Volume 18
Number 4
Summer 2011

www.iijod.com

The Journal of Derivatives

A Publication of Institutional Investor, Inc.

IAFE

The Official Publication of the International Association of Financial Engineers

What Does Implied Volatility Skew Measure?

SCOTT MIXON

SCOTT MIXON
is a director at Lyxor Asset Management Inc. in New York, NY.
scott.mixon@lyxor.com

This article provides theoretical guidance and empirical analysis aimed at differentiating among implied volatility skew measures. Industry analysts and academics use a variety of measures, but most have little formal justification. The author finds that most commonly used skew measures are difficult to interpret without controlling for the levels of both volatility and kurtosis. Many ad hoc measures fail to meet the conditions for a valid skewness ordering. The author's preferred measure is the (25 delta put volatility–25 delta call volatility)/50 delta volatility; among the measures considered, it is the most descriptive and least redundant.

Equity derivatives traders and analysts monitor the implied volatility skew each day, yet there is little practical advice to guide them. What is the best measure of the skew? How should one compare skew for a 25% volatility stock with skew for a 50% volatility stock? Is there any intuition suggesting what "large," "small," or "too much" skew would be? What does the implied vol skew mean, anyway? This article provides theoretical guidance, backed by empirical analysis on S&P 500 Index and single stock options, on these issues.

Exhibit 1 illustrates the practical difficulties facing the derivatives researcher. The chart shows two plausible measures of skew for three-month S&P 500 options, each used by Wall Street strategists. The heavy line shows the 25-delta put volatility minus the 25-delta call volatility. The thin line shows this variable normalized by the 50-delta implied volatility. The series are both stable in 2005 and show a slightly increased skew in 2006. The divergence begins around May 2007, as the normalized series begins suggesting a skew that declines during the next two years, in contrast to the increasing skew of the other measure. At the end of 2009, the normalized measure indicated that the skew was at average levels, whereas the non-normalized version shows a skew that is very steep and well above average ("rich"). A researcher comparing implied volatility for 10% out-of-the-money puts and 10% out-of-the-money calls would also find highly conflicting signals from non-normalized and normalized versions of his series. Both skew measures seem plausible, but which (if either) was telling the truth?

The finance literature has not provided a clear lead in defining a best practice version of skew. This article fills the gap with a rigorous examination of common, ad-hoc skew measures. To prevent confusion, it is worth elaborating on the technical terms used in this study. I always use the terms "skew" or "implied volatility skew" as a measure of the slope of the implied volatility curve for a given expiration date, while "skewness" means the skewness of an option-implied, risk-neutral probability distribution. I do not address the

EXHIBIT 1
Two Measures of Three Month S&P 500 Implied Volatility Skew

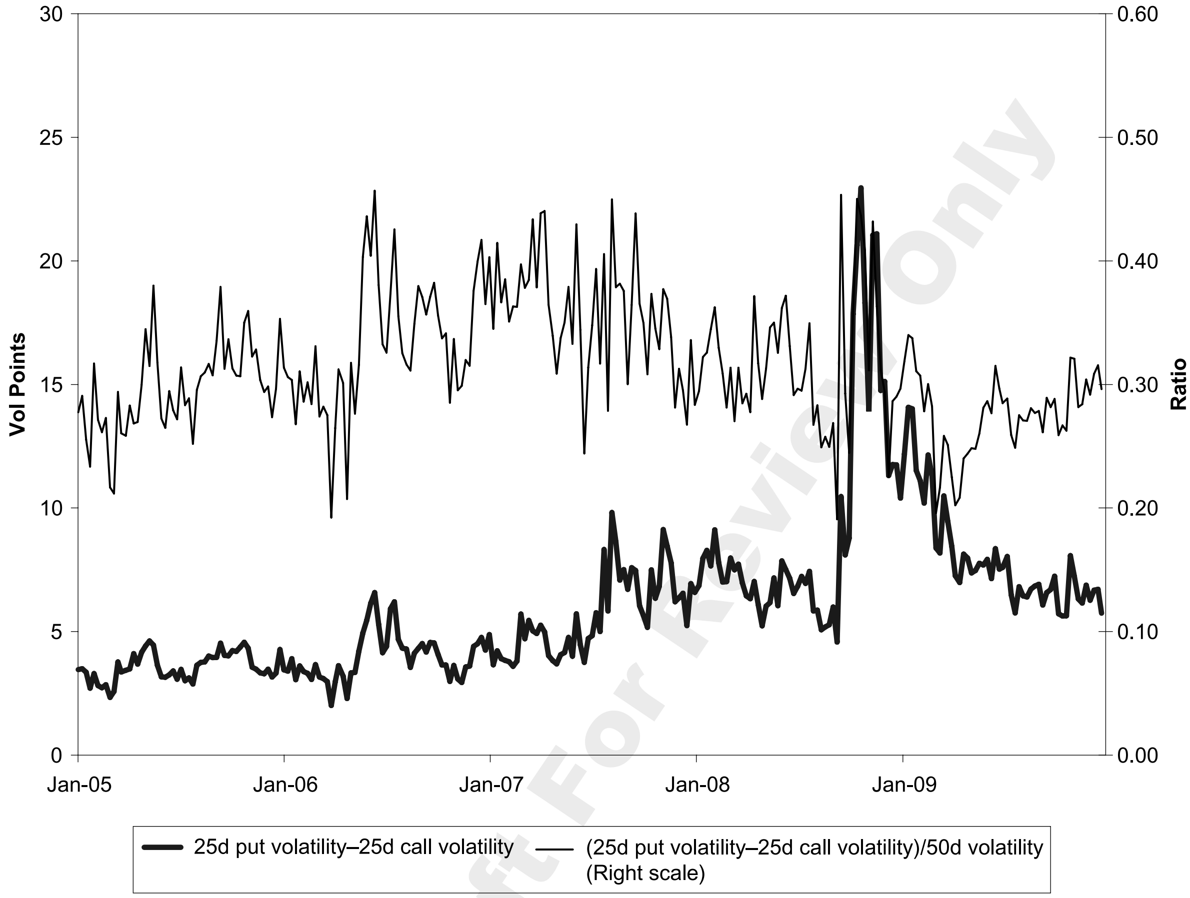

link between implied skewness and the skewness of the underlying asset in the physical world.

A variety of skew measures, most with little theoretical motivation other than plausibility, have been used in previous studies. Some of the methodologies are summarized in Exhibit 2. Researchers have used regression-based values (Bakshi, Kapadia, and Madan [2003]), arithmetic differences in out-of-the-money (OTM) put volatility and OTM call volatility, defined in percentage moneyness, (Bates [1991]), arithmetic differences of put and call volatilities based on delta (Hull, Nelken, and White [2004]), and normalized versions of these last two variables (Toft and Prucyk [1997], Carr and Wu [2007], and Mixon [2009]), among others.

A number of researchers, such as Dennis and Mayhew [2001], have moved toward analyzing a model-free version of the underlying distributional skewness (standardized third central moment) implicit in option prices. The relevant issues are not 100% settled yet: Dennis and Mayhew, for example, reverse the conclusions of Toft and Prucyk by using a different measure of skew. Xing, Zhang, and Zhao [2010] find that future stock returns are predictable based on the implied volatility skew, but returns are not predictable based on option-implied skewness.

Previous researchers have used series expansions and linear approximations to derive analytical expressions relating implied volatility to characteristics of the underlying's distribution. Backus, Foresi, and Wu [2004] find that their approximations work best for options on underlyings with a modestly skewed distribution. Carr and Wu [2005] extend the analysis and approximate skewness as proportional to the difference between volatilities for a call and a put equally out of the money in delta terms (i.e., a risk reversal), divided by the at-the-money implied volatility.

