

Extracting Model-Free Volatility from Option Prices: *An Examination of the VIX Index*

GEORGE J. JIANG AND YISONG S. TIAN

GEORGE J. JIANG

is an assistant professor of finance at Eller College of Management, University of Arizona in Tucson, AZ.
gjiang@eller.arizona.edu

YISONG S. TIAN

is an associate professor of finance at Schulich School of Business, York University in Toronto, Ontario, Canada.
ytian@schulich.yorku.ca

The Chicago Board Options Exchange (CBOE) recently redesigned its widely followed VIX volatility index. While the new VIX is conceptually more appealing than its predecessor, the CBOE's implementation of the index is flawed. Using option prices simulated under typical market conditions, we show that the CBOE procedure may underestimate the true volatility by as much as 198 index basis points or overestimate it by as much as 79 index basis points. As each index basis point is worth \$10 per VIX futures contract, these errors are economically significant. More importantly, these errors exhibit predictable patterns in relation to volatility levels. We propose a simple solution to fix the problems, based on a smooth interpolation-extrapolation of the implied volatility function. This alternative method is accurate and robust across a wide range of model specifications and market conditions.

In 1993, the Chicago Board Options Exchange (CBOE) introduced a stock market volatility index called the VIX. Also known as the *investor fear gauge* (see Whaley [1993, 2000]), the VIX captures the market's aggregate expectation of future volatility over the next 30 days. Based on the implied volatility of at-the-money options on the S&P 100 index (OEX), the VIX is calculated and disseminated to market participants in real time throughout the trading day. Recently, the CBOE made several significant

changes to its widely followed volatility index and introduced a revamped VIX on September 22, 2003.¹ A key change is that the new VIX no longer relies on the Black-Scholes-Merton ([1973], BSM hereafter) model.² Instead, it is based on the *fair value of future variance* developed by Demeterfi et al. [1999a, DDKZ] and extracted directly from option prices. This concept of implied variance has its origin in the pioneering work of Breeden and Litzenberger [1978], who show that the entire risk-neutral distribution of the underlying asset can be extracted from option prices. In addition, the new VIX is based on the S&P 500 index (SPX) instead of the OEX index. This switch is timely as the SPX index has become a more important tracking index for institutional investors and derivatives trading.

In this article, we examine the theoretical underpinning of the new VIX and analyze the CBOE's implementation procedure for the index. While the new VIX is conceptually well grounded and more appealing than its predecessor, we find that the CBOE procedure for constructing the new index is flawed. Using option prices calculated under the stochastic volatility with random jumps (SVJ) model, we demonstrate that the CBOE calculation of the VIX index can underestimate the true volatility by as much as 198 index basis points or overestimate it by as much as 79

index basis points. As each index basis point is worth \$10 per VIX futures contract, these approximation errors translate into dollar values ranging from -\$1,980 to +\$790 per contract. More importantly, the range of these errors remains almost unchanged even if the simple BSM model is used to calculate option prices. This finding is particularly worrisome since the BSM model assumes constant volatility and the implied volatility from a single option (such as the original VIX based on at-the-money implied volatilities) would have provided a correct measure of the true volatility. The failure of the CBOE calculation in this simple setting reveals the serious nature of the problems in the CBOE procedure.

In addition, we also find that the approximation errors exhibit predictable patterns due to the impact of truncation and discretization errors. Truncation errors arise from ignoring strike prices beyond the range of listed strike prices while discretization errors are due to the discreteness of listed strike prices and an unusual numerical integration scheme used in the CBOE procedure. At high (low) volatility levels, truncation (discretization) errors dominate and lead to an underestimation (overestimation) of the true volatility. Such predictable patterns can lead to the mispricing of volatility derivatives (e.g., VIX futures, options and variance swaps) as well as the misestimation of volatility risk premium. As the CBOE procedure is likely to become the industry standard for computing risk-neutral variance for other assets and indices, these errors must be corrected in order to ensure the accurate and consistent pricing in the fast growing market for volatility derivatives.

To fix the problems in the CBOE procedure, we propose a simple *smoothing method* for extracting the model-free implied variance from option prices. It is based on the construction of the implied volatility function using an interpolation-extrapolation scheme. Interpolation is implemented between listed strike prices to minimize discretization errors while extrapolation is employed outside the range of listed strike prices to reduce truncation errors. Unlike existing methods (e.g., Carr and Wu [2004] and Jiang and Tian [2005]), we ensure that the constructed implied volatility function is smooth over the entire range of strike prices.

Using the SVJ model with exactly the same parameters, we show that the approximation error from our smoothing method is mostly within 5 index basis points of the true volatility over a wide range of parameter values. Even in the worst case, the error is only 8 index

basis points. In comparison, the corresponding errors from the CBOE procedure vary from +79 index basis points to -198 index basis points. To ensure robustness of our findings, we conduct additional theoretical experiments under several alternative model specifications including the BSM model, Heston's [1993] stochastic volatility model, and Duffie, Pan and Singleton's [2000] stochastic volatility model with time-varying random jumps. The results suggest that the smoothing method is consistently accurate across model specifications and market conditions. In comparison, the CBOE method performs poorly under all model specifications. Finally, the performance of our smoothing method is also supported by an empirical analysis using daily prices of SPX options over the period from January 1996 to May 2004. We find that the differences in the calculated VIX index values using the CBOE and smoothing methods are consistent with the type and range of errors identified in our theoretical experiments.

THE NEW VIX AND MODEL-FREE IMPLIED VARIANCE

We first examine the theoretical underpinning of the new VIX within the broader context of *model-free implied variance*. The pioneering work of Breeden and Litterberger [1978] has laid the foundation for subsequent research on extracting risk-neutral distributions and model-free variance from option prices. Coined by Britten-Jones and Neuberger [2000], the concept of model-free implied variance arose from the development of variance swaps and appeared in as early as Dupire [1994] and Neuberger [1994]. It is further developed and refined by Carr and Madan [1998], Demeterfi et al. [1999a, 1999b], Britten-Jones and Neuberger [2000], Bakshi, Kapadia and Madan [2003], Carr and Wu [2004], and Jiang and Tian [2005]. We show that the variance measure underlying the new VIX index is theoretically equivalent to the model-free implied variance formulated in Britten-Jones and Neuberger [2000]. This equivalence lends further support for the CBOE's decision to revamp the VIX index.

Specifically, the new VIX is based on the concept of *fair value of future variance* (the DDKZ variance hereafter) developed by Demeterfi et al. [1999a]. Unlike the traditional ex ante volatility measure such as the BSM implied volatility, the DDKZ variance is extracted directly from option prices as follows (see their Equation (26)):

$$V_{ddkz} \equiv \frac{2}{T} \left\{ rT - \left[\frac{S_0}{S_*} \exp(rT) - 1 \right] - \ln(S_*/S_0) \right. \\ \left. + \exp(rT) \int_0^{S_*} \frac{P(T, K)}{K^2} dK \right. \\ \left. + \exp(rT) \int_{S_*}^{\infty} \frac{C(T, K)}{K^2} dK \right\} \quad (1)$$

where S_0 is the current asset price, C and P are call and put prices, respectively, r is the risk-free rate, T and K are option maturity and strike price, respectively, and S_* is an arbitrary stock price (typically chosen to be close to the forward price). In the following proposition, we show that the DDKZ variance is conceptually identical to the model-free implied variance formulated in Britten-Jones and Neuberger [2000]. The proof of this proposition is in the Appendix.

Proposition 1 *The concept of fair value of future variance developed by Demeterfi et al. [1999a] is identical to the model-free implied variance formulated in Britten-Jones and Neuberger [2000], which is defined as:*

$$V_{bjn} \equiv \frac{2 \exp(rT)}{T} \left[\int_0^{F_0} \frac{P(T, K)}{K^2} dK \right. \\ \left. + \int_{F_0}^{\infty} \frac{C(T, K)}{K^2} dk \right] \quad (2)$$

where F_0 is the forward price with maturity coinciding with that of the options.

Proposition 1 is a useful result for two reasons. First, it provides a direct linkage between two variance concepts that have been developed separately for different purposes. It unifies key research results in the study of variance swaps and implied distributions. The theoretical underpinning of the new VIX index is thus rooted in the broader context of model-free implied variance. Secondly, Proposition 1 also implies that prior research findings on the empirical implementation and information content of the model-free implied variance directly apply to that of the DDKZ variance. As demonstrated by Jiang and Tian [2005], the model-free implied variance provides a more efficient forecast for future realized variance and is informationally more efficient than either the BSM implied

variance or historical variance.³ We thus expect the new VIX to be more effective in extracting volatility from option prices than its predecessor.

While the new VIX is conceptually more appealing than its predecessor, the actual construction of the index is also more complex. This is clear from Equations (1) and (2) as the model-free implied variance is defined as an integral of weighted option prices over an infinite range of strike prices. As demonstrated by Carr and Wu [2004] and Jiang and Tian [2005], necessary steps must be taken in order to minimize the implementation errors of the model-free implied variance. However, the actual procedure adopted by the CBOE for the construction of the new VIX does not take these necessary steps and may lead to substantial biases in the calculated index values. Any bias in the calculated index value may have direct economic consequences for variance swaps and other derivatives based on the index. For example, an underestimation (overestimation) of the VIX index leads to a corresponding underpricing (overpricing) of VIX-based variance swaps. If the variance swap rate is naively benchmarked against the VIX index, the bias in the VIX index is translated into pricing errors for the variance swaps leading to potential arbitrage opportunities. In addition, the bias in the VIX index may also lead to a misestimation of the volatility risk premium which measures the difference between realized and risk-neutral variances. As pointed out by Bakshi and Kapadia [2003] and Carr and Wu [2004], a correctly estimated volatility risk premium is important for pricing and hedging derivative securities on underlying assets with stochastic volatility. If the VIX index is used as a proxy for the risk-neutral variance, the errors in the VIX index are subsequently impounded in the volatility risk premium which will lead to biased empirical results.

