda x}\right] = \begin{cases} \exp\left(-\lambda^\alpha \sigma^\alpha \sec\frac{\alpha\pi}{2}\right) & \text{if} \quad \alpha = 1, \\ \exp\left(-\sigma\frac{2}{\pi}\lambda \ln\lambda\right. & \text{if} \quad \alpha = 1, \end{cases} \qquad \mathcal{R}\lambda > 0. \tag{6}$$

We are interested in the distribution of the log return $\ln S_T/S_t$ over the horizon $\tau = T - t$:

$$s_\tau \equiv \ln S_{t+\tau}/S_t = \mu\tau + \sigma\left(L_T^{\alpha,-1} - L_t^{\alpha,-1}\right. \quad ,$$

where μ is determined by the no-arbitrage condition under the measure Q. The conditional characteristic function of the return s_τ can be inferred from (5):

$$\phi_s(u) \equiv \mathbb{E}_t\left[e^{ius_\tau}\right] = \exp\left(iu\mu\tau - \tau(iu\sigma)^\alpha \sec\frac{\pi\alpha}{2}\right), \qquad 1 < \alpha < 2. \tag{7}$$

Furthermore, the nth moment of S_T can be represented as a Laplace transform in (6):

$$\mathbb{E}_t\left[S_T^n\right] \equiv S_t^n \mathbb{E}_t\left[e^{ns_\tau}\right] = S_t^n \exp\left(n\mu\tau - \tau(n\sigma)^\alpha \sec\frac{\pi\alpha}{2}\right), \tag{8}$$

which is finite for all n. Note that $\beta = -1$ is the only α-stable case with $\alpha < 2$ where we have finite moments for S_T.

Setting $n = 1$ in (8), we have:

$$\mathbb{E}_t\left[S_T\right] = S_t \exp\left(\mu\tau - \tau\sigma^\alpha \sec\frac{\pi\alpha}{2}\right).$$

By no-arbitrage, $\mathbb{E}_t\left[S_T\right] = S_t e^{(r-q)\tau}$ under measure Q, and so:

$$\mu = r - q + \sigma^\alpha \sec\frac{\pi\alpha}{2}. \tag{9}$$

Given the characteristic function in (7) and the no-arbitrage condition in (9), we can apply the fast Fourier transform (FFT) method of Carr and Madan (1999) to simultaneously value European calls at a whole spectrum of strike prices.

Our parsimonious FMLS model has only two free parameters σ and α, but it can generate a rich array of risk-neutral densities or volatility smirks. In particular, the model is capable of generating the maturity pattern in the volatility smirk which is observed in US index options. The top panel of Figure 5 illustrates that the implied volatility smirk computed from the FMLS process does not flatten out significantly as maturity increases. In contrast, the bottom panel depicts the case of a typical pure jump model, i.e. the variance-gamma model of Madan, Carr, and Chang (1998). As maturity increases, the implied volatility smirk flattens out significantly.

Figure 6 illustrates that the slope of the smirk is directly linked to the tail index α. The smirk becomes flatter as the tail index α becomes closer to its normal limit $\alpha = 2$. The fact that the tail index α is invariant under time aggregation (self-similarity) is the key reason that volatility smirk does not flatten out as maturity increases.

IV. Model Calibration

In this section, we calibrate the FMLS model to the nonparametrically smoothed implied volatility surface in Figure 1. In particular, we transform the smoothed implied volatility surface to out-of-the-money option prices using a constant interest rate and dividend yield of $r = 7.33\%$ and $q = 1.17\%$ respectively. We then determine the parameters by minimizing the sum of squared pricing errors (sse):

$$\text{sse} = \min_{\Theta} \sum_{i=1}^{N_C} w_i^C \left(C_i - \hat{C}_i(\Theta) \right)^2 + \sum_{i=1}^{N_P} w_i^P \left(P_i - \hat{P}_i(\Theta) \right)^2, \tag{10}$$

where C_i and P_i are realized out-of-the-money call and put prices quoted as a percentage of spot, and $\hat{C}_i(\Theta)$ and $\hat{P}_i(\Theta)$ are their respective counterparts implied from the model given the parameter vector Θ. N_C and N_P are respectively the number of out-of-the-money call and put options.

In a maximum likelihood framework, where errors attached to the observed prices are assumed to be normally distributed with mean zero and variance V_i, the weight would be given by $w_i = 1/V_i$. Unfortunately, the error variance for each price quote is unknown. Bliss

and Panigirtzoglou (2000) provide a detailed discussion on the sources of errors and their implications for the weights. We used equal weights across option prices, assuming equal variance for all pricing errors.

For the estimation, we use option prices on three maturities: 1, 6, and 12 months. For each maturity, we choose 8 strikes which are evenly distributed in $\ln K/S$ and fall roughly within the moneyness range of $[-2, 1]$. In particular,

$$\ln(K/S) = \begin{pmatrix} -0.1841 & -0.3682 & -0.4909 \\ -0.1534 & -0.3068 & -0.3988 \\ -0.1227 & -0.2454 & -0.3068 \\ -0.0920 & -0.1841 & -0.2148 \\ -0.0614 & -0.1227 & -0.1227 \\ -0.0307 & -0.0614 & -0.0307 \\ 0 & 0 & 0.0614 \\ 0.0307 & 0.0614 & 0.1534 \end{pmatrix}.$$

The three columns correspond to the three maturities.

In addition to our two-parameter FMLS model, we also calibrate the popular three-parameter variance-gamma (VG) model of Madan, Carr, and Chang (1998) and the traditional four-parameter jump-diffusion (MJD) model of Merton (1976). Under the VG model, the risk-neutral log price process is:

$$d \ln S_t = (r - q + \mu)dt + \omega d\widetilde{t} + \sigma dW_{\widetilde{t}},$$

where $\widetilde{t}$ denotes a subordinated gamma time, i.e. $d\widetilde{t}$ is independently and identically gamma-distributed: $\gamma(dt : 1, \alpha)$, where α is the variance rate. The parameter μ arises from the concavity of the log:

$$\mu = \frac{1}{\alpha} \ln \left[1 - \left(\omega + \frac{1}{2}\sigma^2 \right) \nu \right].$$

The characteristic function of s_τ is

$$\phi_s(u) = e^{iu(r-q)\tau}\left[1-\left(\omega+\frac{1}{2}\sigma^2\right)\alpha\right]^{-\frac{\tau}{\alpha}}\left[1-\left(iu\omega-\frac{1}{2}\sigma^2u^2\right)\alpha\right]^{-\frac{\tau}{\alpha}}. \tag{11}$$

The stock price process under the MJD model is, under measure Q,

$$dS_t/S_t = (r-q-\alpha g)dt + \sigma dW_t + \left(e^k-1\right.\ \left.\right) dJ(\alpha),$$

where W_t is a Wiener process capturing the continuous part of the price movement while $J(\alpha)$ is a Poisson jump process with intensity α. The jump magnitude is captured by the normally distributed random variable $k \sim N(\omega, \eta^2)$. The parameter $g = \mathbb{E}\left[e^k - 1\right] = e^{\omega+\frac{1}{2}\eta^2}$ is the mean percentage jump conditional on one jump occuring. The characteristic function of $s_\tau \equiv \ln S_{t+\tau}/S_t$ is

$$\phi_s(u) = \exp\left[iu\left(r-q-\alpha g-\frac{1}{2}\sigma^2\right)\tau - \frac{1}{2}u^2\sigma^2\tau + \alpha\left(e^{iu\omega-\frac{1}{2}u^2\eta^2}-1\right.\ \left.\right)\tau\right]. \tag{12}$$

Table I summarizes the calibration results and demonstrates that our FMLS model performs best in having the smallest sum of squared pricing errors (sse), while it has the least number of parameters (two). Figures 7, 8, and 9 depict their performance in matching the out-of-the-money option prices and the implied volatility smirks at different maturities. The Achilles heel of the MJD model and the VG model is that, for a given set of parameters, they cannot simultaneously fit the implied volatility smirk at different maturities. The same weakness holds for all pure jump Lévy processes with finite (but nonzero) skewness and kurtosis for log stock returns, since these moments decrease with maturity by virtue of the central limit theorem. Correspondingly, the volatility smirk flattens out as maturity increases. Our FMLS model relieves the stress in the maturity pattern by adopting a self-similar process. It matches the slope of the volatility smirk by an optimal choice of the tail index α, and guarantees no arbitrage and finite option prices by incorporating maximum skewness to the Lévy α-stable motion.