Exhibit 2

Select Skew Measures in the Literature

Authors	Measure
Bakshi, Kapadia, and Madan [2003]	Slope from regressing log volatility on log moneyness
Bates [2000]	OTM call price/OTM put price – 1
Bates [2000]	4% OTM implied volatility – ATM implied volatility
Cremers, Driessen, Maenhout, and Weinbaum [2008]	8% OTM put volatility – ATM implied volatility
Carr and Wu [2007] Mixon [2009]	(25 delta put volatility – 25 delta call volatility)/50 delta volatility
Gemmill [1996]	2% OTM implied volatility–par volatility –2% OTM implied volatility)/ATM volatility
Hull, Nelken, and White [2004]	25 delta call volatility – 50 delta volatility
Natenberg [1994]	OTM put volatility/ATM volatility
David and Veronesi [2009]	OTM put volatility/OTM call volatility
Toft and Prucyk [1997]	(10% OTM call volatility –10% OTM put volatility)/ATM implied volatility

Bakshi, Kapadia, and Madan (henceforth, BKM) [2003] also connect risk-neutral moments and the implied volatility skew; their empirical tests provide substantial support for their theory. This study builds on BKM's research to provide further intuition on the structure of the skew and its connections with some standard practitioner rules of thumb for measuring the skew. The analysis provides practical suggestions on choosing from and interpreting some common ad-hoc skew measures.

The contributions of this article are as follows. First, I provide a theoretical examination to determine whether several commonly used measures of the implied volatility skew define a valid skewness ordering as defined in the statistics literature. Second, I use regression analysis to test the theory against both index and single-stock option data. I empirically decompose these ad-hoc measures into the portions attributable to the volatility, skewness, and kurtosis of the underlying risk-neutral distribution implicit in option prices. The decomposition provides insight into the characteristics each skew measure is actually capturing.

As a by-product of the analysis, I provide economic intuition about interpreting risk-neutral moments derived from option prices. For example, BKM reported that the average level of skewness implicit in S&P 100 options from 1991–1995 was–1.09, while the average skewness for single stocks over that period was much closer to zero (averaging–0.31 for the 30 stocks they examined). How does one interpret those numbers in terms of economically meaningful constructs? This article provides clear direction on that issue.

I find the following results. Ceteris paribus, the ad-hoc skew measures move in the appropriate direction when there is a change in the skewness for the distribution of the underlying asset, but the relation is generally quite dependent on the level of at-the-money (ATM) volatility. Such measures do not appear to satisfy a key property to be a valid measure of skewness (location and scale invariance). In practical terms, analysis of a time series of an ad-hoc implied volatility skew may provide few clues to the absolute level of skewness or its economic importance. At best, week to week changes in the volatility skew might be a reasonable proxy for changes in skewness; similarly, movements in the level of skew over short intervals might be an acceptable proxy when other factors (e.g., volatility and kurtosis) are roughly constant. It is clear, however, that these measures should be treated with caution and represent the amalgamation of several factors. They are not directly comparable over time (unless volatility is constant) or in the cross-section (unless volatility is the same for all assets).

The dependence of the ad hoc skew measures on the level of ATM volatility also means that regressions incorporating these measures as a proxy for skewness are likely misspecified. If a dependent variable is regressed on, say, 90% strike volatility minus 110% strike volatility as a proxy for skewness, the regression omits the interaction term to control for the dependence of skew on other factors. Including the level of volatility on the right-hand side is not a sufficient control, as the omitted variable is a product of volatility and skewness. The resulting bias is likely to render statistical inference on such regressions unreliable.

Analysts try to work around the difficulties in interpreting these statistics. A common practice is to compare some of the more intuitive measures over a trailing window of time such as one year to accommodate different regimes of volatility or to examine stocks in a single sector at a time. These solutions have practical appeal, but they might not solve the problem.

Other researchers have focused on moment-based measures of skewness, which are theoretically attractive (see, for example, Bank of England [2009, pp. 11-12]). Moment-based measures are unfortunately not very interpretable to most market participants or researchers. More seriously, these measures might be very sensitive to the prices of deep out-of-the-money options that trade infrequently. The resulting estimates might be noisy estimates of the moments of interest.

One solution to this problem is to focus on a central portion of the option-implied distribution and to rely very little on the prices or implied volatilities of illiquid options in the wings of the distribution. Doksum [1975] proposed such central skewness measures in the statistics literature as a robust way to quantify skewness without relying too much on tail behavior. Volatility skew measures relying mostly on liquid options might also be more relevant as statistics for trading strategies that are actually implementable. One such measure is (25 delta put volatility–25 delta call volatility)/50 delta volatility, which emerges as the preferred skew measure based on the theoretical and empirical analysis presented here. This measure has minimal dependence on the level of implied volatility and is therefore the least redundant descriptor of higher order moments among the variables considered. As a description of "market sentiment," it neatly encapsulates relevant information on skewness and kurtosis implicit in the cross section of option prices.

The structure of the article is as follows. First, I provide theoretical results relating various implied volatility skew measures to the level of skewness and the level of at-the-money implied volatility. Numerical and analytical results are based on a flexible, intuitive benchmark model of implied volatility (volatility that is linear in option delta). In the second section, I provide empirical analysis using S&P 500 option skew over the 2005–2009 period and relate several skew measures to the so-called model-free skewness and kurtosis measures considered by BKM. In particular, I use regression models suggested by the theoretical analysis in the first section to differentiate among the various practitioner skew measures. In the third section, I examine the implied volatility skews of a cross-section of stocks to validate the robustness of the results. The final section provides concluding comments.

THEORY

This section demonstrates some theoretical results to set expectations for the empirical analysis in the next section. I assume an intuitive, plausible model for skewness of the option-implied distribution of stock returns and compute the implications for several popular measures of skew. The theoretical framework assumes that the option implied distribution generates an implied volatility function that is linear in delta.

Implications of the Linear Skew in Delta

The model has some very attractive properties. First, it is empirically plausible. For example, when I fit the cross-section of S&P 500 implied volatility each week in 2005–2009 (for the expiry nearest 90 days), the regression R^2 averages 95%.[1] The model can be seen as a more parsimonious version of the implied volatility by delta interpolation used by researchers such as Bliss and Panagirtzoglu [2004] who construct option-implied probability density functions. Similarly, Derman et al. [1999] preferred the linear in delta model over the linear in percentage strike model when computing Taylor series approximations for variance swap strikes. The linear in delta model is convenient for extrapolating beyond observed strikes: implied volatility does not blow up in the wings of the distribution; instead, it tends to asymptote at extreme strikes.

Second, this framework provides enough structure so that one can consider skew (e.g., 90% strike put volatility minus 110% strike call volatility) as a function only of at-the-money implied volatility and the skewness of the implied distribution. For a given level of ATM volatility, there is a one-to-one mapping between skewness and the slope of the skew. Hence, I do not make any assumptions about the dynamic properties of the option-implied or actual distribution of the underlying, but

I can obtain rich, sharp conclusions by varying only two parameters.

Nonetheless, it is worth noting that excess kurtosis, by itself, can also generate non-zero values for these skew measures. Numerical tabulations by Das and Sundaram [1999] and the approximation functions in Backus, Foresi, and Wu [2004] suggest that a 90–110 skew, for example, could be due to kurtosis. BKM conclude that skewness is more important than kurtosis for equity options, and that is obviously the direction followed in this study.