PROBLEMS IN THE CBOE PROCEDURE

As explained in a white paper (see CBOE [2003]), the CBOE calculates the model-free implied variance as follows:

$$\sigma_{vix}^2 = \frac{2}{T} \sum_i \frac{\Delta K_i}{K_i^2} \exp(rT) Q(T, K_i) - \frac{1}{T} \left(\frac{F_0}{K_0} - 1 \right)^2 \quad (3)$$

where F_0 is the forward index level, T is the option maturity, K_i is the strike price of the i -th (out-of-the-money) option (call if $K_i > F_0$ and put otherwise), K_0 is the first

strike price below F_0 ($K_0 \leq F_0$), $Q(T, K_i)$ is the midpoint of the latest available bid and ask prices for the option, r is the risk-free interest rate, and ΔK_i is the strike price increment calculated as:

$$\Delta K_i = \frac{K_{i+1} - K_{i-1}}{2}$$

At the highest and lowest available strike prices, the increment is modified as the difference between the two highest and lowest strike prices, respectively. In addition, a further adjustment at strike price K_0 is made by redefining $Q(T, K)$ as the average price of call and put options. The reason for this averaging and its impact on approximation errors will be discussed in detail subsequently in this section. Once an estimate of the model-free implied variance is obtained, the VIX index is calculated as its square root, multiplied by 100.

For the purpose of estimating the model-free implied variance, the procedure described in Equation (3) may lead to several types of approximation errors including truncation, discretization, expansion and interpolation errors. Some of these errors (e.g., expansion errors) are unlikely to be economically significant. It is nevertheless important to understand their potential impact on the estimated variance (and hence the VIX index value). We next examine each of these approximation errors in detail and then use numerical examples to illustrate their direction, magnitude and economic significance.

Sources of the Approximation Errors

We begin with truncation errors induced by the limited availability of strike prices. The model-free implied variance requires an infinite range of strike prices in its calculation as its definition in either Equation (1) or (2) indicates. Let K_L and K_U be the lowest and highest strike prices listed for a given maturity, respectively. The CBOE procedure introduces a truncation error as an infinite range of strike prices replaced with a finite range $[K_L, K_U]$:

$$\int_0^{K_0} \frac{P(T, K)}{K^2} dK + \int_{K_0}^{\infty} \frac{C(T, K)}{K^2} dK \approx \int_{K_L}^{K_0} \frac{P(T, K)}{K^2} dK + \int_{K_0}^{K_U} \frac{C(T, K)}{K^2} dK$$

The size of the truncation error is given by:

$$\delta_{\text{trunc}} = -\frac{2}{T} \exp(rT) \left[\int_0^{K_L} \frac{P(T, K)}{K^2} dK + \int_{K_U}^{\infty} \frac{C(T, K)}{K^2} dK \right] \quad (4)$$

The negative sign indicates that the truncation error leads to a downward bias in the calculated variance.

Note that the truncation error may vary substantially over time because the truncation interval $[K_L, K_U]$ is not fixed and can change over time. One reason is that the CBOE usually adds new strike prices when the underlying index moves beyond the range of existing strike prices. The added strike prices expand the truncation interval. During periods of rapid market movement, the expansion of the truncation interval can be both frequent and sizeable. Another reason is that the CBOE applies filters to clean up potentially problematic options. Any option with a zero bid price is considered mispriced and the associated strike price is excluded in the VIX calculation. Such options are deemed to have extreme strike prices and are likely to be associated with poor liquidity. In addition, zero bid prices may cluster near the two end points (K_L and K_U) of the truncation interval. If zero bid prices are found for two consecutive strike prices, any option with a strike price outside the two identified strike prices is also deemed problematic and excluded in the VIX calculation.

Another type of approximation error is discretization error due to numerical integration. The CBOE procedure applies a rather unusual numerical integration scheme and calculates the integrals in Equation (1) as follows:

$$\int_{K_L}^{K_0} \frac{P(T, K)}{K^2} dK + \int_{K_0}^{K_U} \frac{C(T, K)}{K^2} dK \approx \sum_i \frac{\Delta K_i}{K_i^2} Q(T, K_i)$$

Although numerical integration errors can be minimized by using a sufficiently fine partition of strike prices, the actual partition used in the CBOE procedure is based on

listed strike prices and typically is quite coarse. The size of the discretization error is:

$$\delta_{disc} = \frac{2}{T} \exp(rT) \left\{ \sum_i \frac{\Delta K_i}{K_i^2} Q(T, K_i) - \left[\int_{K_L}^{K_0} \frac{P(T, K)}{K^2} dK + \int_{K_0}^{K_U} \frac{C(T, K)}{K^2} dK \right] \right\}$$

The third type of approximation error is due to the Taylor series expansion of the log function used in the CBOE procedure. The terms preceding the integrals in Equation (1) can be restated as:

$$\begin{aligned} \frac{2}{T} \left\{ rT - \left[\frac{S_0}{K_0} \exp(rT) - 1 \right] - \ln(K_0/S_0) \right\} \\ = \frac{2}{T} \left[\ln(F_0/K_0) - \left(\frac{F_0}{K_0} - 1 \right) \right] \end{aligned}$$

Applying the Taylor series expansion of the log function (ignoring terms higher than the second order):

$$\ln(F_0/K_0) \approx \left(\frac{F_0}{K_0} - 1 \right) - \frac{1}{2} \left(\frac{F_0}{K_0} - 1 \right)^2 \quad (5)$$

we have the following approximation:

$$\begin{aligned} \frac{2}{T} \left\{ rT - \left[\frac{S_0}{K_0} \exp(rT) - 1 \right] - \ln(K_0/S_0) \right\} \\ \approx -\frac{1}{T} \left(\frac{F_0}{K_0} - 1 \right)^2 \end{aligned}$$

The right-hand side of the above equation is precisely the second term in Equation (3) from the CBOE procedure. The size of this expansion error is

$$\delta_{exp} = \frac{2}{T} \left\{ \left[\left(\frac{F_0}{K_0} - 1 \right) - \frac{1}{2} \left(\frac{F_0}{K_0} - 1 \right)^2 \right] - \ln(F_0/K_0) \right\} \quad (6)$$

Since $K_0 \leq F_0$ by definition, the remainder term of the Taylor series expansion in Equation (5) is positive. This means that the Taylor series expansion error in Equation (6) is negative. Of course, the size of this error is in the order of $(\frac{F_0}{K_0} - 1)^3$ and likely negligible as K_0 is chosen to be the closest strike price to F_0 (but not exceeding F_0).

The final type of approximation error is due to maturity interpolation used in the CBOE procedure. While the VIX index is based on the model-free implied variance with a fixed 30-day maturity, there are generally no options that expire exactly in 30 calendar days. The solution is to find two maturities, T_1 and T_2 , that are closest to the required 30-day maturity. The variance measures are first calculated for the two selected maturities. The model-free implied variance with the required 30-day maturity is then linearly interpolated between the corresponding variances at the two selected maturities using the formula:

$$\begin{aligned} \hat{\sigma}_{vix}^2(T_0) = \frac{1}{T_0} \left[\omega T_1 \sigma_{vix}^2(T_1) \right. \\ \left. + (1 - \omega) T_2 \sigma_{vix}^2(T_2) \right] \end{aligned} \quad (7)$$

where

$$\omega = \frac{T_2 - T_0}{T_2 - T_1}$$

and T_0 is the required maturity of 30 days. The linear interpolation may introduce an approximation error if the model-free implied variance is a nonlinear function of maturity:

$$\delta_{int} = \hat{\sigma}_{vix}^2(T_0) - \sigma_{vix}^2(T_0) \quad (8)$$

where $\sigma_{vix}^2(T_0)$ is the true model-free implied variance with the required 30-day maturity. As documented in the literature (e.g., Xu and Taylor [1994]), the implied variance term structure is neither linear nor monotonic in option maturity. Nevertheless, the maturity interpolation error is more difficult to address empirically as additional contract months are needed in order to capture nonlinearity in the variance term structure. With longer maturities, illiquidity is likely to be a more serious concern and even fewer strike prices are usually listed. These problems are likely to negate any benefit the additional maturity months

might provide. For this reason, we will focus on the first three types of approximation errors hereafter.

How Large Are the Approximation Errors?

We perform Monte Carlo simulations to analyze and illustrate the possible range of approximation errors in the CBOE procedure.⁴ As the true volatility is known in the simulation, the approximation errors in the CBOE procedure can be accurately assessed. We use the simple BSM model to illustrate the approximation errors in the CBOE procedure. If the CBOE procedure does not work well under constant volatility (which is exactly what we find), its performance is likely even worse under more general models (which we subsequently confirm in the next section).

We choose a *base set of parameters* that are consistent with typical market conditions: The initial asset price (S_0) is 100, the volatility of the underlying asset (σ) is 20%, the option maturity (T) is 30 (calendar) days, the available strike prices are between 80 (K_L) and 120 (K_U) in increments of 2.5 (ΔK). We first illustrate the approximation errors for the base case and then show the possible range of errors by varying one or more of the model parameters. Without loss of generality, we assume the risk-free rate is zero and the underlying asset does not pay any dividend.

With the base set of parameters, we first calculate option prices at the specified strike prices using the BSM model and then use Equation (3) to calculate the VIX. Compared to the true volatility of 20%, the CBOE calculation puts the VIX at 20.31(%), overestimating the target volatility by 31 index basis points or by 1.6%. Since one index basis point is worth \$10 per VIX futures contract, the approximation error translates into a \$310 difference per contract.⁵ In this base case, the total approximation error is almost entirely due to discretization errors. The strike price increment is 2.5% of the initial asset price, which is not particularly large but still too coarse for the VIX calculation.