V. Extensions

The FMLS model compares favorably over other jump models in preventing the implied volatility smirk from flattening out as maturity increases. It is a very parsimonious model that meets three requirements needed to explain the US index options market: (1) self-similarity, (2) non-normality and (3) martingality. The martingality condition implies, *inter alia* that the risk-neutral mean of option prices should be finite. The Black-Scholes model delivers self-similarity and martingality, but implies a normal distribution for asset returns. Many jump or jump diffusion models generate the observed non-normality and meet martingality, but fail to deliver self-similarity and thus fail to capture the maturity pattern of the implied smirk. Previous work on α-stable motions violated the martingality condition and hence either generated suspicious results or had to resort to utility specifications to reconcile a reasonable risk premium. The FMLS model is the first to capture all three features simultaneously, and does so parsimoniously.

As the FMLS model is a pure jump model with iid Lévy increments, it does not capture the often-documented "volatility" clustering. Thus, the FMLS model should be regarded as a springboard for further extensions which capture this phenomenon. By treating the α-stable process as a generalization of the standard Brownian driver extensions of our model are straightforward. In what follows, we will briefly discuss some of these potential extensions and their relative merits in practice. To save space, we do not report the the derivations and calibration results of these extended models. However, these results are available upon request.

A. Pure Jump Versus Jump Diffusion

While the FMLS model is a pure jump model, a diffusion component can be added to the model in a very straightforward manner:

$$dS_t/S_t = (r-q)dt + \sigma_1 dL_t^{\alpha,-1} + \sigma_2 dL_t^2,$$

where L_t^2 is a Brownian motion with variance $2t$ and is a degenerate case of an α-stable motion with $\alpha = 2$.

The characteristic function of the log return, $s_\tau \equiv \ln S_{t+\tau}/S_t$, can be obtained as a direct extension of (7):

$$\phi_s(u) \equiv \mathbb{E}_t\left[e^{ius_\tau}\right] = \exp\left(iu\left(r - q + \sigma_1^\alpha \sec\frac{\pi\alpha}{2} - \sigma_2^2\right)\tau - \tau\left(iu\sigma_1\right)^\alpha \sec\frac{\pi\alpha}{2} - \tau(u\sigma_2)^2\right).$$

Option pricing is therefore also straightforward. Repeated calibration exercises, however, invariably generate an estimate for σ_2 that is indistinguishable from zero. This result coincides with the recent empirical evidence of Carr, Geman, Madan, and Yor (2000), who also find that the explanatory power of their pure jump process is not enhanced by the addition of a diffusion component.

These results contrast sharply with empirical results on the traditional Poisson compound jump-diffusion models such as Merton (1976) (MJD). Indeed, in our calibration of the MJD model in Table I, the diffusion component is a significant component and accounts for 17.9% of the variation.[9] The sharp contrast comes from the difference in the fine structure of the jump specification in these models. In both the FMLS and CGMY models, as the jump size approaches zero, the arrival rate for jumps of this size approaches infinity. As a result, the aggregate arrival rate across jumps of all sizes is infinite. In contrast, in the Poisson compound jump model (MJD), the aggregate arrival rate of jumps is constant and finite at α. The fact that there are an "infinite" number of very small jumps in the FMLS and CGMY model relieves their need for a diffusion component. This, however, is not the case for the MJD model. In short, while adding a diffusion component to our FMLS model poses no difficulty in terms of implementation, this component has no additional explanatory power in terms of determining option prices.

[9]Under the MJD model, the variance of the log return is given by

$$\kappa_2 = \left[\sigma^2 + \alpha(\omega^2 + \eta^2)\right]\tau.$$

The fraction of variation explained by the diffusion component is σ^2/κ_2.

B. Refining the Structure of Jumps

Under the FMLS model, the structure of jumps is controlled by a maximum negatively skewed α-stable Lévy motion. The arrival rate of jumps of size x is given by its Lévy measure:

$$\nu(dx) = c_{\pm}|x|^{-\alpha-1},$$

where $c_+ = 0$ for $x \in (0, \infty)$ and $c_- = 1/\Gamma(-\alpha)$ for $x \in (-\infty, 0]$. Although the random component of the process features only negative jumps, the predictable component of the process compensates by so much that the support of the distribution for any positive time is the entire real line.

Our calibration results indicate that the FMLS jump structure adequately captures the non-normality of asset returns, i.e. the implied volatility smirk. Nonetheless, to improve the fit of an implied smirk at any maturity, one can further refine the jump structure through subordination:

$$s_\tau = (r - q + \mu)\,\tau + \omega\tilde{\tau} + \sigma L_{\tilde{\tau}}^{\alpha,-2},$$

where the time length of the α-stable Lévy motion is subordinated by a subordinator $\tilde{\tau}$. As an example, we derive the direct α-stable counterpart of the VG model of Madan, Carr, and Chang (1998), where we replace their gamma subordination of a standard Brownian motion by a gamma subordination of a standard α-stable Lévy motion. That is, the subordinator $\tilde{\tau}$ is assumed to follow an independent gamma process with unit mean rate and variance rate η: $\gamma(\tau; 1, \eta)$. As in Madan, Carr, and Chang (1998), we also add a free drift parameter ω to generate asymmetry. The parameter μ is an adjustment term to guarantee no-arbitrage. Under such a specification, the characteristic function of the log return s_τ is given by

$$\phi_s(u) = e^{iu(r-q+\mu)\tau}\left(1 - \left(iu\omega - (iu\sigma)^\alpha \sec\frac{\pi\alpha}{2}\right)\eta\right)^{-\frac{\tau}{\eta}},$$

with μ given by:

$$\mu = \frac{1}{\eta}\ln\left[1 - \left(\omega - \sigma^\alpha \sec\frac{\pi\alpha}{2}\right)\eta\right].$$

We calibrate this gamma-subordinated FMLS model to the implied volatility smirk. The results (not reported) indicate that while its capability of fitting one implied volatility smirk at a given maturity is improved, the self-similarity structure is lost and hence the stability across maturities is actually not as good as the benchmark model.

C. Stochastic "Volatility"

Under the FMLS model, the variance of the log return is infinite. However, an analogous parameter σ^α is referred to as the α-dispersion and this statistic captures the "variability" of the log return. Through the use of stochastic time change, a stochastic dispersion process can be readily incorporated into the FMLS model in order to capture the "volatility clustering" phenomenon.

As an example, we incorporate the extensively cited square-root model of Heston (1993) into the FMLS framework:

$$\begin{aligned} dS_t/S_t &= (r-q)dt + v_t^{1/\alpha} dL_t^{\alpha,-1}; \\ dv_t &= \kappa(\theta - v_t)dt + \sqrt{\beta v_t} dW_t. \end{aligned} \tag{13}$$

The characteristic function of the log return s_τ is now given by:

$$\phi_s(u) = \exp\left[iu(r-q)\tau - b(\tau)v_t - c(\tau)\right]$$

with

$$\begin{aligned} b(\tau) &= \frac{2\lambda\left(1-e^{-\eta\tau}\right)}{2\eta - (\eta-\kappa)\left(1-e^{-\eta\tau}\right)}, \\ c(\tau) &= \frac{\kappa\theta}{\beta}\left[2\ln\left(1 - \frac{\eta-\kappa}{2\eta}\left(1-e^{-\tau\eta}\right)\right) + (\eta-\kappa)\tau\right], \end{aligned}$$

with

$$\eta = \sqrt{\kappa^2 + 2\beta\lambda}, \quad \lambda = \left[(iu)^\alpha - iu\right]\sec\frac{\pi\alpha}{2}.$$

The model has four free parameters: $\Theta \equiv (\alpha, \kappa, \theta, \beta)$ and two state variables (S_t, v_t). Since one cannot directly observe v_t, one needs to infer it from the option price data each day by minimizing the sum of squared pricing errors, given a choice of the parameters.

We also calibrate this model to the nonparametrically smoothed mean implied volatility surface. The results (not reported) indicate that, even with two more free parameters, the model does not significantly outperform the FMLS benchmark in fitting the mean implied volatility surface across dates. however, its greater flexibility does show up in fitting the variation of the implied volatility smirk across maturity at each date.

VI. Final Thoughts

Modeling U.S. equity index returns with our FMLS process is very appealing because the model-implied skewness and kurtosis of the index return are not finite, and hence do not decrease over maturity. It thus provides a parsimonious way of capturing the implied volatility surface along both the moneyness and maturity dimensions.

Further research should go into documenting the maturity pattern for other indices and for individual stocks. Our preliminary investigations suggest that stochastic time change can be profitably used to improve the fit of our model even further.