Exhibit 3 uses the model to put observed values of skewness into human terms. What sort of intuition should we get for option valuation when we know that skewness is, say, –0.5 or –1.5? The chart demonstrates that a given value for skewness in this model leads to different implications for options as the level of ATM volatility changes. I plot the probability of an extreme outlier—a return less than –2.33 times the ATM implied volatility—on the vertical axis and the skewness of the distribution on the horizontal axis. The chart shows this relation for ATM values ranging from 10% to 50%. All of the probabilities are computed under the linear vol skew in delta model with a three-month maturity. Interest rates are set to 4.17% and dividend yields are set to 2%; these are representative values for the 2005–2009 data used in the empirical analysis. For ease of comparison, the means of the distributions have been adjusted by a value of $\sigma^2/2$ so that all of the distributions have a 1% probability of such a crash if there is no skewness.

The shapes of the level curves are intuitive. Moving from a scenario with zero skewness to a scenario with a skewness of −0.5 doubles the likelihood of an extreme outlier to 2%. A scenario with a skewness of −1.5 shows much greater dependence on the level of ATM volatility. At a 10% volatility, the crash probability is around 5.5%, but the probability increases as volatility is assumed to be at a higher level. With a 30% ATM volatility, the probability is 6%; a 50% volatility means an 8% probability of an extreme outlier.

EXHIBIT 3

Crash Probability under Linear Volatility Skew in Delta

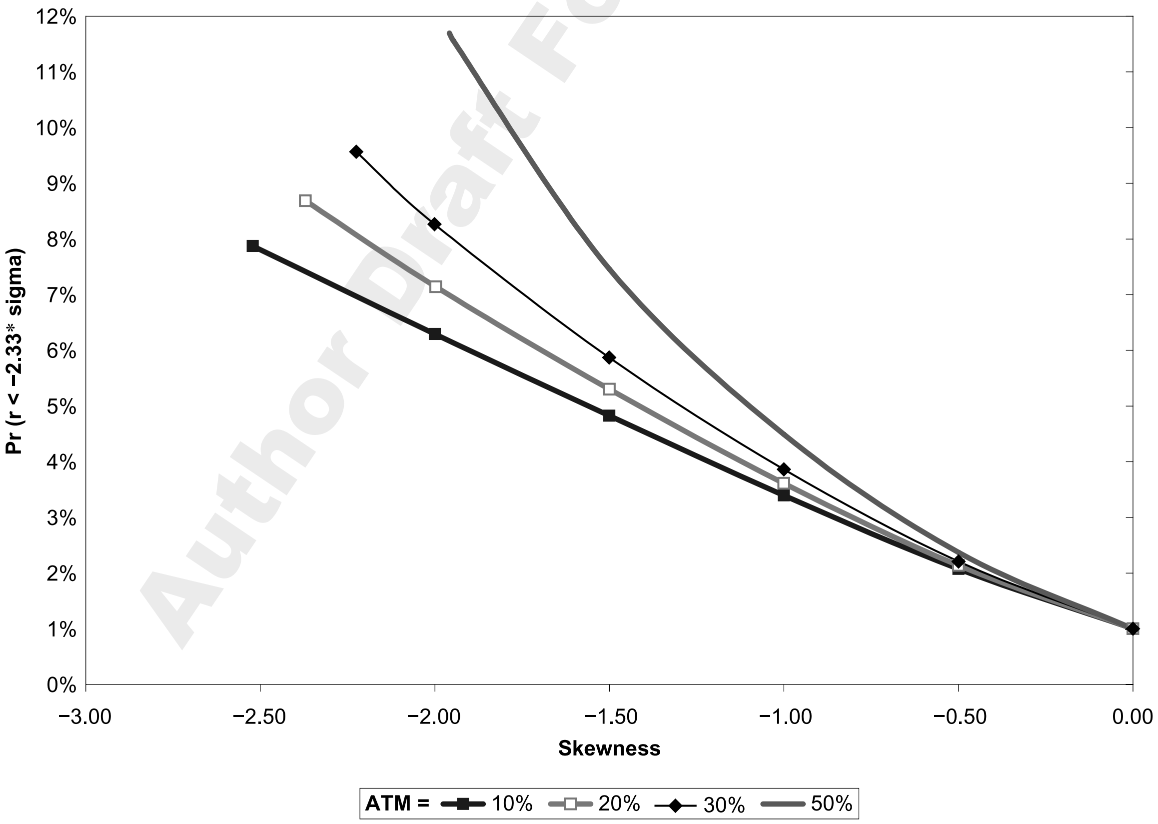

Exhibit 4 illustrates implications of the theory across several ad-hoc skew measures. I assume an ATM volatility and use the linear skew in delta model to compute implied volatility skew as a function of distributional skewness and the fixed ATM level. The option maturity is set to three months, and interest rates and dividend yields are the same as in Exhibit 3. I repeat this for ATM values of 10%, 20%, 30%, and 50%. Each solid line in the exhibit maps out an ad-hoc skew measure as a function of skewness for a given ATM value.

The top two charts in Exhibit 4 display the results for the 90% volatility minus the 110% volatility (left hand chart) and 90% volatility minus 110% volatility, divided by 100% volatility (right hand chart). The bottom two charts display results for delta versions of the charts at the top. The left chart is for 25-delta put volatility minus 25-delta call volatility, whereas the right chart is for (25-delta put volatility minus 25-delta call volatility)/50-delta volatility. (For the moment, ignore the black dots on the chart, as these are discussed in the next section).

All of the charts show a similar headline story, in that the level curves for given ATM values all slope downward. Other factors held constant, more left skewness means a higher value for these skew measures. Yet the slope of the level curves shows considerable variation and dependence on the level of ATM volatility. For example, consider the 90–110 skew measure (top left panel of Exhibit 4) when the underlying skewness is

EXHIBIT 4

Implied Skewness vs. Implied Volatility Skew Measures under Linear Volatility Skew in Delta