To illustrate the results over a wide range of parameter values, Exhibit 1 presents approximation errors as the option maturity varies between 15 and 45 days, the increment in strike price between 2.5 and 0.5, and the range of strike prices between [95, 105] and [70, 130]. Other parameters are identical to those in the base set. As shown in the exhibit, the total error may vary between +62 and -275 index basis points or between +3.1% and

-13.8% of the true volatility. The truncation errors are generally negative (as expected) and can be as large as 15.0% of the true volatility. In comparison, the discretization errors are typically positive and can be as large as 6.3% of the true volatility. The approximation errors also tend to increase with option maturity.

The rather large variations in approximation errors illustrated in Exhibit 1 indicate the need for further examination of these errors. To provide insight on the determinants of approximation errors, we plot truncation, discretization, expansion, and total errors in Exhibits 2-4. All approximation errors are stated as a percentage of the true volatility. In each figure, we illustrate the effect of changes in one parameter while keeping other parameters fixed.

Exhibit 2 illustrates approximation errors as the truncation interval ($K_U - K_L$) increases. We vary the length of the truncation interval from 10 to 40 by moving the truncation points (K_U and K_L) away from the initial asset price (S_0). To focus on truncation errors, we use a relatively small strike price increment of $\Delta K = 0.5$ (i.e., 0.5% of the initial asset price). All other parameters are identical to those in the base set. As expected, the total approximation error declines as the truncation interval increases. It is clear that the truncation error accounts for most of the total error in this case. Consistent with the results reported in Exhibit 1, truncation errors lead to an underestimation of the true volatility. As shown in Jiang and Tian [2005], truncation errors are negligible if the truncation points are at least three standard deviations from the initial asset price. Applying to the case illustrated in Exhibit 2, this means that the minimum truncation interval should be [84.1, 118.6]. Indeed, the truncation error is only -0.04% of the true volatility when the truncation interval is [84, 116].

Exhibit 3 illustrates approximation errors as the volatility of the underlying asset varies from 10% to 45%. The strike price increment is fixed at 0.5 in order to control for discretization errors. All other parameters are identical to those in the base set. The total approximation error rises drastically as volatility increases from 20% to 45%. Nearly all of the increase is due to the rise in truncation errors, which leads to an underestimation of the true volatility. This is expected since we use a fixed truncation interval of [80, 120]. As the volatility increases, a larger probability mass is pushed to the tails beyond the truncation interval. Three standard deviations from the initial asset price cover an interval of [84.1, 118.6] when the

EXHIBIT 1

Errors in the CBOE Calculation of the Model-Free Implied Volatility


Option prices are simulated from the BSM model using the following parameters: The initial asset price (S_0) is 100, the volatility of the underlying asset (σ) is 20%, option maturity (T) is either 15, 30 or 45 (calendar) days, the highest available strike price (K_U) varies between 105 and 130, the lowest available strike price (K_L) varies between 70 and 95, and the strike price increment (ΔK) varies between 0.5 and 2.5. The risk-free rate and dividend yield are both zero.

Maturity (days)	ΔK	K_U	K_L	VIX	Total Error		Type of Errors		
					(points)	(%)	Truncation (%)	Discretization (%)	Expansion (%)
15	2.5	105	95	20.26	0.2597	1.30	-4.51	5.81	0.00
15	2.5	110	90	20.62	0.6157	3.08	-0.14	3.22	0.00
15	2.5	120	80	20.62	0.6238	3.12	0.00	3.12	0.00
15	2.5	130	70	20.62	0.6238	3.12	0.00	3.12	0.00
15	1.0	105	95	19.45	-0.5516	-2.76	-4.51	1.75	0.00
15	1.0	110	90	20.08	0.0832	0.42	-0.14	0.55	0.00
15	1.0	120	80	20.10	0.1011	0.51	0.00	0.50	0.00
15	1.0	130	70	20.10	0.1011	0.51	0.00	0.50	0.00
15	0.5	105	95	19.25	-0.7465	-3.73	-4.51	0.78	0.00
15	0.5	110	90	20.00	0.0028	0.01	-0.14	0.15	0.00
15	0.5	120	80	20.03	0.0253	0.13	0.00	0.13	0.00
15	0.5	130	70	20.03	0.0253	0.13	0.00	0.13	0.00
30	2.5	105	95	19.11	-0.8879	-4.44	-10.48	6.04	0.00
30	2.5	110	90	20.19	0.1937	0.97	-1.20	2.17	0.00
30	2.5	120	80	20.31	0.3139	1.57	0.00	1.57	0.00
30	2.5	130	70	20.31	0.3143	1.57	0.00	1.57	0.00
30	1.0	105	95	18.35	-1.6544	-8.27	-10.48	2.21	0.00
30	1.0	110	90	19.86	-0.1359	-0.68	-1.20	0.52	0.00
30	1.0	120	80	20.05	0.0499	0.25	0.00	0.25	0.00
30	1.0	130	70	20.05	0.0506	0.25	0.00	0.25	0.00
30	0.5	105	95	18.12	-1.8825	-9.41	-10.48	1.07	0.00
30	0.5	110	90	19.80	-0.1998	-1.00	-1.20	0.20	0.00
30	0.5	120	80	20.01	0.0118	0.06	0.00	0.06	0.00
30	0.5	130	70	20.01	0.0127	0.06	0.00	0.06	0.00
45	2.5	105	95	18.27	-1.7346	-8.67	-14.96	6.29	0.00
45	2.5	110	90	19.87	-0.1268	-0.63	-2.80	2.16	0.00
45	2.5	120	80	20.21	0.2053	1.03	-0.04	1.07	0.00
45	2.5	130	70	20.21	0.2100	1.05	0.00	1.05	0.00
45	1.0	105	95	17.49	-2.5085	-12.54	-14.96	2.42	0.00
45	1.0	110	90	19.57	-0.4288	-2.14	-2.80	0.65	0.00
45	1.0	120	80	20.03	0.0266	0.13	-0.04	0.18	0.00
45	1.0	130	70	20.03	0.0337	0.17	0.00	0.17	0.00
45	0.5	105	95	17.25	-2.7537	-13.77	-14.96	1.19	0.00
45	0.5	110	90	19.50	-0.5012	-2.51	-2.80	0.29	0.00
45	0.5	120	80	20.00	0.0004	0.00	-0.04	0.05	0.00
45	0.5	130	70	20.01	0.0084	0.04	0.00	0.04	0.00

EXHIBIT 2

Errors in the CBOE Calculation as the Truncation Interval Changes

Option prices are simulated from the BSM model using the following parameters: The initial asset price (S_0) is 100, the volatility of the underlying asset (σ) is 20%, option maturity (T) is 30 (calendar) days, the strike price increment (ΔK) is 0.5, and the truncation interval $[K_L, K_U]$ increases from [95, 105] to [80, 120]. The risk-free rate and dividend yield are both zero.




volatility is 20% but a much wider interval of [67.3, 146.0] when it is 45%. The former is well within the truncation interval used while the latter is not. This means that the approximation error is likely to rise sharply if the market volatility experiences a sudden surge, precisely when a more accurate measure of the VIX is warranted. Unless additional strike prices are added to cover a larger range of strike prices, the approximation errors due to a volatility spike are likely substantial.

Exhibit 4 illustrates approximation errors as the strike price increment (ΔK) varies from 2.5 to 0.1. To focus on discretization errors, we set all other parameters identical to those in the base set. As shown in Exhibit 4, the total approximation error declines as the strike price increment shrinks. The truncation and expansion errors are negligible and represent no more than 0.01% of the true volatility. It is also clear that discretization errors lead to an overestimation of the true volatility. The

EXHIBIT 3

Errors in the CBOE Calculation as the Index Volatility Changes

Option prices are simulated from the BSM model using the following parameters: The initial asset price (S_0) is 100, option maturity (T) is 30 (calendar) days, the highest available strike price (K_U) is 120, the lowest available strike price (K_L) is 80, the strike price increment (ΔK) is 0.5, and the volatility of the underlying asset (σ) varies from 10% to 60%. The risk-free rate and dividend yield are both zero.



discretization error is approximately 0.06% of the true volatility when the strike price increment is 0.5. The error is equivalent to approximately one index basis point or \$10 per VIX futures contract. Discretization errors are thus negligible when the strike price increment is 0.5 or smaller. As shown in Exhibit 4, the errors can be much larger, however, if a larger strike price increment is used.

To understand why discretization errors lead to an overestimation of the VIX, we need to take a closer look at the numerical integration scheme adopted in the CBOE

procedure. The VIX requires the computation of the following integral:

$$\int_0^{+\infty} \frac{Q(T, K)}{K^2} dK$$


where

$$Q(T, K) = \begin{cases} P(T, K), & \text{if } K \leq F_0, \\ C(T, K), & \text{otherwise.} \end{cases}$$

EXHIBIT 4

Errors in the CBOE Calculation as the Strike Price Increment Changes

Option prices are simulated from the BSM model using the following parameters: The initial asset price (S_0) is 100, the volatility of the underlying asset (σ) is 20%, option maturity (T) is 30 (calendar) days, the highest available strike price (K_U) is 120, the lowest available strike price (K_L) is 80, and the strike price increment (ΔK) varies between 0.1 and 2.5. The risk-free rate and dividend yield are both zero.