References

Abadir, K., and M. Rockinger, 1997, Density-embedding functions, Working paper, HEC.

Abken, Peter A., Dilip B. Madan, and Sailesh Ramamurite, 1996a, Estimation of risk-neutral and statistical densities by hermite polynomial approximation: with an application to eurodollar futures options, Working paper, Federal Reserve Bank of Atlanta.

Abken, Peter A., Dilip B. Madan, and Sailesh Ramamurite, 1996b, Pricing s&p 500 index options using a hilbert space basis, Working paper, Federal Reserve Bank of Atlanta.

Ait-Sahalia, Yacine, and Andrew Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance* 53, 499–547.

Aparicio, S., and S. Hodges, 1998, Implied risk-neutral distribution: A comparison of estimating methods, Working paper, Warwick University.

Backus, David, Silverio Foresi, and Liuren Wu, 1997, Accounting for biases in black-scholes, Working paper, New York University.

Bates, David, 1991, The crash of 87: Was it expected? the evidence from option markets, *Journal of Finance* 46, 1009–1044.

Black, Fisher, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.

Bliss, Robert R., and Nikolaos Panigirtzoglou, 2000, Testing the stability of implied probability density functions, manuscript, Federal Reserve bank of Chicago and Bank of England.

Bookstaber, Richard M., and James B. McDonald, 1987, A general distribution for describing security price returns, *Journal of Business* 60, 401–424.

Brenner, Micahel, and Y. Eom, 1997, No-arbitrage option pricing: New evidence on the validity of the martingale property, Working paper, New York University.

Carr, Peter, and Dilip B. Madan, 1999, Option valuation using the fast fourier transform, *Journal of Computational Finance* 2, 61–73.

Carr, Peter P., Hélyette Geman, Dilip B. Madan, and Macr Yor, 2000, The fine structure of asset returns: an empirical investigation, manuscript, Banc of America Securities, Université Paris IX Dauphine, University of Maryland, and Université Paris VI.

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, 175–192.

Corrado, Charles J., and Tie Su, 1997, Implied volatility skews and stock return skewness and kurtosis implied by stock option prices, *European Journal of Finance* 3, 73–85.

Das, Sanjiv Ranjan, and Rangarajan Sundaram, 1999, Of smiles and smirks: A term structure perspective, *Journal of Financial and Qunatitative Analysis* 34, 211–240.

Duan, Jin-Chuan, 1995, The garch option pricing model, *Mathematical Finance* 5, 13–32.

Duan, Jin-Chuan, G. Gauthier, and Jean-Guy Simonato, 1999, An analytical approximation for the garch option pricing model, *Journal of Computational Finance* 2, 75–116.

Duan, Jin-Chuan, and Jean-Guy Simonato, 2000, American option pricing under garch by a markov chain approximation, *Journal of Economic Dynamics and Control* forthcoming.

Duan, Jin-Chuan, and J. Wei, 1999, Pricing foreign currency and cross-currency options under garch, *Journal of Derivatives* Fall, 51–63.

Fama, Eugene F., 1965, The behavior of stock market prices, *Journal of Business* 38, 34–105.

Feller, William, 1971, *An Introduction to Probability Theory and Its Applications* vol. II. (John Wiley & Sons New York) 2nd edn.

Heston, Stephen, 1993, Closed-form solution for options with stochastic volatility, with application to bond and currency options, *Review of Financial Studies* 6, 327–343.

Hull, John, and Alan White, 1987, The pricing of options on assets wityh stochastic volatilities, *Journal of Finance* 42, 281–300.

Jackwerth, Jens Carsten, and Mark Rubinstein, 1996, Recovering probability distributions from contemporary security prices, *Journal of Finance* 51, 347–369.

Janicki, Aleksander, and Aleksander Weron, 1994, *Simulation and Chaotic Behavior of α-Stable Stochastic Processes*. (Marcel Dekker New York).

Janicki, Aleksander W., Ivilina Popova, Peter H. Ritchken, and W.A. Woyczynski, 1997, Option pricing bounds in an α-stable security market, *Communication in Statistics – Stochastic Models* 13, 817–839.

Jarrow, Robert, and Andrew Rudd, 1982, Approximate option valuation for arbitrary stochastic processes, *Journal of Financial Economics* 10, 347–369.

Jondeau, E., and M. Rockinger, 1998, Estimating gram-charlier expansions under positivity constraints, Working paper, HEC.

Konikov, Mikhail, and Dilip B. Madan, 2000, Pricing options of all strikes and maturities using a generalization of the vg model, Working paper, University of Maryland.

Kou, Steve G., 1999, A jump-diffusion model for option pricing with three properties: Leptokurtic feature, volatility smile, and analytical tractability, manuscript, Columbia University.

Lévy, Pierre, 1925, *Calcul des Probabilités*. (Gauthier-Villars Paris).

Longstaff, Francis, 1995, Option pricing and martingale restriction, *Review of Financial Studies* 8, 1091–1124.

Madan, Dilip B., Peter P. Carr, and Eric Chang, 1998, The variance gamma process and option pricing, *European Financial Review* 2, 79–105.

Madan, Dilip B., and F. Milne, 1991, Option pricing with vg martingale components, *Mathematical Finance* 1, 39–56.

Mandelbrot, Benoit B., 1963a, New methods in statistical economics, *Journal of Polictical Economy* pp. 421–440.

Mandelbrot, Benoit B., 1963b, The variation of certain speculative prices, *Journal of Business* 36, 394–419.

McCulloch, John H., 1987, Foreign exchange option pricing with log-stable uncertainty, in S. J. Khoury, and A. Ghosh, eds.: *Recent Developments in INternational Banking and Finance* (Lexington Books, Lexington, MA).

McCulloch, John H., 1996, Financial aplications of stable distributions, in G. S. Maddala, and C. R. Rao, eds.: *Statistical Methods in Finance* (Elsevier Science, Amsterdam).

McCulloch, John M., 1978, Continuous time process with stable increments, *Journal of Business* 51, 601–619.

Melick, W., and C. Thomas, 1997, Recovering an asset's implied pdf from option prices: An application to crude oil during the gulf crisis, *Journal of Financial and Quantative Analysis* 32, 91–115.

Merton, Robert C., 1976, Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 125–144.

Mittnik, Stefan, Svetlozar T. Rachev, and Marc S. Paolella, 1998, Stable paretian modeling in finance: some empirical and theoretical aspects, in Robert J. Adler, Raisa E. Feldman, and Murad S. Taqqu, eds.: *A Practical Guide To Heavy Tails* (Birkhäuser, Boston/Basel/Berline).

Popova, Ivilina, and Peter Ritchken, 1998, On bounding option prices in paretian stable markets, *The Journal of Derivatives* pp. 32–43.

Posner, Steven E., and Moshe Arye Milevsky, 1998, Valuing exotic options by approximating the spd with higher moments, *Journal of Financial Engineering* 7, 109–125.

Potters, M., R. Cont, and J. Bouchaud, 1998, Financial markets as adaptive systems, *European Letters* 41, 239–244.

Ritchey, Robert J., 1990, Call option valuation for discrete normal mixtures, *Journal of Financial Research* 13, 285–295.

Romano, Marc, and Nizar Touzi, 1997, Contingent claims and market completeness in a stochastic volatility model, *Mathematical Finance* 7, 279–302.

Rubinstein, Mark, 1994, Implied binomial trees, *Journal of Finance* 49, 771–818.

Rubinstein, Mark, 1998, Edgeworth binomial trees, *Journal of Derivatives* 5, 20–27.

Samorodnitsky, Gennady, and Murad S. Taqqu, 1994, *Stable non-Gaussian random processes : stochastic models with infinite variance*. (Chapman & Hall New York).

Samuelson, Paul A., 1965, Rational theory of warrant pricing, *Industrial Management Review* 6, 13–31.

Sherrick, B., P. Garcia, and V. Tirupattur, 1995, Recovering probabilistic information from option markets: Tests of distributional assumptions, Working paper, University of Illinois at Urban-Champaign.

Sherrick, B., S. Irwin, and D. Forster, 1992, Option-based evidence of the nonstationarityof the expected s&p 500 futures price distributions, *Journal of Futures Markets* 12, 275–290.

Sherrick, Bruce J., Scott H. Irwin, and D. Lynn Forster, 1996, An examination of option-implied s&p 500 futures price distributions, *Financial Review* 31, 667–694.

Simonoff, Jeffrey S., 1996, *Smoothing Methods in Statistics*. (Springer-Verlag New York).