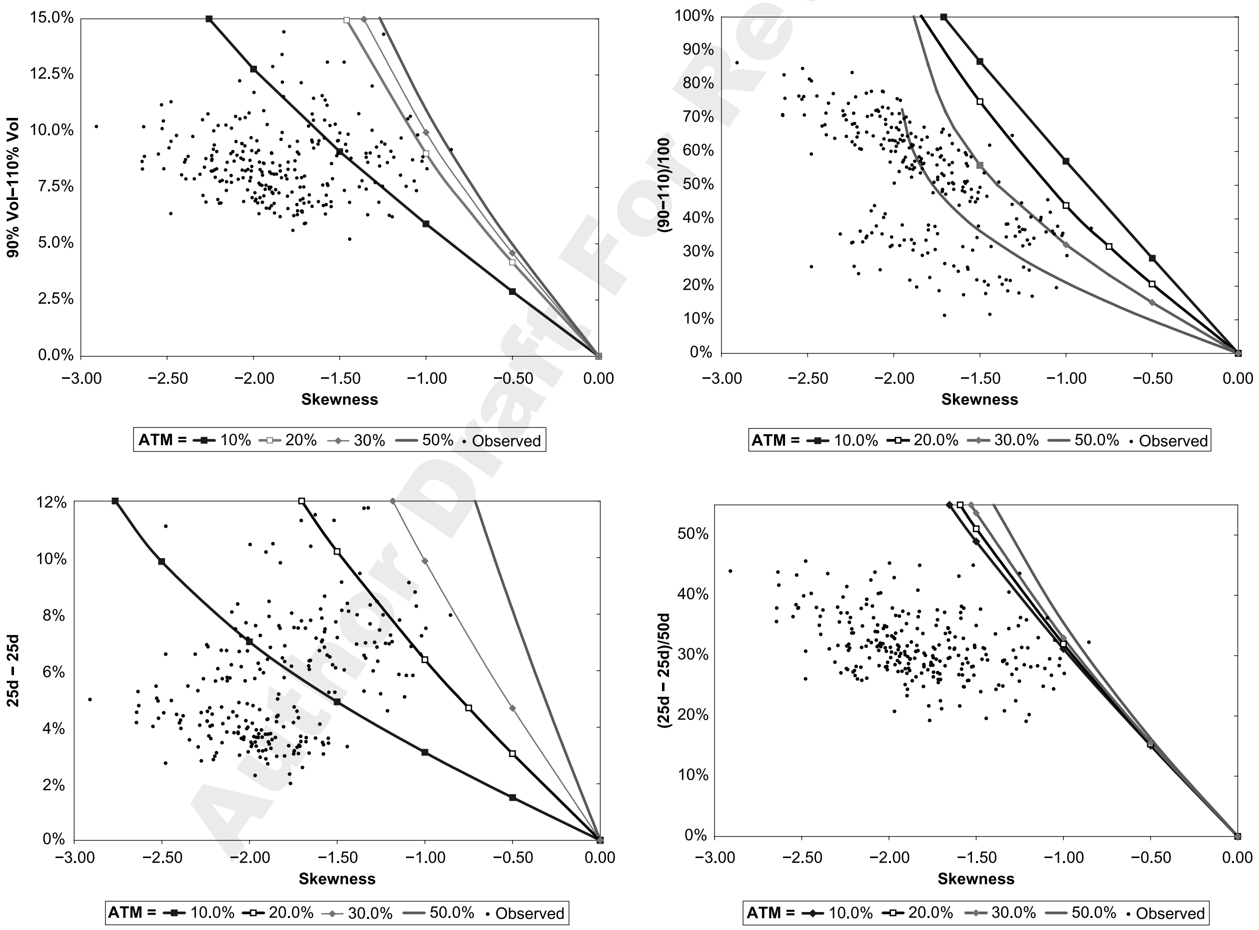

Notes: Solid lines map out theoretical values of four popular volatility skew measures as a function of option-implied skewness, as predicted by the linear-in-delta volatility skew model. Each panel plots the function for at-the-money (ATM) volatility levels of 10%, 20%, 30%, 40%, and 50%. Dots represent observed (skewness, skew) pairs for S&P 500 three-month options, 2005–2009.

held fixed at −1. The 90–110 skew is six volatility points for an ATM volatility of 10%; skew must *increase* to nine volatility points to be consistent with the same skewness and ATM volatility of 20%. Moving from the top left to the top right chart, the opposite behavior is seen. At a 10% volatility and skewness of −1, the (90–110)/100 skew is 0.55, but it must *decline* to 0.40 to be consistent with the same skewness at an ATM volatility of 20%. Neither measure suggests a direct link between volatility skew and skewness. For a given asset, the implied volatility skew could increase, decrease, or stay the same when skewness changes, depending on movements in volatility. Similarly, a cross-sectional comparison of skew would not correspond to a cross-sectional ranking of skewness if the underlying assets exhibit differing levels of volatility. The same level of skew means something different at different volatility levels.

Perhaps the most interesting chart is the bottom right one. For the (25-delta put volatility–25-delta call volatility)/50-delta volatility, there is virtually no dependence on the ATM volatility. Roughly speaking, more left skewness means a higher value for this measure of skew, irrespective of the level of volatility. This ease of interpretation makes it a very attractive measure according to the theory.

Importing Rigor from the Statistics Literature

We can be more systematic about differentiating among various skew measures. Van Zwet [1964] introduced the notion of ordering two distributions with respect to skewness, and an acceptable measure should respect that ordering. For random variables with continuous CDFs $F(x)$ and $G(x)$ and PDFs $f(x)$ and $g(x)$ having interval support, $G(x)$ is defined to be at least as skew to the right as $F(x)$ if $G^{-1}\{F(x)\} = R(x)$ is convex. The notation is that $F <_c G$ ("*F* c-precedes *G*"). Oja [1981] shows that a sufficient condition for $F <_c G$ is that the standardized distribution functions F_S and G_S cross each other exactly twice, with the final sign of $F_S(x) - G_S(x)$ being positive.

Groeneveld and Meeden [1984, pp. 392-393] follow van Zwet [1964] and Oja [1981] and define four properties that a valid skewness functional γ should satisfy. The four properties are

- P1. A scale or location change for a random variable does not alter γ. Thus, if $Y = cX + d$ for $c > 0$ and $-\infty < d < \infty$, then $\gamma(X) = \gamma(Y)$.
- P2. For a symmetric distribution $\gamma = 0$.
- P3. If $Y = -X$ then $\gamma(Y) = -\gamma(X)$.
- P4. If F and G are cumulative distribution functions for X and Y, respectively, and $F <_c G$, then $\gamma(X) \leq \gamma(Y)$.

Van Zwet [1964] showed that the standardized third central moment ("the" skewness coefficient) satisfies P1 to P4 and is therefore a valid skewness measure. It is difficult to interpret, however. In the analysis that follows, I use the standardized third moment as a benchmark for the simpler practitioner measures.

Almost all of the measures I examine display a dramatic failure of P1. The results, foreshadowed in the tabulations above, are based on the following test. I first fix a level of skewness for the underlying distribution and then compute the skew measures for a range of risk-neutral volatility levels. For ease of comparison, each measure is scaled by its value at 10% ATM volatility. Each rescaled measure is therefore indexed to 100% at 10% volatility and, if the measure satisfies P1, should equal 100% when evaluated at any level of volatility. Results are displayed in Exhibit 5 for a skewness of −0.5; additional analysis available from the author shows that the results are robust to different levels of skewness.

The 90–110 skew nearly doubles as volatility increases from 10% to 50%, despite the constancy of the underlying skewness. The value rises monotonically, tracing out a concave function of volatility. The value at 50% volatility is approximately 75% larger than the value at 10%. The (90–110)/100 skew measure fares just as poorly, declining by two-thirds as volatility moves from 10% to 50%. This measure declines monotonically, tracing out a convex function of volatility.

The 25-delta put volatility minus 25-delta call volatility fails in even more spectacular fashion. This measure rises almost linearly, with the value evaluated at 50% volatility a full five times the value at 10% volatility. As noted earlier, however, when this delta-based measure is divided by 50-delta implied volatility, it exhibits minimal variation as a function of volatility. When this statistic is computed at 50% volatility, it is approximately 1.05 times the value exhibited at 10% volatility. Although none of the measures strictly conform to P1, this normalized version is extremely close to constant

EXHIBIT 5

Relative Values of Implied Volatility Skew Statistics at Various Levels of Volatility