The integrand is illustrated in Panel A of Exhibit 5 using the base set of parameters described previously. The integrand function reaches a peak at the strike price $K = F_0$ (equal to 100 in the example) and declines sharply (and monotonically) on both sides as the strike price either increases or decreases from F_0 . Due to the kink at $K = F_0$, the accuracy and robustness of any numerical integration scheme depend critically on how the kink is handled. Numerical integration schemes may not perform well when a kink occurs in the middle of the integration

interval. A typical solution is to divide the integral into two parts and have them computed separately on the two sides of the kink:

$$\int_0^{+\infty} \frac{Q(T, K)}{K^2} dK = \int_0^{F_0} \frac{P(T, K)}{K^2} dK + \int_{F_0}^{+\infty} \frac{C(T, K)}{K^2} dK$$


A practical problem with this approach is that the forward price F_0 usually does not coincide with any listed

EXHIBIT 5

Numerical Integration Scheme Used in the CBOE Method

Panel A illustrates integrand in the model-free implied variance for option prices simulated from the BSM model using the base set of parameters: The initial asset price (S_0) is 100, the volatility of the underlying asset (σ) is 20%, option maturity (T) is 30 (calendar) days, the highest available strike price (K_U) is 120, the lowest available strike price (K_L) is 80, and the strike price increment (ΔK) is 0.5. The risk-free rate and dividend yield are both zero. Panel B plots the average call and put price over the strike price range between 90 and 110 centered around the forward price F_0 using the same parameters as in Panel A. Panel C plots the numerical integration scheme used in the CBOE method to calculate the model-free implied variance. The area represented by the vertical bars is the variance calculated using the CBOE method while the true variance is the area under the integrand.

Panel A



Panel B





EXHIBIT 5 (Continued)

Panel C



strike price and the option value (or equivalently the integrand) at F_0 has to be estimated. A relatively small error in the estimated option value may lead to a sizeable approximation error in the VIX since the integrand peaks sharply at F_0 . The CBOE procedure implements an *unusual adjustment* by substituting the average price of call and put options at strike price K_0 (the first strike price below F_0) for the missing option price at F_0 :

$$\hat{Q}(T, F_0) = \frac{P(T, K_0) + C(T, K_0)}{2}$$

Panel B of Exhibit 5 illustrates the average call and put price over the strike price range between 90 and 110 centered around the forward price F_0 using the base set of parameters. It is clear that $\hat{Q}(T, F_0)$ tends to overestimate

the option price $Q(T, F_0)$. The greater the gap between K_0 and F_0 , the larger the overestimation. This is true across a wide range of parameter values for the BSM model.⁶

The impact of the unusual numerical integration scheme adopted in the CBOE procedure is illustrated in Panel C of Exhibit 5. The initial asset price (S_0) is set equal to 99 while the strike prices remain centered at 100 and spread in both directions at an increment of 2.5. All other parameters are identical to those in the base set. Since the risk-free rate is assumed to be zero, we have $F_0 = S_0 = 99$ while $K_0 = 97.5$. The integrand is the curve with a kink at strike price $F_0 = 99$ and the area under the integrand represents the value of the integral. The numerical integration scheme in the CBOE procedure approximates the integral by:

$$\sum_i \frac{\Delta K_i}{K_i^2} Q(T, K_i)$$

which is plotted as the area covered by the shaded bars. The overestimation in $\hat{Q}(T, F_0)$ leads to a center bar that is higher than the peak of the integrand at F_0 . It is clear that the area covered by the bars is larger than the area under the integrand. This unusual numerical integration scheme tends to overestimate the true integral, which explains the overestimation errors illustrated in Exhibit 4.

The analysis thus far has focused on truncation intervals that are symmetric to the initial asset price (S_0) which is also equal to the forward price (F_0) due to the assumption of zero interest rate and dividend yield. How do the approximation errors change if the truncation interval is asymmetric? Asymmetric truncation is quite common due to movement in the underlying index and asymmetry in listed strike prices. Over the period from 1996 to 2004, the range of strike prices below the index ($S_0 - K_L$) is on average 2.2 times the range of strike prices above the index ($K_U - S_0$) for the near month contracts of SPX options and the ratio varied from 9.9 to 0.4. This type of mostly one-sided symmetry is mainly due to investor interest in downside risk protection. To analyze the effect of asymmetric truncation, we keep the length of the truncation interval fixed but move the initial asset price away from the center. We recalculate the VIX using the CBOE procedure and report the approximation errors in Exhibit 6. To focus on truncation errors, we minimize the impact of discretization errors by using a relatively small strike price increment of 0.5. We also fix the volatility of

the underlying asset at 20%. As expected, asymmetry can have a large impact on truncation errors because the implied variance derives much of its value in the central region of the strike price distribution (surrounding the initial asset price S_0) and very little from the two tails. Consider the [90, 110] truncation interval, for example. If the initial asset price is located at the center of the truncation interval (i.e., $S_0 = 100$), the truncation error is -1.5% of the true volatility. If the asset price moves up (down) from 100 to 104 (96), the truncation error jumps to -3.9% (-3.1%) representing a 160% (107%) increase in magnitude. The impact of the asymmetry can be larger or smaller depending on the degree of the asymmetry, the length of the truncation interval, and volatility.

Finally, we analyze the approximation errors in the CBOE calculation under a more realistic setting that mimics actual market structure and conditions. All previous examples are constructed by assuming that option prices are available at the required maturity and across all strike prices at a fixed increment between K_L and K_U . Neither is true in the marketplace. Consider the calculation of the VIX on April 1, 2004. The choice of the actual date is mostly inconsequential as we merely glean the contract maturities and strike prices listed on that date. On April 1, 2004, the near and next term contracts (maturing in April and May, respectively) have 15 and 50 days remaining to maturity, respectively. The April contracts have quotes for 44 strike prices while the May contracts have quotes for 23 strike prices.⁷ Due to zero bid prices, we drop three extreme strike prices for the April contracts and two extreme prices for the May contracts. The remaining strike prices range from 775 to 1,275 for the April contracts and from 800 to 1,250 for the May contracts. The S&P 500 index closed at 1,132.17 on that date but fluctuated between 1,126.20 and 1,135.67 during the day. We continue to assume that the BSM model is the correct model and use it to simulate option prices at the listed strike prices for the April and May contracts. Model prices instead of market prices are used in order to assess the approximation errors. To cover a wide range of market conditions, we vary the volatility of the underlying index from 10% to 45% and the index value from 1,126 to 1,144. The range of index volatilities used here is consistent with historical high and low VIX values (45.7% and 9.3%, respectively) calculated by the CBOE during the ten-year period between 1992 and 2002 (see the table on p. 9 of CBOE [2003]). The VIX index closed at 16.65 on April 1, 2004, well within the range of

EXHIBIT 6

Errors in the CBOE Calculation of the Model-Free Implied Volatility under Asymmetric Strike Price Distribution

Option prices are simulated from the BSM model using the following parameters: The volatility of the underlying asset (σ) is 20%, option maturity (T) is 30 (calendar) days, the highest available strike price (K_U) varies between 105 and 130, the lowest available strike price (K_L) varies between 70 and 95, and the strike price increment (ΔK) is 0.5, and the initial asset price (S_0) varies within the truncation interval. The risk-free rate and dividend yield are both zero.

Asset Price				Total Error		Type of Errors		
S_0	K_U	K_L	VIX	(points)	(%)	Truncation (%)	Discretization (%)	Expansion (%)
103	105	95	16.94	-3.0564	-15.28	-16.76	1.48	0.00
102	105	95	17.60	-2.3998	-12.00	-13.24	1.24	0.00
101	105	95	17.98	-2.0170	-10.09	-11.19	1.10	0.00
100	105	95	18.12	-1.8825	-9.41	-10.48	1.07	0.00
99	105	95	18.01	-1.9908	-9.95	-11.08	1.12	0.00
98	105	95	17.64	-2.3570	-11.79	-13.07	1.28	0.00
97	105	95	16.98	-3.0198	-15.10	-16.67	1.57	0.00
106	110	90	18.54	-1.4582	-7.29	-8.00	0.71	0.00
104	110	90	19.29	-0.7053	-3.53	-3.93	0.41	0.00
102	110	90	19.67	-0.3326	-1.66	-1.92	0.25	0.00
100	110	90	19.80	-0.1998	-1.00	-1.20	0.20	0.00
98	110	90	19.75	-0.2548	-1.27	-1.52	0.25	0.00
96	110	90	19.46	-0.5377	-2.69	-3.09	0.41	0.00
94	110	90	18.79	-1.2057	-6.03	-6.77	0.74	0.00
112	120	80	19.59	-0.4086	-2.04	-2.29	0.25	0.00
108	120	80	19.93	-0.0671	-0.34	-0.43	0.10	0.00
104	120	80	20.00	0.0018	0.01	-0.06	0.06	0.00
100	120	80	20.01	0.0118	0.06	0.00	0.06	0.00
96	120	80	20.01	0.0128	0.06	-0.01	0.07	0.00
92	120	80	20.00	0.0017	0.01	-0.08	0.09	0.00
88	120	80	19.88	-0.1201	-0.60	-0.79	0.19	0.00
118	130	70	19.89	-0.1105	-0.55	-0.66	0.10	0.00
112	130	70	20.00	0.0031	0.02	-0.04	0.05	0.00
106	130	70	20.01	0.0111	0.06	0.00	0.06	0.00
100	130	70	20.01	0.0127	0.06	0.00	0.06	0.00
94	130	70	20.01	0.0143	0.07	0.00	0.07	0.00
88	130	70	20.02	0.0163	0.08	0.00	0.08	0.00
82	130	70	20.01	0.0146	0.07	-0.03	0.10	0.00

volatility levels used here. We again assume zero interest rate and dividend yield on the index.

We next calculate the model-free implied variance for the near maturity month (April). Option prices are simulated from the BSM model for all available strike prices. The forward index level is calculated from at-the-money options using the put-call parity:

$$F_0 = K_* + \exp(-rT) [C(K_*, T) - P(K_*, T)]$$

The at-the-money strike price (K_*) is determined as the strike price at which the difference between the call and put prices is the smallest. Given the calculated forward index value, we next mark the first strike price (K_0) below the forward index value. Out-of-the-money

options are selected by choosing call options with strike price greater than K_0 and put options with strike price less than or equal to K_0 . Equation (3) is then used to calculate the model-free implied variance for the April maturity. The same procedure is repeated for the next maturity month (i.e., May) to calculate another model-free implied variance. Finally, the required 30-day model-free implied variance is linearly interpolated from the variances of the near and next maturity months using Equation (7).⁸ The calculated VIX values are summarized in Exhibit 7.