Zolotarev, Vladimir M., 1986, *One-Dimensional Stable Distributions*. (American Mathematical Society Providence, Rhode Island).

Table I
Estimates of Jump(-Diffusion) Models

Entries are estimates (in the parentheses are standard errors) and sum square pricing errors (sse) of the parameters of models: (i) the FMLS model, (ii) the VG model, and (iii) the MJD model. The characteristic functions of the log stock return under these models are, respectively,

$$\begin{aligned}
\phi_{FMLS}(u) &= \exp\left[iu\left(r - q + \sigma^2 \sec\frac{\pi\alpha}{2}\right)\tau - \tau\left(iu\sigma\right)^{\alpha}\sec\frac{\pi\alpha}{2}\right];\\
\phi_{VG}(u) &= e^{iu(r-q)\tau}\left[1 - \left(\omega + \frac{1}{2}\sigma^2\right)\alpha\right]^{-\frac{\tau}{\alpha}}\left[1 - \left(iu\omega - \frac{1}{2}\sigma^2 u^2\right)\alpha\right]^{-\frac{\tau}{\alpha}};\\
\phi_{MJD}(u) &= \exp\left[iu\left(r - q - \alpha\left(e^{\omega+\frac{1}{2}\eta^2} - 1 \right. - \frac{1}{2}\sigma^2\right)\tau - \frac{1}{2}u^2\sigma^2\tau + \alpha\left(e^{iu\omega-\frac{1}{2}u^2\eta^2} - 1 \right. \tau\right].
\end{aligned}$$

Parameters	FMLS	VG	MJD
α	1.6145	0.7221	1.8604
	(0.0110)	(0.1384)	(0.3523)
σ	0.1486	0.1642	0.1003
	(0.0010)	(0.0168)	(0.0079)
ω	—	−0.1834	−0.0930
	—	(0.0527)	(0.0117)
η	—	—	0.1271
	—	—	(0.0196)
sse($\times 10^5$)	5.8004	18.2319	11.5350

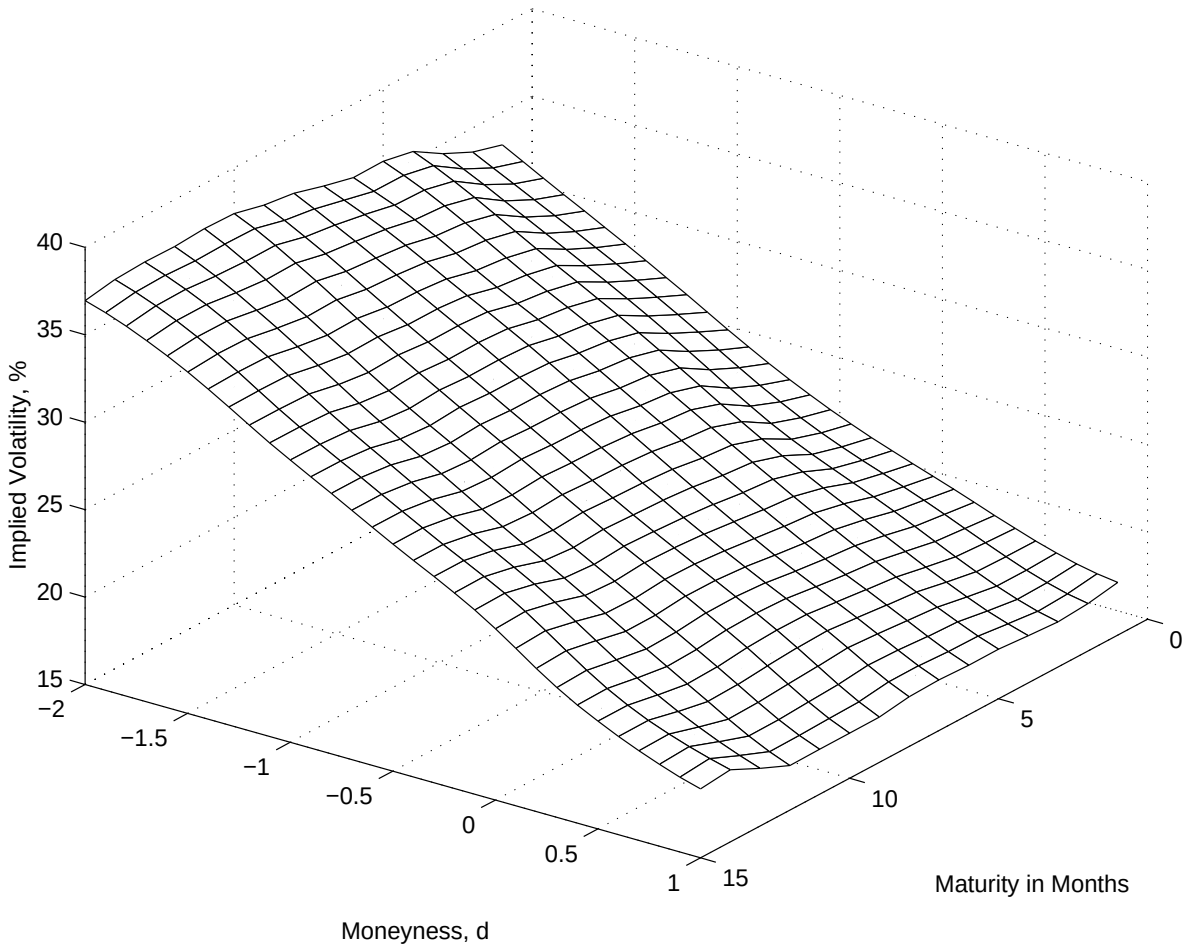

Figure 1. Volatility Smirk in S&P 500 Index Options. The surface is obtained via nonparametric smoothing of daily closing implied volatility quotes (mid point of bid and ask) on S&P 500 index options from April 4th, 1999 to May 31st, 2000 (62,950 observations). Independent Gaussian kernels are used with bandwidth $h_j = \sigma_j n^{-1/6}$. Moneyness, d, is defined as $d = \ln(K/F)/\sigma\sqrt{\tau}$, where $\sigma = 27.4\%$ is the average of all implied volatility quotes.