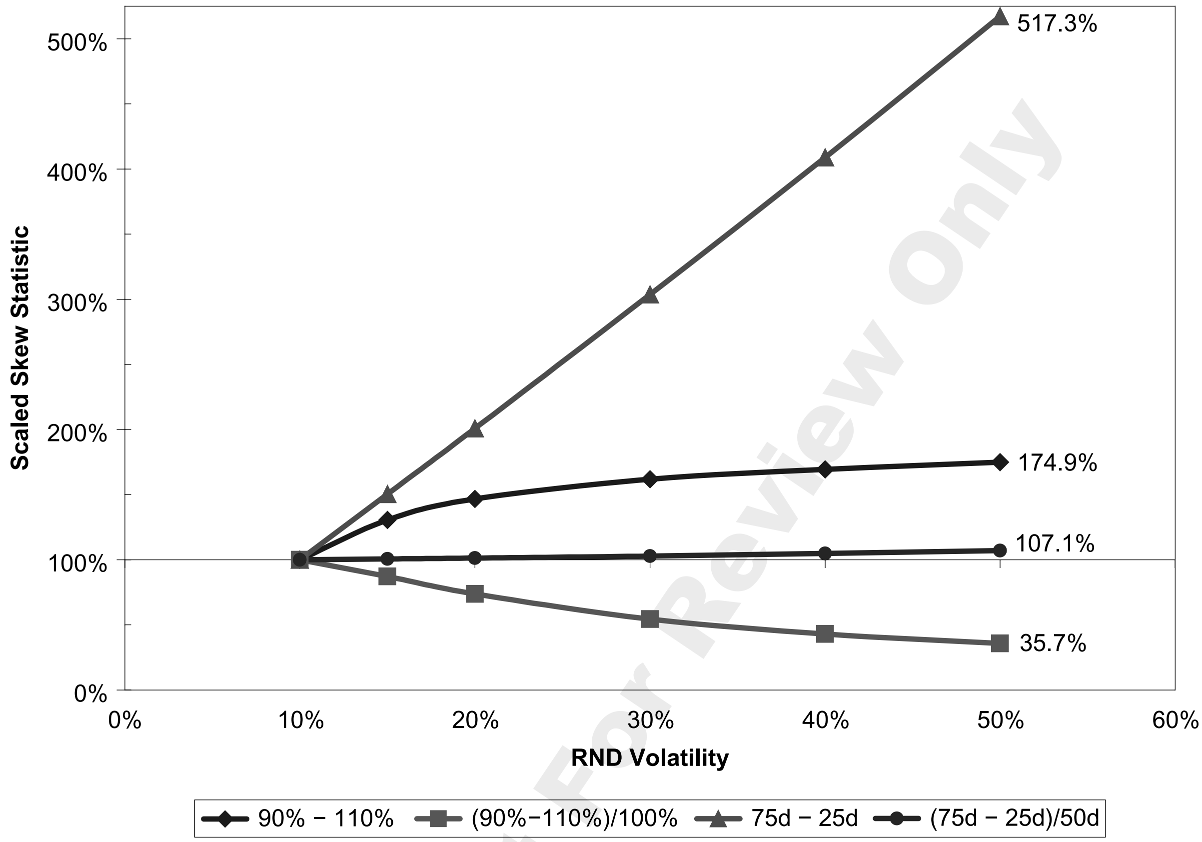

when compared to the other measures. By this criterion, only the normalized, delta-based measure passes the test.

All of the skew measures appear to exhibit the proper characteristics to pass the test for P2 and P3, and so I do not dwell on them. Property P4, however, poses a bigger challenge. It is straightforward to find examples of standardized distributions F_S and G_S such that $F <_c G$ (i.e., F is more left-skewed than G), but the volatility skew measures suggest that G is more skewed to the left than F. The intuition is that, for some measures, the strong correlation with the level of volatility can overpower the impact of the asymmetry. For example, a distribution that is significantly left-skewed can exhibit a very small absolute value for the 90–110 skew if the level of volatility is quite high. The 10% below spot/10% above spot strike range measured by the 90–110 skew can be enormously wide for a low-volatility stock but relatively small for a high-volatility stock, resulting in an apples to oranges comparison.

One can readily use Oja's [1981] sufficient condition to show that $F <_c G$ for some parameter configurations that generate such misleading results for skew measures; I provide some examples in an unpublished appendix. The bottom line is that comparing most percentage skew measures at very different levels of ATM volatility can give misleading results. The (25 delta put volatility–25 delta call volatility)/50 delta volatility measure is the only one considered here that appears relatively immune to this problem.

Discussion Based on Analytical Approximations

The linear in delta skew model for implied volatility offers analytical tractability to explore these results.

The final part of this section sketches the approximation derivations.

First, parameterize the implied volatility as $\sigma_m = \sigma_{50\Delta} + b(N(d_1^m) - \frac{1}{2})$, where σ_m is implied volatility for moneyness m, $\sigma_{50\Delta}$ is the 50 delta implied volatility, and d_1^m is the value of d_1 corresponding to moneyness m. The variable d_1 is defined in the usual manner as $d_1 = \frac{\ln(S/K)+(r-q+\sigma^2/2)T}{\sigma\sqrt{T}}$ for strike price K, riskless interest rate r and dividend yield q. Note that all the delta measures are expressed in terms of call deltas in this analytical section. Whereas the previous text talks about, say, the volatilities of 25 delta calls and 25 delta puts, the expressions in this section refer to volatilities of 25 delta calls and 75 delta calls and writes them as $\sigma_{25\Delta}$ and $\sigma_{75\Delta}$. This approximation is convenient and should not impact the qualitative results.

For moneyness defined by call delta, we can simplify the skew measures easily: $\sigma_{75\Delta} - \sigma_{25\Delta} = b(N(d_1^{75\Delta}) - N(d_1^{25\Delta})) = (0.5)b$. Using the linear approximation for the standardized normal cumulative distribution $N(z)\frac{1}{2}+\frac{1}{\sqrt{2\pi}}(z)$, which is quite good in the region $z = -0.75$ and $z = 0.75$, yields the approximate relation $\sigma_{90\%} - \sigma_{110\%} = \frac{1}{\sqrt{2\pi}}\frac{b\ln(11/9)}{\sigma_{50\Delta}\sqrt{T}}$ for percentage moneyness.

Note that the parameter b represents a cross-sectional fit parameter on a given date; there is no assumption that it is constant over time. One simple way to proceed is to note that the variable b can be numerically approximated by the function $b_1\sigma_{50\Delta}$ for a fixed level of skewness.[2] We can therefore approximate $\sigma_{75\Delta} - \sigma_{25\Delta} = (0.5)b_1\sigma_{50\Delta}$, and it is obvious that this skew measure based on the arithmetic difference between equally spaced in- and out-of-the-money options exhibits a strong, approximately linear dependence on the level of volatility. Dividing this measure by $\sigma_{50\Delta}$ removes the dependence. This is precisely the result demonstrated in Exhibit 5 for the two delta-based measures.

Perhaps more interestingly, we also obtain the approximation $\sigma_{90\%} - \sigma_{110\%} = \frac{1}{\sqrt{2\pi}}\frac{\ln(11/9)}{\sqrt{T}}(b_1)$ from these results. This expression suggests that there is very little dependence on the level of volatility for the fixed percentage strike skew measure, and that it is perhaps equally as good as the normalized delta-based measure along this dimension. Further reflection shows why this result obtains and why it is misleading. The first order approximation to the standardized normal cumulative distribution used in this derivation is good as long as we approximate the distribution from about three-quarters of a standard deviation below the mean to about three quarters of a standard deviation above the mean. For the three-month option maturity used in these examples, this region of good approximation includes 10% below the mean and 10% above the mean only if the annualized volatility is above roughly 27% (because $\pm 10\% \approx (\pm 0.75) \times 27\%/\sqrt{4}$). At lower levels of volatility, these two strikes are too far from the mean for a linear approximation to the normal CDF to be adequate. Examination of Exhibit 5 confirms this explanation, as this skew measure shows very little dependence on the level of volatility, as long as volatility is above around 27%. For lower values, the function is clearly increasing. It might still be useful to have analytical expressions for these volatilities, and we could replace the linear approximation to the standardized cumulative normal distribution with an ad hoc approximation such as $N(z) \approx \frac{1}{1+e^{-1.702z}}$ (Bowling et al. [2009]). Unfortunately, the resulting expressions are complicated and might be more useful for simplifying applied work than for generating theoretical insights.