As shown in Exhibit 7, the total approximation error from the CBOE procedure varies from +25 to -197 index basis points or from +2.6% to -4.4% of the true volatility. These errors translate into value differences ranging from +\$250 to -\$1,970 per VIX futures contract and are clearly economically significant. More importantly, the approximation errors exhibit predictable patterns in relations to volatility levels. Although the magnitude of the approximation errors appears to be small when the true volatility is close to 20%, it may increase sharply if the volatility declines or rises from 20%. (Positive) discretization and (negative) truncation errors appear to mostly offset each other at 20% volatility, leading to negligible total errors ranging from 1 to 3 index basis points. As the volatility increases (decreases), the truncation (discretization) error becomes more dominant leading to underestimation (overestimation) of the true volatility. In other words, the CBOE procedure tends to underestimate (overestimate) the true volatility when the underlying volatility is high (low). In addition, similar patterns in approximation errors are also found when the SVJ model is used to simulate option prices (as will be seen). The magnitude of approximation errors is also greater in that case.

A SIMPLE SOLUTION

To fix the problems in the CBOE procedure, we propose a simple *smoothing method* for extracting the model-free implied variance from option prices. It is based on the construction of the implied volatility function using an interpolation-extrapolation scheme. Care is taken to ensure that the constructed implied volatility function is smooth over the entire range of strike prices. Smoothness is an important requirement since it is directly implied by no-arbitrage conditions (e.g., Breeden and Litzenberger [1978]).

The Smoothing Method

Suppose N strike prices are listed for trading for a given maturity T . Market prices for call and put options are $C^M(K_i, T)$ and $P^M(K_i, T)$, respectively, for $i = 1, 2, \dots, N$.⁹ Without loss of generality, let $0 < K_L = K_1 < K_2 < \dots < K_N = K_U < +\infty$. Only out-of-the-money options are used in the construction of the implied volatility function. An option is at the money if its strike price is equal to the forward asset price F_0 . Let K_{i_0} be the first strike price below the forward price:

$$K_{i_0} \leq F_0 < K_{i_0+1}$$

The implied volatilities ($\sigma(K_i, T)$ for $i = 1, 2, \dots, N$) at the N available strike prices are then calculated from the corresponding call or put options using the BSM model.

The implied volatility function over the entire strike price interval $(0, +\infty)$ must be constructed from the N known implied volatilities. Previous research has considered various methods including both parametric and nonparametric ones (e.g., Bahra [1997], Jackwerth [1999], Bliss and Panigirtzoglou [2000], Bondarenko [2000], and Anagnou et al. [2002]). To be consistent with the basic philosophy of the CBOE procedure, the constructed implied volatility function should fit the N known implied volatilities exactly. As most parametric methods are not flexible enough to fit an arbitrary number of observed implied volatilities, a nonparametric approach is adopted. Interpolation is first implemented between listed strike prices to construct a smooth function that exactly fits the N known implied volatilities. Extrapolation is then implemented outside the range of listed strike prices to construct an extension of the implied volatility function in the two tails of the strike price distribution.

In the interpolation step, we wish to find a differentiable function $f(K)$ in the interval $[K_L, K_U]$ such that

$$f(K_i) = \sigma(K_i, T)$$

for $i = 1, 2, \dots, N$. Following prior research (e.g., Bates [1991, 2000], Campa, Chang, and Reider [1998], and Jiang and Tian [2005]), we use cubic splines to fit the known implied volatilities. In order to obtain an exact fit, all interior strike prices (from K_2 to K_{N-1}) are used as knot points in fitting the cubic splines. In other words, we use *natural cubic splines* with a different cubic function fitted in the interval between any two consecutive strike

EXHIBIT 7

Errors in the CBOE Calculation of the VIX Index

The VIX is calculated using the CBOE method on a typical trading day (April 1, 2004). Option prices at listed strike prices are simulated using the Black-Scholes model for both the near and next maturity months. The index volatility (σ) is either 10, 20, 40 or 60% and the index value varies between 1,126 and 1,144. The risk-free rate and dividend yield are both zero.

Index	True Volatility (%)	VIX	Total Error		Type of Errors		
			(points)	(%)	Truncation (%)	Discretization (%)	Expansion (%)
1126	10	10.19	0.1871	1.87	0.00	1.87	0.00
1128	10	10.19	0.1945	1.94	0.00	1.94	0.00
1130	10	10.20	0.2018	2.02	0.00	2.02	0.00
1132	10	10.21	0.2095	2.09	0.00	2.10	0.00
1134	10	10.22	0.2171	2.17	0.00	2.18	-0.01
1136	10	10.22	0.2247	2.25	0.00	2.26	-0.01
1138	10	10.23	0.2323	2.32	0.00	2.34	-0.02
1140	10	10.24	0.2398	2.40	0.00	2.42	-0.02
1142	10	10.25	0.2475	2.47	0.00	2.51	-0.03
1144	10	10.25	0.2549	2.55	0.00	2.60	-0.05
1126	20	20.03	0.0322	0.16	-0.29	0.45	0.00
1128	20	20.03	0.0314	0.16	-0.30	0.46	0.00
1130	20	20.03	0.0303	0.15	-0.33	0.48	0.00
1132	20	20.03	0.0291	0.15	-0.35	0.49	0.00
1134	20	20.03	0.0276	0.14	-0.38	0.51	0.00
1136	20	20.03	0.0258	0.13	-0.40	0.53	0.00
1138	20	20.02	0.0236	0.12	-0.43	0.55	0.00
1140	20	20.02	0.0210	0.11	-0.46	0.57	-0.01
1142	20	20.02	0.0181	0.09	-0.49	0.59	-0.01
1144	20	20.01	0.0147	0.07	-0.52	0.61	-0.01
1126	30	29.65	-0.3467	-1.16	-1.28	0.12	0.00
1128	30	29.64	-0.3616	-1.21	-1.34	0.13	0.00
1130	30	29.62	-0.3772	-1.26	-1.39	0.14	0.00
1132	30	29.61	-0.3933	-1.31	-1.46	0.14	0.00
1134	30	29.59	-0.4103	-1.37	-1.52	0.15	0.00
1136	30	29.57	-0.4280	-1.43	-1.59	0.16	0.00
1138	30	29.55	-0.4465	-1.49	-1.65	0.17	0.00
1140	30	29.53	-0.4659	-1.55	-1.73	0.18	0.00
1142	30	29.51	-0.4861	-1.62	-1.80	0.19	0.00
1144	30	29.49	-0.5073	-1.69	-1.88	0.20	-0.01
1126	45	43.45	-1.5492	-3.44	-3.54	0.10	0.00
1128	45	43.41	-1.5909	-3.54	-3.63	0.10	0.00
1130	45	43.37	-1.6340	-3.63	-3.74	0.11	0.00
1132	45	43.32	-1.6783	-3.73	-3.84	0.11	0.00
1134	45	43.28	-1.7240	-3.83	-3.95	0.12	0.00
1136	45	43.23	-1.7710	-3.94	-4.08	0.14	0.00
1138	45	43.18	-1.8194	-4.04	-4.19	0.15	0.00
1140	45	43.13	-1.8694	-4.15	-4.31	0.16	0.00
1142	45	43.08	-1.9208	-4.27	-4.44	0.17	0.00
1144	45	43.03	-1.9738	-4.39	-4.56	0.18	0.00

prices, and the cubic functions in any two adjacent intervals have identical first- and second-order derivatives at the common knot point. The output of the cubic spline fitting is a smooth implied volatility function $f(K)$ in the interval $[K_L, K_U]$, together with the first- and second-order derivatives $f'(K)$ and $f''(K)$ at every strike price in the interval. For further details on fitting natural cubic splines, see Press et al. [1996], pp. 107–110.

The next step in the construction of the implied volatility function is the extrapolation procedure. As no option value is known beyond the range of listed strike prices, the tail segments of the implied volatility function must be extrapolated from known implied volatilities at listed strike prices. Both Carr and Wu [2004] and Jiang and Tian [2005] implement a flat extrapolation scheme, assuming the implied volatility function is flat beyond listed strike prices. This flat extrapolation scheme has two drawbacks. First, it tends to underestimate implied volatilities beyond listed strike prices due to the commonly observed volatility smile or skew. As implied volatility tends to rise as the strike price moves away from the stock price (more so on the left side than on the right side), flat extrapolation underestimates the implied volatility function at the two tails. Secondly, the flat extrapolation scheme leads to an implied volatility function with kinks at K_L and K_U . This violates no-arbitrage conditions as the kinks in the implied volatility function are associated with negative local risk-neutral density.

In order to avoid such drawbacks, we impose a smooth pasting condition at the minimum (K_L) and maximum (K_U) strike prices but maintain a linear extrapolation structure. This is done by adjusting the slope of the extrapolated segment (on both sides) to match the corresponding slope of the interior segment at the minimum (K_L) or maximum (K_U) strike price. The implied volatility function constructed this way is smooth everywhere and fits the known implied volatilities exactly. This modified extrapolation scheme conforms more closely with the commonly observed volatility smile or skew.

Once the implied volatility function is constructed, option prices at any required strike prices can be translated from the corresponding implied volatilities using the BSM model. To calculate the model-free implied variance, we use the following partition of strike prices:

$$\ln \hat{K}_i = \ln F_0 + i\eta$$

where η is a small positive number and $i = 0, \pm 1, \pm 2, \dots$. Let $C^{EX}(\hat{K}_i, T)$ and $P^{EX}(\hat{K}_i, T)$ be the call and put prices, respectively, calculated from the constructed implied volatility function at strike price $\hat{K}_i$ and maturity T . Using a trapezoidal rule to implement the numerical integration in Equation (2), the model-free implied variance is calculated as:

$$V_{\text{bm}} = \frac{2}{T} \exp(rT) \left\{ \sum_{i \leq 0} \frac{\Delta \hat{K}_i}{2} \left[\frac{P^{EX}(\hat{K}_i, T)}{\hat{K}_i^2} + \frac{P^{EX}(\hat{K}_{i-1}, T)}{\hat{K}_{i-1}^2} \right] + \sum_{i > 0} \frac{\Delta \hat{K}_i}{2} \left[\frac{C^{EX}(\hat{K}_i, T)}{\hat{K}_i^2} + \frac{C^{EX}(\hat{K}_{i-1}, T)}{\hat{K}_{i-1}^2} \right] \right\} \quad (9)$$

where

$$\Delta \hat{K}_i = \hat{K}_i - \hat{K}_{i-1}$$

We can minimize the numerical integration error by choosing a sufficiently small η .