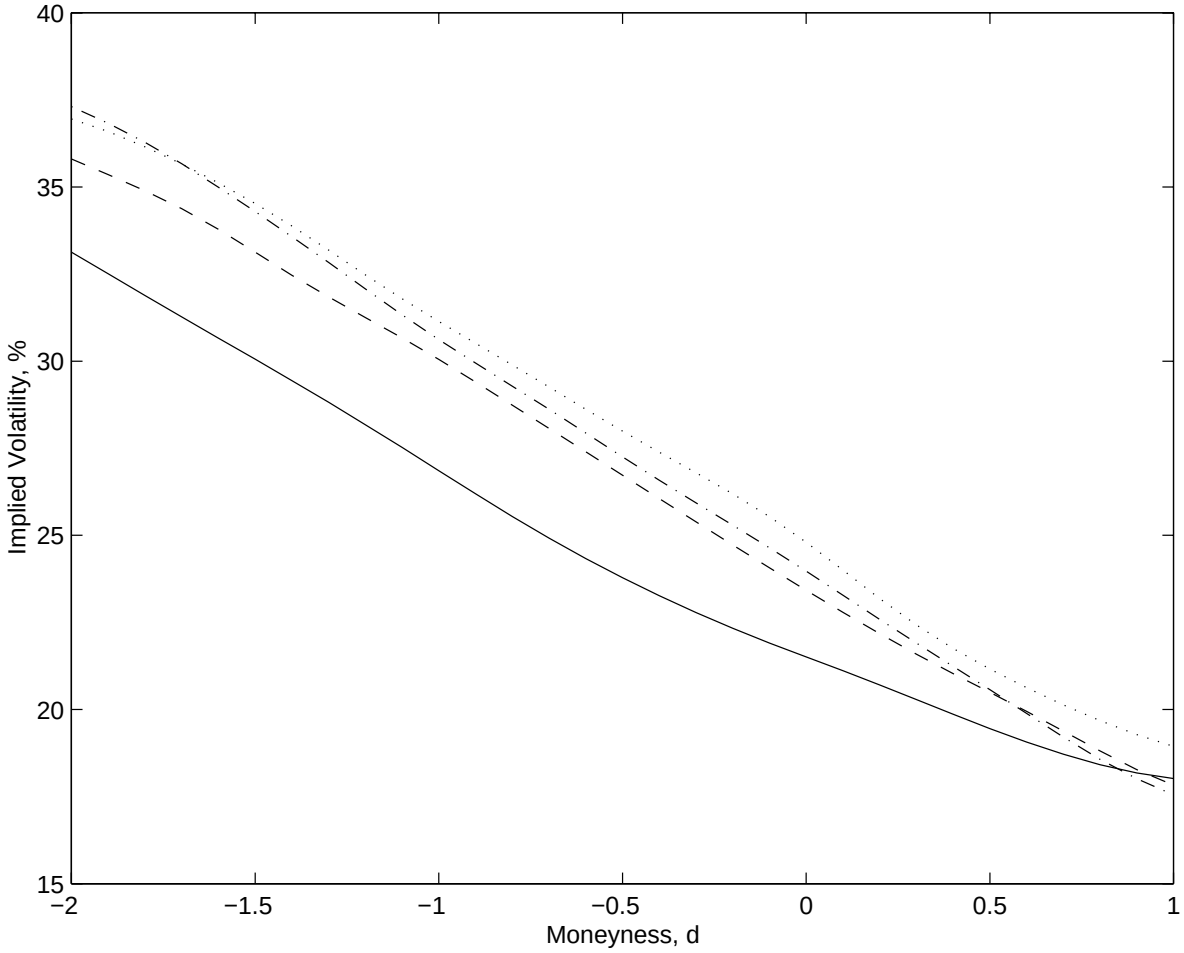

Figure 2. Volatility Smirk in S&P 500 Index Options. Lines are two-dimensional cuts from Figure 1 at maturities of 1 month (solid line), 6 months (dashed line), 12 months (dash-dotted line), and 15 months (dotted line).

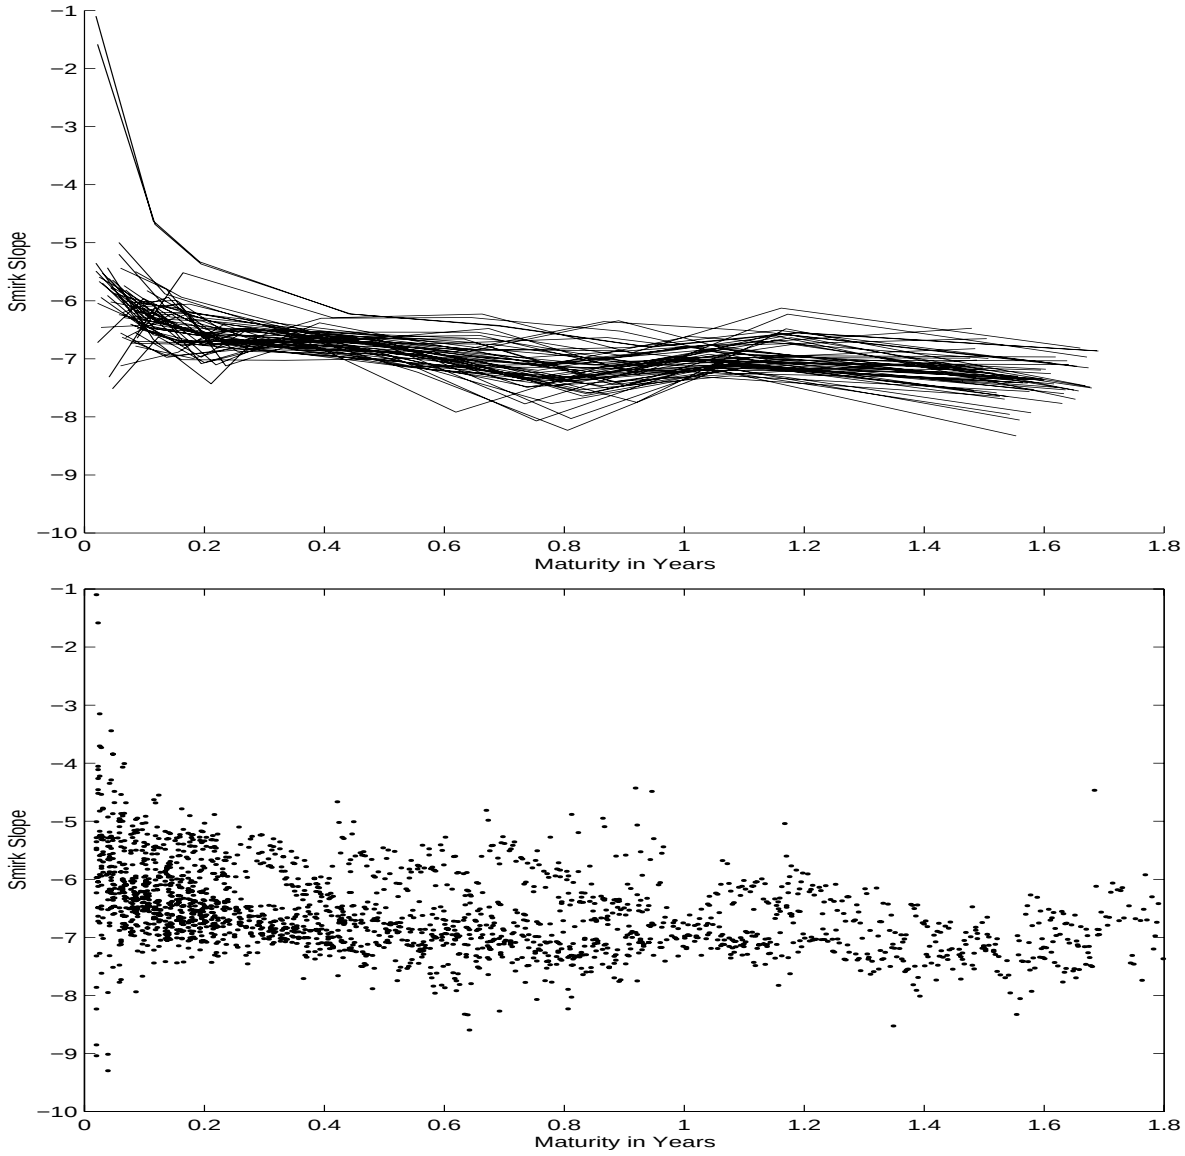

Figure 3. Maturity Effect on Volatility Smirk. Lines in the the top panel are the term structure of the estimated smirk slopes at each day. The dashed line in the bottom panel depicts a linear regression fit to the stacked data of the term structure. At each maturity and each day, the smirk slope is estimated by the following regression:
$IV_j = a_j + b_j d_j + e_j$,
where the moneyness d_j is defined as in Figure 1.