The bottom line from this theoretical analysis is that (25-delta put volatility–25-delta call volatility)/50-delta volatility gives useful results, whereas the other measures can give highly misleading results. The analysis gives sharp results but is predicated on a plausible model describing the cross-section of implied volatility on a given date. The next section takes the predictions of the model and tests them against real world option data without imposing this constraint.

SKEW IN THE TIME SERIES: S&P 500

BKM describe a method for computing skewness and kurtosis of the option implied distribution, and I implement their method in this section. The idea is to compare these "true" statistics with ad-hoc skew measures computed using the same S&P 500 option data. The results are highly consistent with predictions of the model developed in the previous section.

Exhibit 6 displays the BKM skewness computed for the S&P 500 over the period 2005–2009, as well as the 50-delta implied volatility. The underlying data are end of week values, and both series are interpolated from nearby expiration dates to form constant maturity 90-day horizons. I use the standard expiration days and remove data for weekly or quarter-end options. The chart clearly shows the low-volatility regime from 2005–2006, the high-volatility regime in 2007 and much of 2008, and

Exhibit 6

Three-Month Skewness and Implied Volatility for the S&P 500

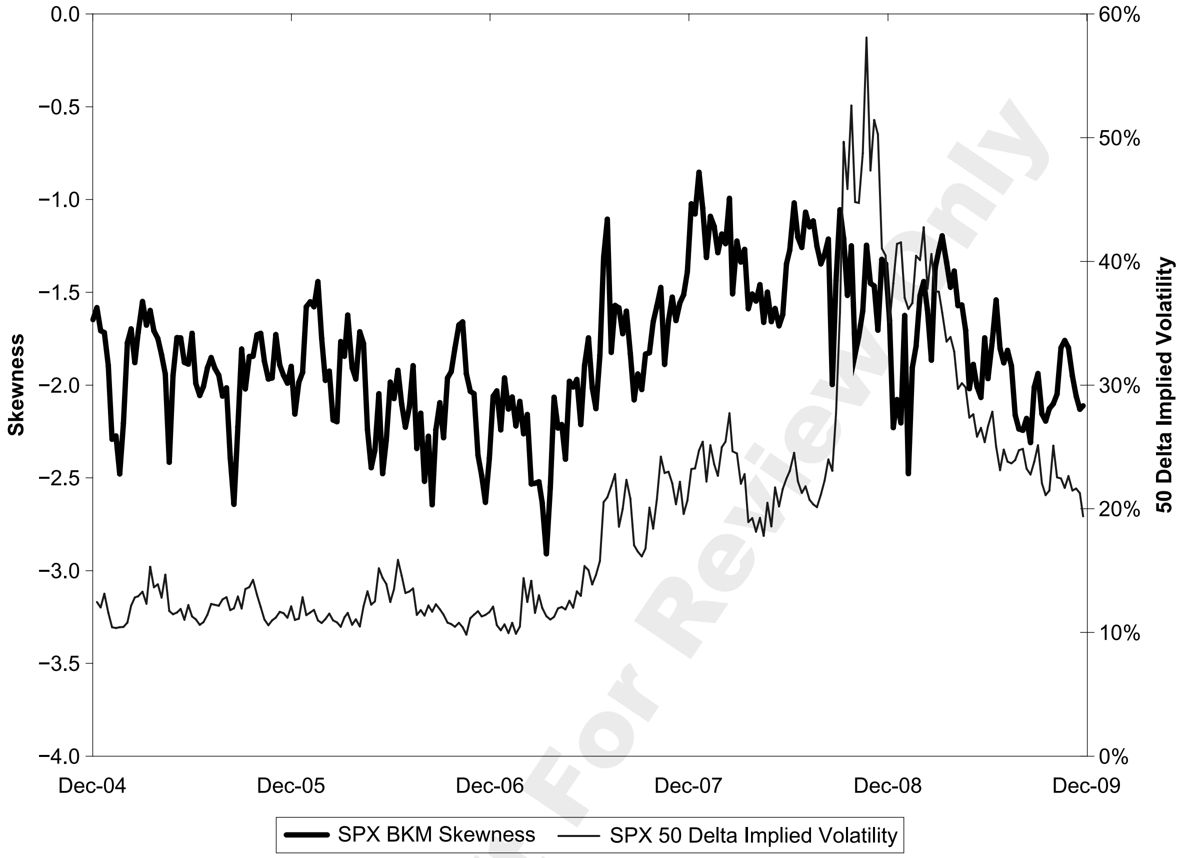

the ultra-high-volatility regime in late 2008. Similarly, the skewness was range bound around −2.0 during the first few years of the sample, rising to a range around −1.5 as volatility increased.[3]

The linear in delta benchmark model and its predictions plotted in Exhibit 3 help to give intuition to these values. So, for example, when S&P implied volatility was around 10% during 2005–2006 and BKM skewness was approximately −2.0, the linear volatility skew in delta model suggests a the probability of an extreme outlier (i.e., a crash) around 6.5%. And when S&P volatility moved to the 20% range from mid-2007 to mid-2008, skewness had moved to −1.5. The crash probability declined to around 5.5%. To the extent that kurtosis makes those crash scenarios even more likely, given skewness and volatility values, the probabilities would be even higher under a model incorporating kurtosis. Nonetheless, these probabilities appear to be an economically large jump from the 1% value implied by a symmetric, normal distribution.

Exhibit 7 answers the question "Do implied volatility skew measures track skewness?" Each row shows the results of regressing BKM skewness for the S&P 500 on an ad-hoc measure of skew for the 2005–2009 period.[4] The first three columns of numbers display results from a regression on levels. The first two rows are for the 90–110 skew and (90–110)/100 skew. The next row is a common measure of skew describing a cost of protection: the cost (as a percentage of spot) of the put struck at 95% of spot minus the cost of the call struck at 105% of spot. The last two rows are for the 25-delta put volatility minus the 25-delta call volatility and the normalized version of this measure, respectively.

Results from the levels regression are striking: The only two measures with the expected negative slope coefficient are the two normalized measures. The slope coefficients are positive for the non-normalized

EXHIBIT 7

Results from Regressing BKM Skewness on Various Implied Volatility Skew Measures

	Levels Regression			First Differences Regression		
Independent Variable	**Slope**	**(*t*-stat)**	**R^2 (%)**	**Slope**	**(*t*-stat)**	**R^2 (%)**
90% vol–110% vol	0.01	(0.5)	0.2	−0.04	(1.9)	5.0
(90% vol–110% vol)/100% vol	−1.30	(4.6)	35.7	−1.77	(8.6)	22.8
95% put price–105% call price	0.19	(1.6)	2.6	−0.19	(1.4)	3.6
25d put vol – 25d call vol	0.04	(3.7)	11.1	−0.02	(1.3)	1.6
(25d put vol – 25d call vol)/50d vol	−2.73	(3.4)	14.7	−2.01	(6.5)	16.3

Notes: Values in parentheses are t-statistics computed using Newey-West standard errors with 13 lags. The data are end-of-week values for the S&P 500 from January 7, 2005, to December 24, 2009.

measures (significantly so for the delta version), and the slope is positive and significant for the put price minus call price measure. The regression results suggest that most of the skew measures are below average when skewness is high in absolute value and above average when the distribution is more symmetric. Contrary to the typical interpretation, these measures move in the wrong direction if they are supposed to represent the skewness of the underlying distribution.