Accuracy and Robustness of the Smoothing Method

We now illustrate the accuracy and robustness of the smoothing method and compare its performance with that of the CBOE method. This is first implemented using simulated option prices. We then apply the two methods to daily SPX options data and see how the calculated VIX index values differ under actual market conditions.

We begin with a simulation experiment using the SVJ model:

$$\begin{aligned} dS_t/S_t &= V_t^{1/2} dW_t + J_t dN_t - \mu_j \lambda dt, \\ dV_t &= (\theta_v - \kappa_v V_t) dt + \sigma_v V_t^{1/2} dW_t^v, \\ dW_t dW_t^v &= \rho dt \end{aligned} \quad (10)$$

where V_t is the instantaneous variance, θ_v , κ_v , σ_v , λ , μ_j and σ_j are model parameters, ρ is the correlation coefficient between the two diffusion processes, and $N_t \sim \text{iid Poisson}(\lambda)$ and $\ln(1 + J_t) \sim \text{iid N}[\ln(1 + \mu_j) - \frac{1}{2}\sigma_j^2, \sigma_j]$ specify the jump component of the asset price process. We

again make the convenient assumption that both interest rate and dividend yield on the underlying asset are zero. To cover the wide range of volatility levels experienced in the historical SPX index returns, we calibrate the SVJ model such that the annualized integrated volatility (σ_1) vary from 10% to 45%. Other model parameters are chosen to reflect typical estimates from previous empirical studies (e.g., Bakshi, Cao, and Chen [1997]): $\kappa_v = 1$, $\sigma_v = 0.25$, $\rho = 0$, $\lambda = 0.5$, $\mu_j = -0.15\sigma_j$, $\sigma_j = 0.15\sigma_v$, $\theta_v = 0.85\kappa_v\sigma_v^2$, and $V_0 = \theta_v/\kappa_v$. Note that the choice of initial volatility level (V_0) also ensures a flat term structure of the model-free implied variance. The maturity interpolation error is thus absent in our simulation experiment here. The jump parameters (λ , μ_j and σ_j) are set such that the jump volatility ($\sqrt{\lambda(\mu_j^2 + \sigma_j^2)}$) accounts for 15% of the total volatility. Exhibit 8 illustrates the volatility smile for one-month and six-month options using these parameters with total volatility varying from 10% to 45%. It is clear that a variety of volatility smile patterns are represented in these plots.

We then use the above SVJ model to simulate option prices across all strike prices listed on April 1, 2004. As discussed in the previous section, the April and May contracts have 15 and 50 days remaining to maturity on that date, respectively. For each of the two contract months, we calculate the model-free implied variance first using the smoothing method and then the CBOE method. The model-free implied variance with the required 30-day maturity is linearly interpolated from the two variances from April and May contracts for each method. The results are summarized in Exhibit 9.

As shown in Exhibit 9, the VIX index calculated using the smoothing method is consistently accurate across all volatility and index levels. The maximum error is only 8 index basis points while most errors are within 5 index basis points of the true volatility. In comparison, the VIX index calculated using the CBOE procedure exhibits errors ranging from +79 to -198 index basis points. As expected, the approximation errors from the CBOE procedure are larger under the SVJ model (Exhibit 9) than under the BSM model (Exhibit 7). More importantly, the CBOE procedure continues to exhibit the same predictable patterns of errors first reported in Exhibit 7. In other words, the CBOE procedure leads to positive discretization errors and negative truncation errors. At the modest level of 20% volatility, the two errors appear to mostly offset each other, leading to a small total error in the CBOE method (ranging from 0 to 2 index basis

points). At lower volatility levels (e.g., 10%), however, the discretization error dominates and the CBOE method overestimates the true volatility. At higher volatility levels (e.g., 30%), the truncation error dominates and the CBOE method underestimates the true volatility.

Finally, we compare the two methods using daily data of the SPX options from January 1996 to May 2004. Daily closing prices of SPX options are obtained from the market data service of Commodity Systems, Inc. Daily one-month and three-month Treasury-bill yields (our risk-free rates) are obtained from the Federal Reserve Bulletin. When the option maturity is shorter (longer) than either of the two Treasury-bill maturities, the one-month (three-month) yield is used as the risk-free rate. For other option maturities, the risk-free rate is linearly interpolated between the two yields.


To select the final sample of SPX options, we apply several commonly used data filters. First, we exclude options with less than seven days to maturity because of liquidity and market microstructure concerns. This filter is adopted in the CBOE procedure and in most prior research. Second, in-the-money options are excluded from the sample. In-the-money options are more expensive and generally less liquid than at-the-money or out-of-the-money options (see Ait-Sahalia and Lo [1998]). Following the CBOE procedure, we use the put-call parity to imply the forward index value from option prices. Call options with strike prices below or equal to the implied forward index value are excluded while put options with strike prices above the implied forward index value are excluded. Finally, we exclude a small subset of options (approximately 8% of the sample) that appear to be mispriced. These are typically options that are extremely out of the money and have BSM implied volatilities below zero or above 100%.¹⁰

Exhibit 10 illustrates the differences in the two daily VIX time series calculated using the CBOE procedure and the smoothing method. As shown in Exhibit 10, the differences in the two series are nearly all negative indicating underestimation by the CBOE procedure. This is consistent with the theoretical analysis in the previous section and the simulation results presented previously in Exhibit 9. In addition, Exhibit 11 provides detailed summary statistics of the two VIX series and those of the discretization, truncation and total errors from the CBOE calculation. As the true value of the model-free implied variance is unknown, we assess the errors of the CBOE calculation relative to the smoothing method.

EXHIBIT 8

The Volatility Smile of Option Prices Simulated from the Stochastic Volatility with Random Jump (SVJ) Model

The volatility smile is illustrated for options with one-month and six-month maturities for four sets of model parameters. The following SVJ model is used to simulated option prices: $dS_t/S_t = V_t^{1/2}dW_t + J_t dN_t - \mu_f \lambda dt$, $dV_t = (\theta_v - \kappa_v V_t)dt + \sigma_v V_t^{1/2}dW_t^v$, and $dW_t dW_t^v = \rho dt$ where S_t is the price of an underlying asset, V_t is the instantaneous variance, θ_v , κ_v , σ_v , λ , μ_f and σ_f are non-negative constant, ρ is the correlation coefficient between the two diffusion processes, and $N_t \sim \text{iid Position}(\lambda)$ and $\ln(1 + J_t) \sim \text{iid } N[\ln(1 + \mu_f) - \frac{1}{2}\sigma_f^2, \sigma_f]$ specify the jump component of the asset price process. The annualized integrated volatility (σ_f) varies from 10 to 60% across the four panels. Other model parameters are chosen to reflect typical estimates from empirical studies: $\kappa_v = 1$, $\sigma_v = 0.25$, $\rho = 0$, $\lambda = 0.5$, $\mu_f = -0.15\sigma_f$, $\sigma_f = 0.15\sigma_f$, $\theta_v = 0.85 \kappa_v \sigma_f^2$, and $V_0 = \theta_v/\kappa_v$. The jump parameters (λ , μ_f and σ_f) are chosen such that the jump volatility ($\sqrt{\lambda(\mu_f^2 + \sigma_f^2)}$) accounts for 15% of the total volatility.



As shown in Exhibit 11, the mean and median VIX index value calculated from the CBOE procedure are 22.85 and 21.64, respectively. The corresponding numbers from the smoothing method are 23.90 and 22.70, respectively. These statistics indicate the CBOE method on

average underestimates the VIX by approximately 105 or 106 index basis points. Such magnitude of underestimation is well within the range of possible errors found in our Monte Carlo simulation (Exhibit 9). More importantly, the underestimation of the CBOE method is persistent

EXHIBIT 9

Comparison of the CBOE Method and the Smoothing Method


The VIX is calculated using the CBOE method and the smoothing method on a typical trading day (April 1, 2004). Option prices at listed strike prices are simulated using the SVJ model for both the near and next maturity months. The index volatility (σ) is either 10, 20, 40 or 60% and the index value varies between 1,126 and 1,144. The risk-free rate and dividend yield are both zero.