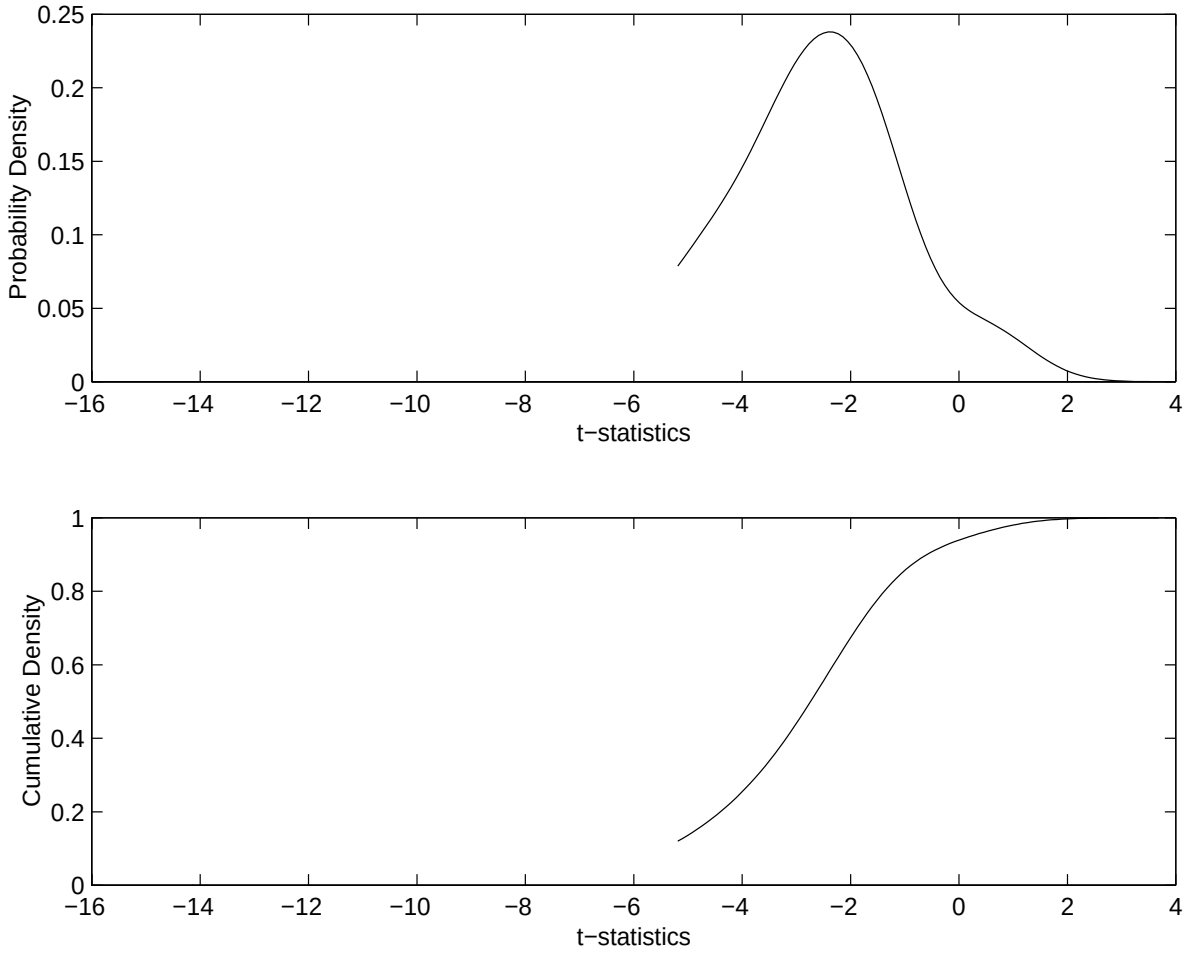

Figure 4. The Term Structure of Implied Volatility Smirk. Solid lines capture the nonparametrically estimated (with Gaussian kernel) distribution (probability density at the top panel and cumulative density at the bottom) of the t-statistics of the β estimates for the 290 days from January 6th, 1999 to May 31sdt, 2000. The dot-dashed lines are implied by the normal density with comparable mean and variance. β captures the term structure of the implied volatility smirk slope b_j:
$b_j = \alpha + \beta\tau_j + e_j$,
where the slope b_j is estimated from the regression:
$IV_j = a_j + b_j d_j + e_j$,
at each day and maturity j.

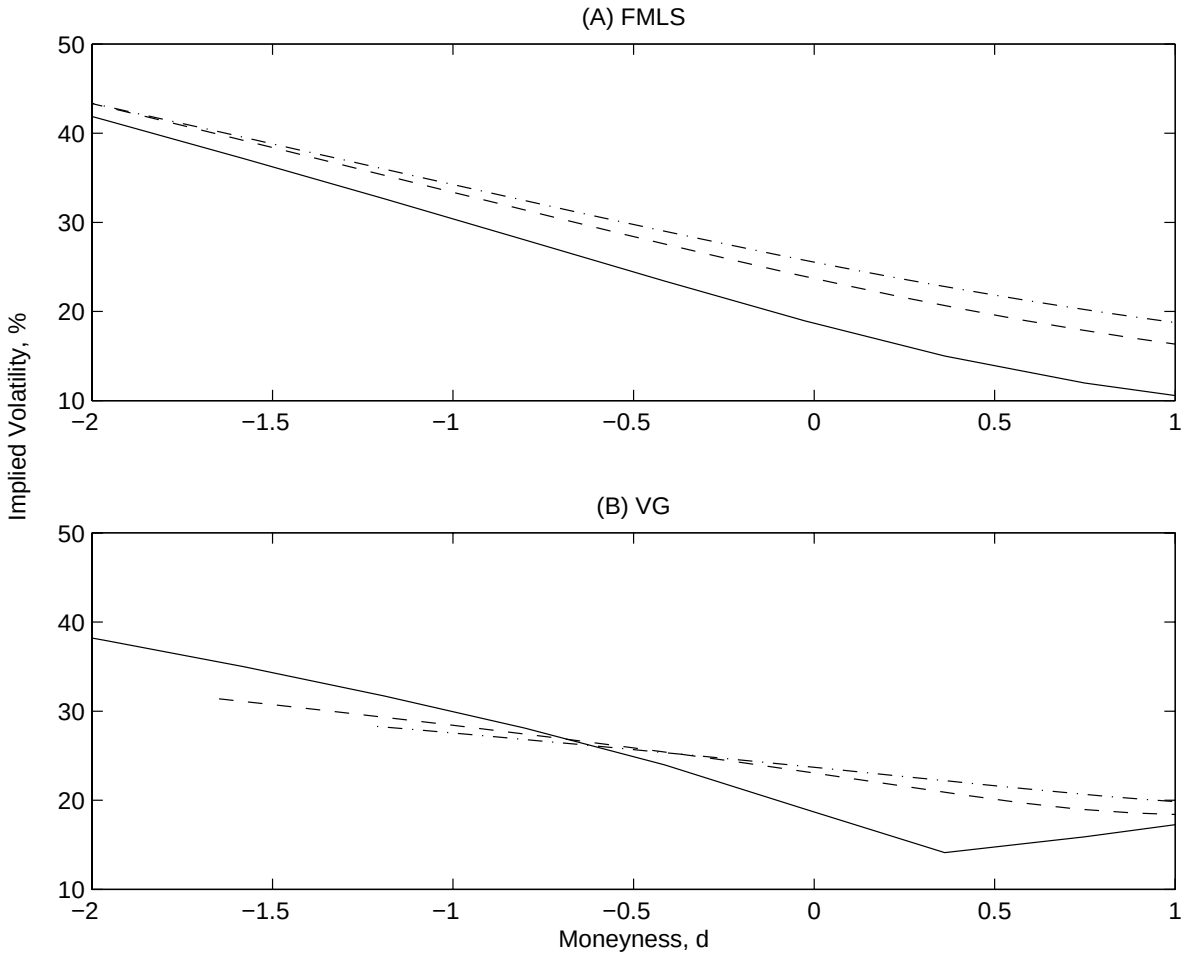

Figure 5. Volatility Smirk Under (a) FMLS Model and (b) VG Model. Implied volatilities are computed from the FMLS model (the top panel) with parameters: $\alpha = 1.4$ and $\sigma = 14\%$ and from the VG model (the bottom panel) with parameters: $\sigma = 14\%$, $\eta = 0.3$, and $\omega = -0.3$. We also set interest rate $r = 7.33\%$ and dividend yields $q = 1.17\%$. Maturity τ equals 1 month (solid line), 6 months (dashed line), and 12 months (dash-dotted line). Moneyness is defined as in Figure 1.