The last three columns in the table are for regressions on first differences. Every regression shows a negative slope coefficient, suggesting that the implied volatility skew gets steeper when skewness becomes more negative. The regressions suggest that changes in skew do move in the appropriate direction. However, the regressions on non-normalized versions, when compared to the regressions on the normalized values, have R-squareds that most analysts would view as pitiful. Regressions on non-normalized skew measures or on the put minus call price yield R-squareds below 5%, whereas the two normalized versions have R-squareds of 15% to greater than 25%. To the extent that differencing focuses the attention on skewness and eliminates much of the dependence on changes in volatility and kurtosis, all of the measures at least move in the correct direction as an indicator of skewness. Analysts focusing on, say, week-to-week changes in these measures may be able to infer skewness changes from any of them, but they could easily be misled unless everything else really does stay constant.

The regressions in Exhibit 7 are naive. They may be helpful with a visual inspection of a chart, but they do not provide deep insight into the sources of the skew. They do not rely on any of the theory developed in earlier sections and consequently omit relevant variables. The regression specifications in Exhibit 8 rely on, and extrapolate from, the theory developed earlier. The headline conclusion is that the implied volatility skew generally impounds information about the volatility, skewness, and kurtosis of the underlying implicit distribution, and the linkages are consistent with the model outlined in the previous section.

Before moving to these more comprehensive regressions, it is worthwhile to revisit the charts in Exhibit 4. The black dots in the charts are a scatterplot of observed data for S&P 500 three-month options using weekly data from 2005–2009. One key observation is that the majority of the observed data lie below the model's predicted values, irrespective of the skew measure. The linear-in-delta volatility model cannot generate skewness levels as negative as the BKM values seen for the S&P 500 in the real world. This suggests a significant role for kurtosis in exaggerating the impact of skewness relative to the benchmark model.

The foregoing analysis suggests the following regression specification:

$$skew = a_0 + (b_0 + b_1 \times volatility + b_2 \times kurtosis) \times (skewness)$$

Based on the intuition outlined above, we expect the following signs for slope coefficients. The coefficient b_0 is expected to be negative for all of the implied skew measures; more negatively skewed distributions have higher levels of implied volatility skew. The coefficient b_1 is expected to be negative for most cases, meaning that the slope of the level curves are steeper at higher levels of

EXHIBIT 8

Results from Regressing Various Implied Volatility Skew Measures on BKM Skewness and BKM Skewness Interacted with ATM Volatility and BKM Kurtosis

Dependent Variable	Skewness	(*t*-stat)	Volatility × Skewness	(*t*-stat)	Kurtosis × Skewness	(*t*-stat)	Adjusted R^2 (%)
90% vol–110% vol	0.23	(0.45)					−0.2
	0.20	(0.38)	−4.53	(2.34)			14.3
	−5.73	(3.34)	−0.12	(0.05)	0.20	(3.79)	31.3
(90% vol–110% vol)/100% vol	−0.27	(6.44)					35.4
	−0.27	(11.25)	0.77	(7.85)			81.3
	−0.42	(4.88)	0.88	(6.54)	0.01	(2.02)	82.5
95% put price–105% call price	0.14	(1.75)					2.3
	0.13	(3.21)	−1.51	(5.44)			55.4
	−0.34	(1.81)	−1.16	(4.02)	0.02	(2.49)	58.8
25d put vol – 25d call vol	2.89	(3.19)					10.8
	2.75	(5.05)	−17.65	(7.13)			79.4
	1.92	(1.13)	−17.03	(5.28)	0.03	(0.63)	79.4
(25d put vol – 25d call vol)/50d vol	−0.05	(3.28)					14.4
	−0.05	(3.32)	0.01	(0.10)			14.0
	−0.20	(4.07)	0.11	(1.16)	0.00	(3.43)	25.3

Notes: Values in parentheses are t-statistics computed using Newey–West standard errors with 13 lags. Cells shaded gray have t-statistics with an absolute value greater than 2. The data are end of week values for the S&P 500 from January 7, 2005, to December 24, 2009.

volatility. The expected sign for b_2 is positive, as higher kurtosis is expected to flatten out the level curves.

Estimated slope coefficients and robust t-statistics are displayed in Exhibit 8. Shaded cells in the table represent coefficients with t-statistics with absolute value greater than two. The results generally indicate that skew is related to skewness, but there is a substantial role for the interaction with volatility and kurtosis. Incorporating all three variables raises the adjusted R^2 values from 10 to 70 percentage points. The implied volatility skew reflects more than skewness.

These regression results can be used to distinguish among the skew candidates. I first ask if the slope coefficients conform to the sign implied by the theory, and I focus on results for the full regression models. The model predicts a negative coefficient on skewness, and this is generally verified quite strongly. The 25-delta put volatility minus 25–delta call volatility is the exception, but the dominant effect in this case is the relation of the measure with the level of volatility, which is consistent with the prediction of the theoretical model. The generally negative coefficient for the volatility interaction term is verified in the empirical results. The most interesting case is for the (25-delta put volatility–25-delta call volatility)/50-delta volatility, which has virtually no role for this variable. Consistent with the model's predictions, the level of volatility has very little impact on this skew. Finally, the hypothesized positive coefficient for the kurtosis interaction term is strongly verified in almost all cases.

Some additional tests confirm the conclusions from the table. I also ran the regressions including the kurtosis variable in addition to the interaction terms. Regression R^2 values generally rise just a few percentage points, the significant variables from the regression without the non-interacted kurtosis term generally remain significant, and the kurtosis term is highly significant. The signs on the kurtosis terms are negative, and the resulting regression fits can generate skew measures that are positive or negative even if skewness is zero.

The overall message from the empirical exercise is that the predictions of the linear-in-delta model for implied volatility are upheld in S&P 500 option data. Common skew measures reflect the volatility, skewness, and kurtosis of the underlying distribution. The (25-delta put volatility−25-delta call volatility)/50-delta volatility measure is the main exception, as it appears not to exhibit a statistically significant dependence on the level of ATM volatility.

SKEW IN THE CROSS-SECTION: SINGLE STOCKS

In practice, comparisons of the implied volatility skew are often carried out on individual stocks. This is typically a cross-sectional comparison at a given point in time to determine which stocks have very high or very low levels of skew. While further analysis is conducted to examine specific trading opportunities or to investigate anomalies, the simple ranking of stocks by skew might be an informative first screen. The difficulty appears to be finding an appropriate measure of skew by which to rank the stocks. Thus motivated, this section examines the skew measures for single stock options.

To provide context, Exhibits 9 and 10 plot BKM skewness and kurtosis levels, respectively. Both charts show quarterly values for the S&P 500 and for constituents of the S&P 100. For the S&P 100, values for the 25th, 50th, and 75th percentiles of the cross section are plotted for each quarter.