Index	True Vol (%)	The Smoothing Method			The CBOE Method					
		VIX	Total Error		VIX	Total Error		Type of Errors		
			(points)	(%)		(points)	(%)	Trunca- tion (%)	Discreti- zation (%)	Expan- sion (%)
1126	10	10.05	0.0535	0.53	10.16	0.1616	1.62	-0.02	1.63	0.00
1128	10	10.05	0.0510	0.51	10.32	0.3173	3.17	-0.01	3.19	0.00
1130	10	10.05	0.0499	0.50	10.46	0.4629	4.63	-0.01	4.64	0.00
1132	10	10.05	0.0498	0.50	10.59	0.5898	5.90	-0.01	5.91	0.00
1134	10	10.06	0.0602	0.60	10.69	0.6902	6.90	-0.06	6.97	-0.01
1136	10	10.05	0.0486	0.49	10.76	0.7593	7.59	0.00	7.60	-0.01
1138	10	10.06	0.0624	0.62	10.79	0.7941	7.94	-0.05	8.01	-0.02
1140	10	10.07	0.0719	0.72	10.79	0.7940	7.94	-0.10	8.06	-0.02
1142	10	10.08	0.0766	0.77	10.76	0.7608	7.61	-0.11	7.76	-0.03
1144	10	10.08	0.0772	0.77	10.70	0.6978	6.98	-0.11	7.14	-0.05
1126	20	20.02	0.0229	0.11	20.02	0.0212	0.11	-0.34	0.45	0.00
1128	20	20.02	0.0232	0.12	20.02	0.0202	0.10	-0.36	0.47	0.00
1130	20	20.02	0.0235	0.12	20.02	0.0188	0.09	-0.39	0.48	0.00
1132	20	20.02	0.0236	0.12	20.02	0.0174	0.09	-0.41	0.50	0.00
1134	20	20.03	0.0257	0.13	20.02	0.0156	0.08	-0.44	0.52	0.00
1136	20	20.03	0.0261	0.13	20.01	0.0135	0.07	-0.46	0.53	0.00
1138	20	20.03	0.0264	0.13	20.01	0.0111	0.06	-0.49	0.55	0.00
1140	20	20.03	0.0267	0.13	20.01	0.0082	0.04	-0.52	0.57	-0.01
1142	20	20.03	0.0270	0.14	20.01	0.0050	0.03	-0.55	0.59	-0.01
1144	20	20.03	0.0273	0.14	20.00	0.0014	0.01	-0.59	0.61	-0.01
1126	30	30.01	0.0117	0.04	29.64	-0.3567	-1.19	-1.32	0.13	0.00
1128	30	30.01	0.0117	0.04	29.63	-0.3716	-1.24	-1.37	0.13	0.00
1130	30	30.01	0.0118	0.04	29.61	-0.3872	-1.29	-1.43	0.14	0.00
1132	30	30.01	0.0119	0.04	29.60	-0.4035	-1.34	-1.49	0.15	0.00
1134	30	30.01	0.0119	0.04	29.58	-0.4204	-1.40	-1.55	0.15	0.00
1136	30	30.01	0.0118	0.04	29.56	-0.4382	-1.46	-1.62	0.16	0.00
1138	30	30.01	0.0113	0.04	29.54	-0.4567	-1.52	-1.69	0.17	0.00
1140	30	30.01	0.0112	0.04	29.52	-0.4761	-1.59	-1.76	0.18	0.00
1142	30	30.01	0.0112	0.04	29.50	-0.4963	-1.65	-1.84	0.19	0.00
1144	30	30.01	0.0113	0.04	29.48	-0.5175	-1.72	-1.91	0.20	-0.01
1126	45	45.00	-0.0048	-0.01	43.44	-1.5561	-3.46	-3.56	0.10	0.00
1128	45	45.00	-0.0047	-0.01	43.40	-1.5977	-3.55	-3.66	0.11	0.00
1130	45	44.99	-0.0055	-0.01	43.36	-1.6406	-3.65	-3.76	0.12	0.00
1132	45	44.99	-0.0054	-0.01	43.32	-1.6847	-3.74	-3.87	0.12	0.00
1134	45	44.99	-0.0053	-0.01	43.27	-1.7302	-3.84	-3.98	0.13	0.00
1136	45	45.00	-0.0038	-0.01	43.22	-1.7771	-3.95	-4.09	0.14	0.00
1138	45	45.00	-0.0037	-0.01	43.17	-1.8254	-4.06	-4.20	0.15	0.00
1140	45	45.00	-0.0036	-0.01	43.12	-1.8752	-4.17	-4.32	0.16	0.00
1142	45	45.00	-0.0035	-0.01	43.07	-1.9264	-4.28	-4.44	0.16	0.00
1144	45	45.00	-0.0033	-0.01	43.02	-1.9792	-4.40	-4.57	0.17	0.00

EXHIBIT 10

Time Series Plot of Differences in the Daily Closing VIX Calculated Using the CBOE Method and the Smoothing Method

The VIX is calculated using the CBOE method and the smoothing method using daily closing prices of SPX options from January 1996 to May 2004. On each trading day in the sample period, the two nearest contract months are identified and strike prices are selected following the CBOE procedure.



and occurs in more than 90% of the sample as indicated by the percentiles reported for the total errors in the CBOE calculation.

Another interesting finding is that the CBOE calculation leads to positive discretization errors and negative truncation errors, consistent with our analysis and simulation results in the previous section. In particular, the discretization error is positive in more than 95% of the cases over the sample period with an average value of +24 index basis points. In comparison, the truncation error is all negative over the sample period with an average error

of -129 index basis points. The size of the truncation error clearly dominates the size of the discretization error in most cases, leading to an overall underestimation of the VIX by the CBOE calculation. Both types of errors appear to be within the range of errors reported in our simulation experiment (Exhibit 9).

CONCLUSION

The VIX index is designed to capture the market's aggregate expectation of future volatility in real time. It

EXHIBIT 11

Empirical Comparison of the CBOE Method and the Smoothing Method

The VIX is calculated using the CBOE method and the smoothing method using daily closing prices of SPX options from January 1996 to May 2004. On each trading day in the sample period, the two nearest contract months are identified and strike prices are selected following the CBOE procedure.

	VIX (Smoothing)	VIX (CBOE)	Errors in the CBOE Method (Index Points)		
			Truncation	Discretization	Total
N	2117	2117	2117	2117	2117
Mean	23.90	22.85	-1.29	0.24	-1.05
St. Dev	5.77	5.71	1.10	0.20	1.08
Skewness	1.03	1.05	-2.51	2.60	-2.29
Kurtosis	1.27	1.14	11.24	25.41	9.53
99%	41.63	40.38	-0.11	0.97	0.18
95%	35.52	34.41	-0.18	0.55	0.03
90%	31.89	30.93	-0.26	0.44	-0.04
Q3	26.68	25.57	-0.51	0.31	-0.27
Median	22.70	21.64	-1.05	0.20	-0.81
Q1	20.12	19.06	-1.73	0.14	-1.52
10%	17.49	16.54	-2.53	0.09	-2.32
5%	16.16	15.59	-3.10	0.05	-2.88
1%	14.54	13.61	-5.37	-0.08	-4.96

has been closely followed by market participants and the financial media. Futures contracts on the VIX began trading on the CBOE in early 2004 while option contracts followed one year later, adding further importance to the widely followed volatility index. Given its importance and economic significance, it is essential that the VIX index is rooted on solid theoretical ground and its calculation is accurate and robust.

After analyzing the CBOE procedure for constructing the VIX index, we find that it can lead to substantial biases in the calculated index value. Both truncation and discretization errors can be economically significant and have predictable patterns. Truncation errors lead to an underestimation of the true volatility as tails of the volatility smile are ignored while discretization errors lead to an overestimation due to the use of an unusual numerical integration scheme. Simulation experiments and empirical estimation indicate that truncation errors typically dominate discretization errors and the CBOE

procedure tends to underestimate the model-free implied variance.

To fix the problems in the CBOE procedure, we propose a simple *smoothing method* for extracting the model-free implied variance from option prices. Based on an interpolation-extrapolation scheme, our method has the advantage that the constructed implied volatility function is smooth and fits the known implied volatilities exactly. Simulation results indicate that the smoothing method is accurate and robust across a wide range of model specifications and market conditions. Empirical estimates using daily SPX option prices in the period from January 1996 to May 2004 provide further support for the smoothing method. As the market for VIX futures, options and variance swaps is growing rapidly, a reliable VIX index is critically important in maintaining investor confidence, market stability and growth. Our research should help the CBOE to fine-tune its important volatility index.

APPENDIX

In this appendix, we provide a proof for Proposition 1. As shown in Jiang and Tian [2005], the model-free implied variance formulated in Britten-Jones and Neuberger [2000] can be written as:

$$V_{bjn} = \frac{2}{T} \int_0^{\infty} \frac{\exp(rT) C(T, K) - \max[0, S_0 \exp(rT) - K]}{K^2} dK$$

Partitioning the integral into two segments at $F_0 = S_0 \exp(rT)$, we have:

$$V_{bjn} = \frac{2 \exp(rT)}{T} \left[\int_0^{F_0} \frac{C(T, K) - S_0 + K \exp(-rT)}{K^2} dK + \int_{F_0}^{\infty} \frac{C(T, K)}{K^2} dK \right]$$

Using the put-call parity, a more compact expression for the model-free implied variance follows:

$$V_{bjn} = \frac{2 \exp(rT)}{T} \left[\int_0^{F_0} \frac{P(T, K)}{K^2} dK + \int_{F_0}^{\infty} \frac{C(T, K)}{K^2} dK \right]$$

Rewrite the above equation as:

$$V_{bjn} = \frac{2 \exp(rT)}{T} \left[\int_0^{S_*} \frac{P(T, K)}{K^2} dK + \int_{S_*}^{\infty} \frac{C(T, K)}{K^2} dK + \int_{S_*}^{F_0} \frac{P(T, K) - C(T, K)}{K^2} dK \right]$$

Using the put-call parity once more, we have:

$$V_{bjn} = \frac{2 \exp(rT)}{T} \left[\int_0^{S_*} \frac{P(T, K)}{K^2} dK + \int_{S_*}^{\infty} \frac{C(T, K)}{K^2} dK + \int_{S_*}^{F_0} \frac{K \exp(-rT) - S_0}{K^2} dK \right]$$

Integrating out the third term inside the brackets, we have an equivalent definition for the model-free implied variance:

$$V_{bjn} = \frac{2}{T} \left\{ rT - \left[\frac{S_0}{S_*} \exp(rT) - 1 \right] - \ln(S_*/S_0) + \exp(rT) \int_0^{S_*} \frac{P(T, K)}{K^2} dK + \exp(rT) \int_{S_*}^{\infty} \frac{C(T, K)}{K^2} dK \right\}$$

This alternative expression for the model-free implied variance is exactly the same as the DDKZ variance in Equation (1), which establishes their theoretical equivalence.