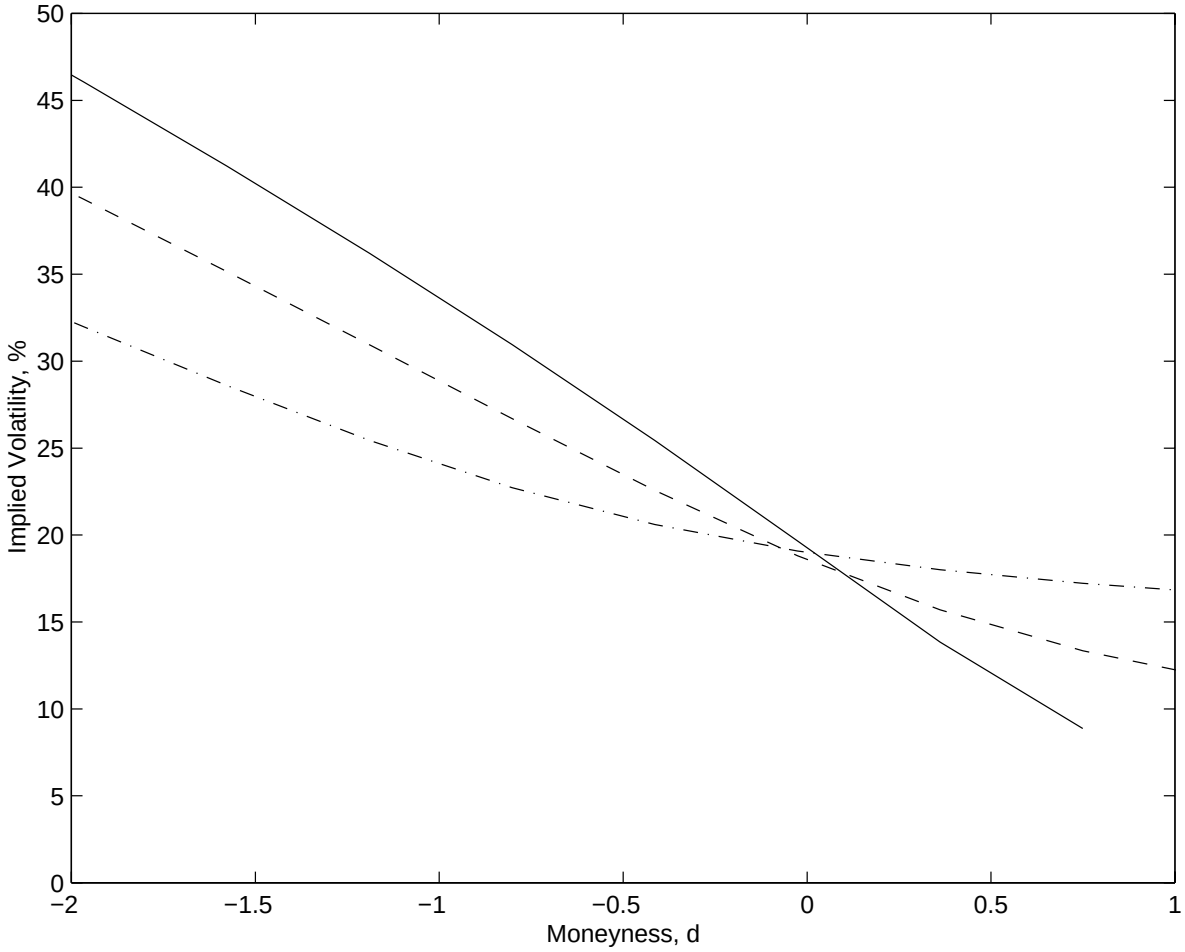

Figure 6. Volatility Smirk Under FMLS Model. Implied volatilities are computed from the FMLS model with α being 1.2 (solid line), 1.5 (dashed line), and 1.8 (dashed-dotted line). We set interest rate $r = 7.33\%$ and dividend yields $q = 1.17\%$. Maturity is 1 month. Money is defined as in Figure 1.

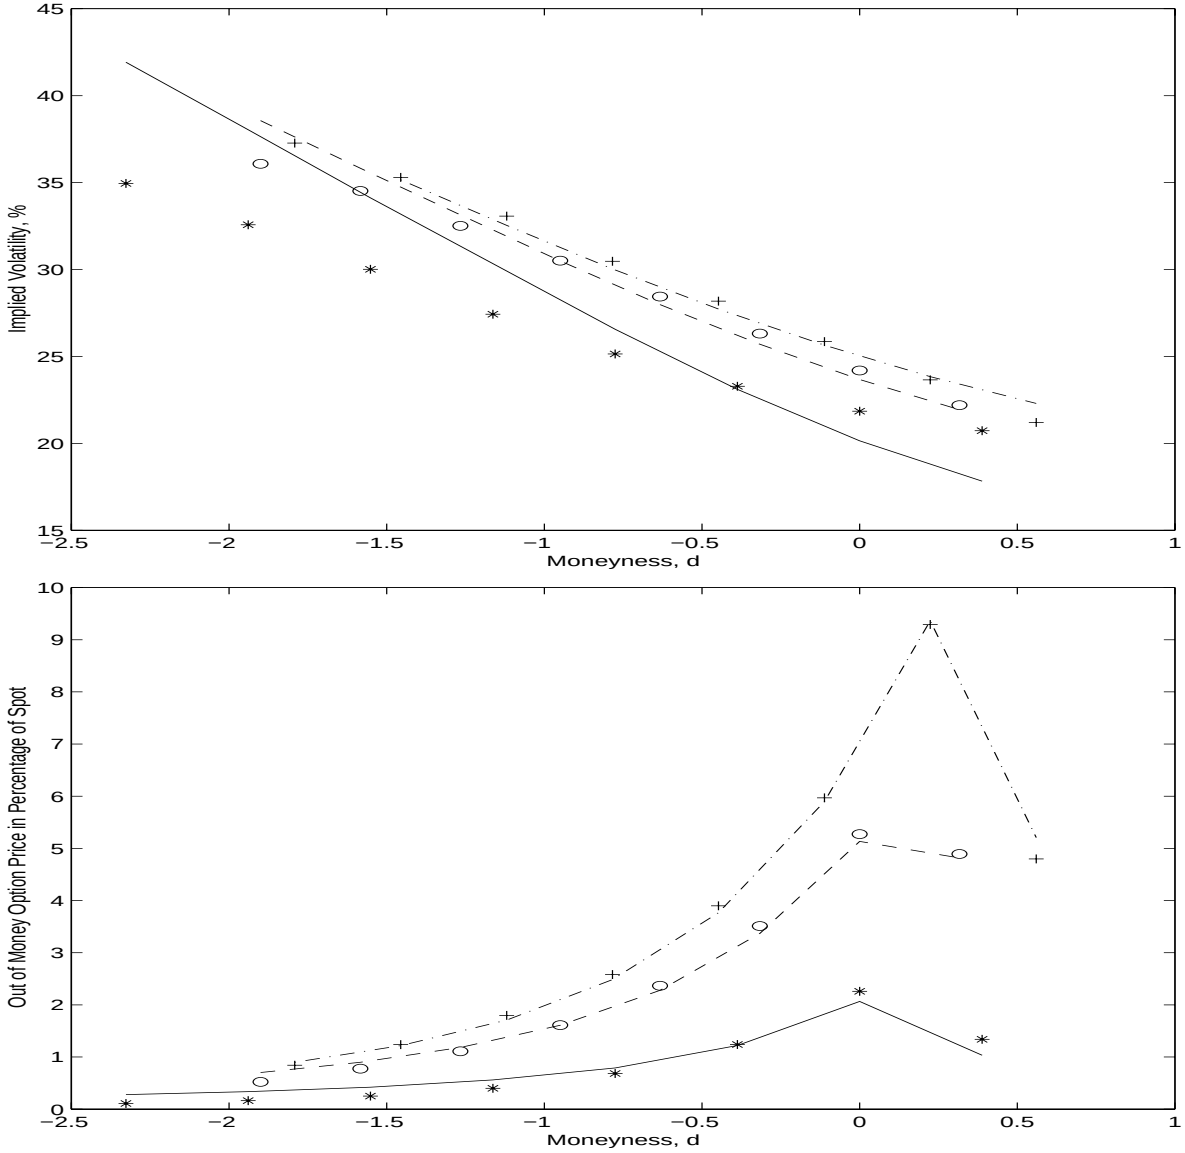

Figure 7. Empirical Performance of the FMLS Model. Stars, circles, and crosses are smoothed implied volatility quotes (top panel) and its corresponding call option prices (in percentage of spot, bottom panel) with maturities of, respectively, 1 month, 6 months, and 12 months. Lines (solid, dashed, and dash-dotted, respectively) are implied by the FMLS model with parameter estimates given in Table I. Interest rate and dividend yield are assumed to be constant across maturities: $r = 7.33\%$ and $q = 1.17\%$. Money is defined as in Figure 1.

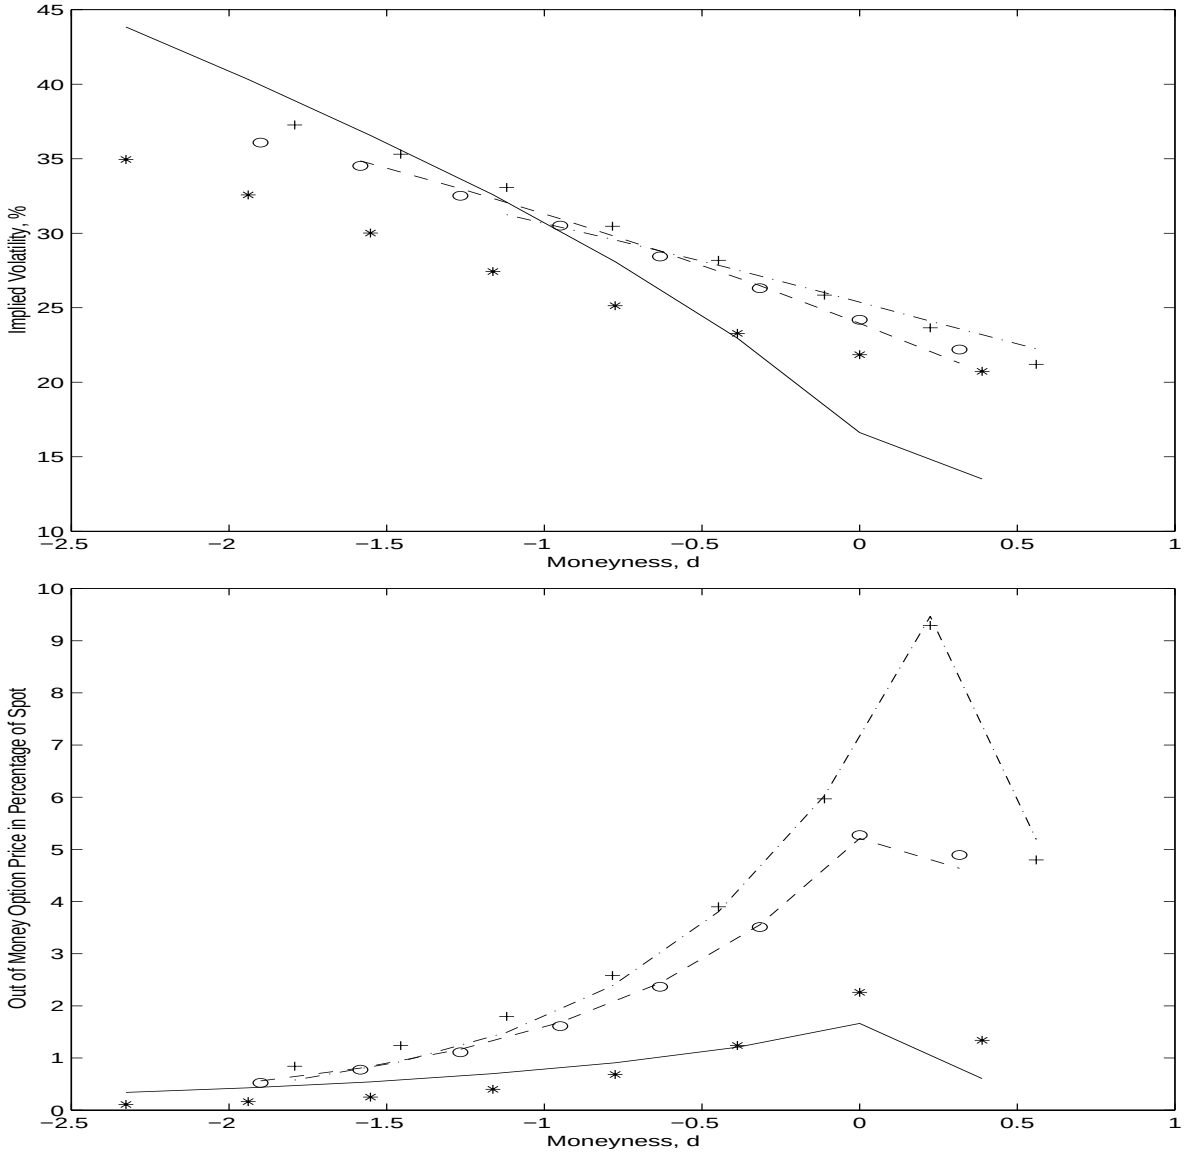

Figure 8. Empirical Performance of the VG Model. Stars, circles, and crosses are smoothed implied volatility quotes (top panel) and its corresponding call option prices (in percentage of spot, bottom panel) with maturities of, respectively, 1 month, 6 months, and 12 months. Lines (solid, dashed, and dash-dotted, respectively) are implied by the VG model with parameter estimates in Table I. Interest rate and dividend yield are assumed to be constant across maturities: $r = 7.33\%$ and $q = 1.17\%$. Money is defined as in Figure 1.

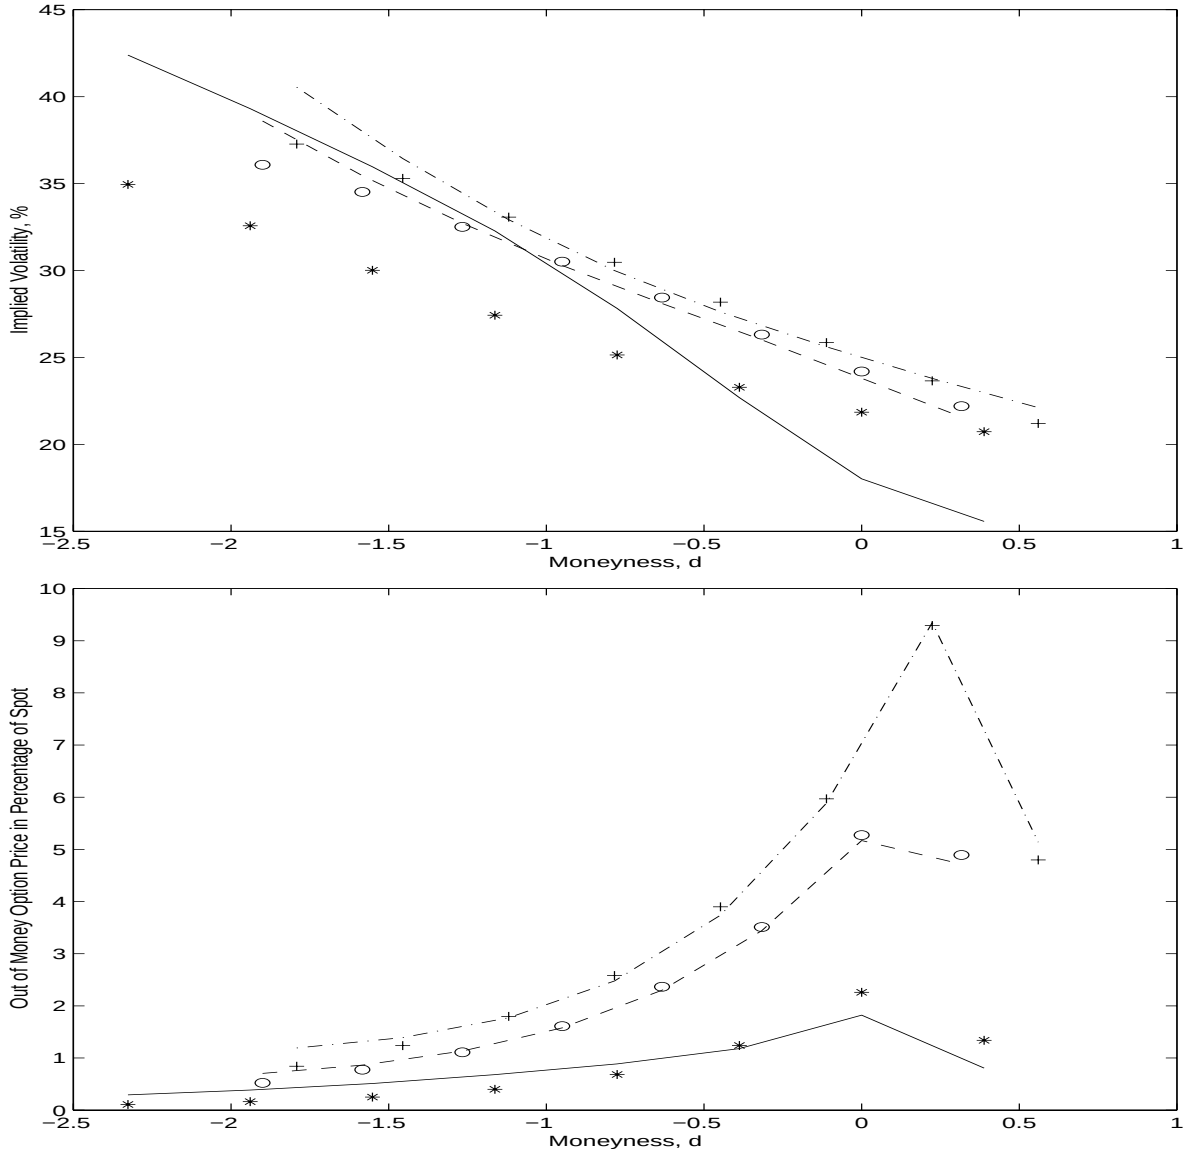

Figure 9. Empirical Performance of the MJD Model. Stars, circles, and crosses are smoothed implied volatility quotes (top panel) and its corresponding call option prices (in percentage of spot, bottom panel) with maturities of, respectively, 1 month, 6 months, and 12 months. Lines (solid, dashed, and dash-dotted, respectively) are implied by the MJD model with parameter estimates given in Table I. Interest rate and dividend yield are assumed to be constant across maturities: $r = 7.33\%$ and $q = 1.17\%$. Money is defined as in Figure 1.