ENDNOTES

The authors thank Torben Andersen, Peter Carr, Stephen Figlewski (the editor and referee), Moshe Milevsky, Catherine Shalen of the CBOE, and participants in the 2005 *Northern Finance Association* annual meeting for helpful comments and suggestions. This article was previously circulated under the title "Gauging the 'investor fear gauge': Implementation problems in the CBOE's new volatility index and a simple solution." Tian acknowledges the financial support of the *Social Sciences and Humanity Research Council of Canada*.

¹The new VIX retains the original name (i.e., the VIX) while the original VIX is now known as the VXO. Both volatility indices are currently calculated and disseminated in real time to market participants by the CBOE.

²Numerous empirical studies demonstrate that the BSM model has well known pricing biases (e.g., Black and Scholes [1972], Galai [1977], MacBeth and Merville [1979], Rubinstein [1985], Chance [1986], Shastri and Tandon [1986], Lauterbach and Schultz [1990], and Bakshi, Cao, and Chen [1997]).

³For a survey of the broader literature on volatility forecasting and informational efficiency, see Figlewski [1997], Poon and Granger [2003], and Andersen, Bollerslev, Christoffersen and Diebold [2005].

⁴What we do here is not really Monte Carlo simulation. A specific model (e.g., the BSM model) is used to calculate option prices which are then used to demonstrate approximation errors. For convenience, we use/abuse the term (Monte Carlo) simulation in all such cases as it clearly indicates the theoretical nature of the results.

⁵The contract size of the VIX futures is \$1,000 $\times$ VIX.

⁶Unreported results indicate that the average price of call and put options exhibits similar patterns across a wide range of parameter values for the SVJ model as well.

⁷The 44 listed strike prices for the April contracts are 700 to 975 at a 25-point increment, 995, 1,005, 1,025 to 1,040 at a 5-point increment, 1,050 to 1,070 at a 10-point increment, 1,075, 1,080 to 1,100 at a 10-point increment, 1,105 to 1,180 at a 5-point increment, 1,190, 1,200, and 1,225 to 1,275 at a 25-point increment. In comparison, the 23 listed strike prices for the May contracts are 700 to 800 at a 50-point increment, 825 to 975 at a 25-point increment, 995, 1,005, 1,025 to 1,125 at a 25-point increment, 1,145, 1,150, 1,190, and 1,200 to 1,250 at a 25-point increment.

⁸Note that the maturity interpolation used here does not introduce any approximation error since the volatility term structure is flat in the BSM model. The true model-free implied volatility is identical for both maturities.

⁹We implicitly assume that the market prices of call and put options are internally consistent and do not admit any arbitrage opportunities. In a market with frictions such as transaction costs, option prices may violate theoretical no-arbitrage

restrictions and yet do not allow for profitable arbitrage after transaction costs.

¹⁰To ensure robustness, we also included the subset of options that appear to be mispriced and repeated the VIX calculation. Unreported results indicate that the differences in the VIX index value calculated using the two methods remain quantitatively and qualitatively similar.

REFERENCES

- Aït-Sahalia, Y., and A.W. Lo. "Nonparametric Estimation of State-price Densities Implicit in Financial Asset Prices." *Journal of Finance*, 53 (1998), pp. 499–547.
- Andersen, T.G., T. Bollerslev, P.F. Christoffersen, and F.X. Diebold. "Volatility Forecasting." Working paper, *Penn Institute for Economic Research*, 2005.
- Anagnou, I., M. Bedendo, S. Hodges, and R. Tompkins. "The Relation Between Implied and Realised Probability Density Functions." Working paper, *University of Warwick*, 2002.
- Bahra, B. "Implied Risk-neutral Probability Density Functions from Option Prices: Theory and Applications." Working paper, *Bank of England*, 1997.
- Bakshi, G., C. Cao, and Z. Chen. "Empirical Performance of Alternative Option Pricing Models." *Journal of Finance*, 52 (1997), pp. 2003–2049.
- Bakshi, G., and N. Kapadia. "Delta-hedged Gains and the Negative Market Volatility Risk Premium." *Review of Financial Studies*, 16 (2003), pp. 527–566.
- Bakshi, G., N. Kapadia, and D. Madan. "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options." *Review of Financial Studies*, 16 (2003), pp. 101–143.
- Bates, D.S. "The Crash of '87: Was It Expected? The Evidence from Options Markets." *Journal of Finance*, 46 (1991), pp. 1009–1044.
- . "Post-'87 Crash Fears in the S&P 500 Futures Options Market." *Journal of Econometrics*, 94 (2000), pp. 181–238.
- Black, F., and M. Scholes. "The Valuation of Option Contracts and a Test of Market Efficiency." *Journal of Finance*, 27 (1972), pp. 399–418.
- . "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81 (1973), pp. 637–659.

- Bliss, R.R., and N. Panigirtzoglou. "Testing the Stability of Implied Probability Density Functions." Working paper, *Bank of England*, 2000.
- Bondarenko, O. "Recovering Risk-neutral Densities: A New Nonparametric Approach." Working paper, *University of Illinois at Chicago*, 2000.
- Breedon, D.T., and R.H. Litzenberger. "Prices of State-contingent Claims Implicit in Option Prices." *Journal of Business*, 51 (1978), pp. 621-651.
- Britten-Jones, M., and A. Neuberger. "Option Prices, Implied Price Processes, and Stochastic Volatility." *Journal of Finance*, 55 (2000), pp. 839-866.
- Campa, J.M., P.H.K. Chang, and R.L. Reider. "Implied Exchange Rate Distributions: Evidence from OTC Option Markets." *Journal of International Money and Finance*, 17 (1998), pp. 117-160.
- Carr, P., and D. Madan. "Towards a Theory of Volatility Trading." In Jarrow, Robert A. ed.: *Volatility: New Estimation Techniques for Pricing Derivatives*. London, Risk Publications, 1998.
- Carr, P., and L. Wu. "Variance Risk Premia." Working paper, *Courant Institute and Baruch College*, 2004.
- CBOE. "The VIX White Paper." *Chicago Board Options Exchange*, URL: <http://www.cboe.com/micro/vix/vixwhite.pdf>, 2003.
- Chance, D.M. "Empirical Tests of the Pricing of Index Call Options." *Advances in Futures and Options Research*, 1 (1986), pp. 141-166.
- Demeterfi, K., E. Derman, M. Kamal, and J. Zou. "More Than You Ever Wanted to know About Volatility Swaps." *Quantitative Strategies Research Notes*, Goldman Sachs, 1999a.
- . "A Guide to Volatility and Variance Swaps." *Journal of Derivatives*, 6 (1999b), pp. 9-32.
- Duffie, D., J. Pan, and K.J. Singleton. "Transform Analysis and Asset Pricing for Affine Jump-diffusions." *Econometrica*, 68 (2000), pp. 1343-1376.
- Dupire, B. "Pricing With a Smile." *Risk*, 7 (1994), pp. 18-20.
- Figlewski, S. "Forecasting Volatility." *Financial Markets, Institutions, and Instruments*, 6 (1997), NYU Stern School of Business.
- Galai, D. "Tests of Market Efficiency and the Chicago Board Options Exchange." *Journal of Business*, 50 (1977), pp. 167-197.
- Heston, S.L. "A Closed-form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options." *Review of Financial Studies*, 6 (1993), pp. 327-343.
- Jackwerth, J.C. "Option Implied Risk-neutral Distributions and Implied Binomial Trees: a Literature Review." *Journal of Derivatives*, 7 (1999), pp. 66-82.
- Jiang, G.J., and Y.S. Tian. "Model-free Implied Volatility and Its Information Content." *Review of Financial Studies*, 18 (2005), pp. 1305-1342.
- Lauterbach, B., and P. Schultz. "Pricing Warrants: An Empirical Study of the Black-Scholes Model and its Alternatives." *Journal of Finance*, 4 (1990), pp. 1181-1210.
- MacBeth, J.D., and L.J. Merville. "An Empirical Examination of the Black-Scholes Call Option Pricing Model." *Journal of Finance*, 34 (1979), pp. 1173-1186.
- Merton, R.C. "The Theory of Rational Option Pricing." *Bell Journal of Economics and Management Science*, 4 (1973), pp. 141-183.
- Neuberger, A. "The Log Contract: a New Instrument to Hedge Volatility." *Journal of Portfolio Management*, Winter 1994, pp. 74-80.
- Poon, S.-H., and C.W.J. Granger. "Forecasting Volatility in Financial Markets: A Review." *Journal of Economic Literature*, 41 (2003), pp. 478-539.
- Press, W.H., S. A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical recipes in Fortran 77: The Art of Scientific Computing*, 2nd edition. Cambridge University Press, 1996.
- Rubinstein, M. "Nonparametric Tests of Alternative Option Pricing Models Using All Reported Trades and Quotes on the 30 Most Active CBOE Option Classes from August 23, 1976 through August 31, 1978." *Journal of Finance*, 40 (1985), pp. 455-480.
- Shastri, K., and K. Tandon. "Valuation of Foreign Currency Options: Some Empirical Tests." *Journal of Financial and Quantitative Analysis*, 21 (1986), pp. 57-74.
- Whaley, R.E. "Derivatives on Market Volatility: Hedging Tools Long Overdue." *Journal of Derivatives*, 1 (1993), pp. 71-84.

———. “The Investor Fear Gauge.” *Journal of Portfolio Management*, 26 (2000), pp. 12–17.

Xu, X., and S.J. Taylor. “The Term Structure of Volatility Implied by Foreign Exchange Options.” *Journal of Financial and Quantitative Analysis*, 29 (1994), pp. 57–74.

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@ijournals.com or 212-224-3675

©Euromoney Institutional Investor PLC. This material must be used for the customer's internal business use only and a maximum of ten (10) hard copy print-outs may be made. No further copying or transmission of this material is allowed without the express permission of Euromoney Institutional Investor PLC.

The most recent two editions of this title are only ever available at <http://www.euromoneyplc.com>