Exhibit 11 presents the results of regressions of the various implied volatility skew measures on BKM skewness and BKM skewness interacted with the level of ATM volatility and with BKM kurtosis. The table replicates the same regressions shown in Exhibit 8, but for single stocks in the S&P 100. The data are end-of-quarter snapshots from December 2004 to December 2009 (21 quarters). Because the errors in the

EXHIBIT 9

BKM Skewness for the S&P 500 and Constituents of the S&P 100, Quarterly from December 2004 to December 2009

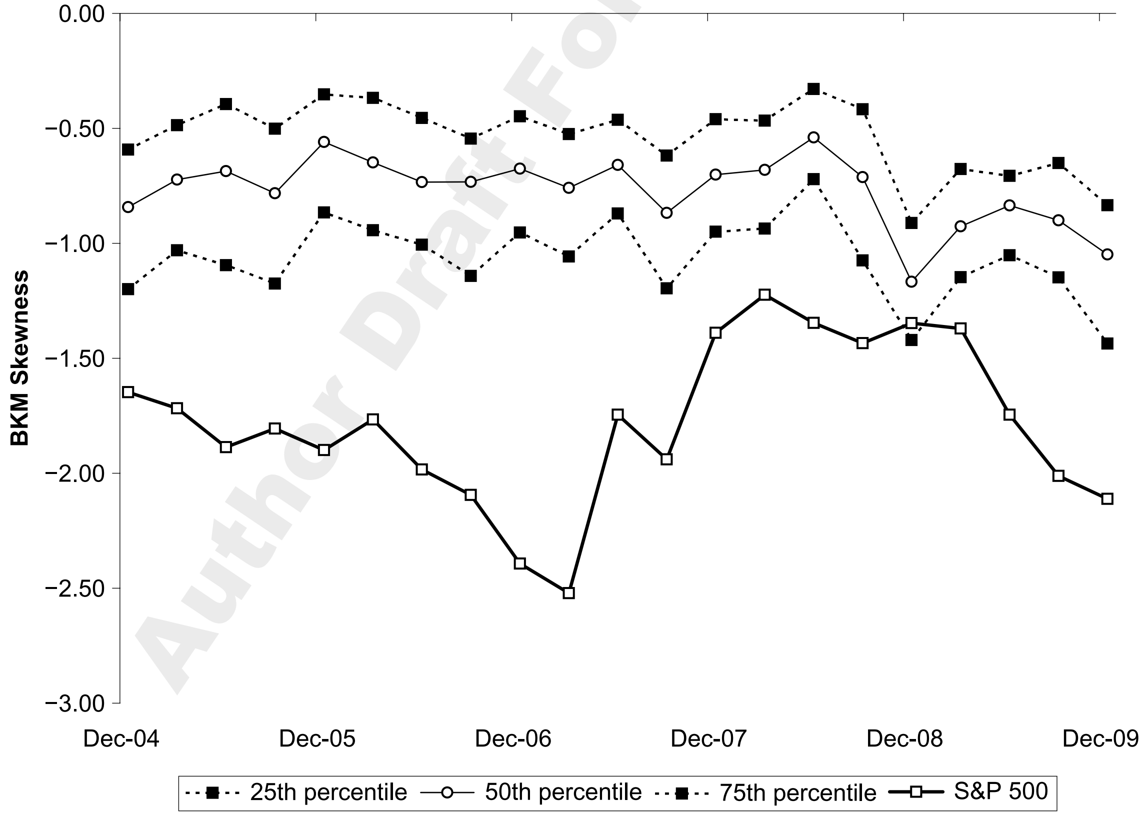

Note: For the S&P 100 constituents, the skewness values at the 25th, 50th, and 75th percentiles are plotted.

Exhibit 10

BKM Kurtosis for the S&P 500 and Constituents of the S&P 100, Quarterly from December 2004 to December 2009

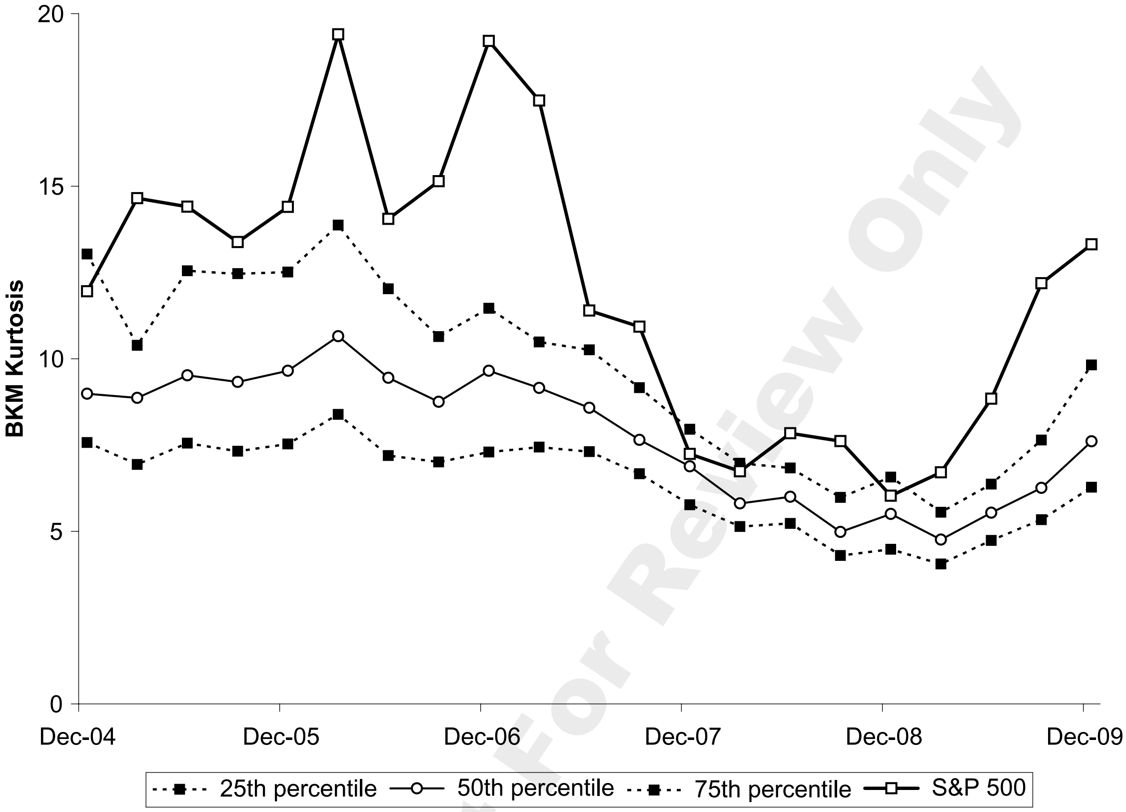

Notes: For the S&P 100 constituents, the kurtosis values at the 25th, 50th, and 75th percentiles are plotted.

regressions are quite likely to be correlated across stocks, the Fama–MacBeth [1973] two-pass regression technique is employed. Regressions are computed for each quarterly cross section of data, and the coefficients for each regression are averaged across time. The t-statistics for each coefficient are computed from the time series of estimated coefficients.

The results are qualitatively similar to those for the index option regressions. The coefficient on skewness is mostly negative, reflecting the correlation of implied volatility skewness with skew, and the coefficient on the interaction term between volatility and skewness is significant for several of the regressions. The coefficients on the interaction terms for kurtosis and skewness are significant in the regression for the (90–110)/100 measure, as well as for the regression of the (25 delta put volatility–25 delta call volatility)/50 delta volatility measure. There is often a sharp increase in the explanatory power of the regressions when the interaction terms are included, suggesting that they are indeed important to explain the variation in most of the measures. The adjusted R^2 for the 90–110 regression, for example, increases from 23.4% to 31.4% when the volatility interaction term is included, although this fit is still not as good as the fit for the regression of the (25 delta put volatility–25 delta call volatility)/50 delta volatility skew measure on skewness alone (a 33.4% R^2). As with the index option regressions, we find that the (25 delta put volatility–25 delta call volatility)/50 delta volatility skew fits the data and the theory quite well. This measure features no significant coefficient on the volatility/skewness interaction variable, and the comprehensive regression incorporating the kurtosis/skewness interaction has an R^2 just over 42%.

Despite the regression results, the significant correlation documented in the regressions between the 90–110

EXHIBIT 11

Results from Fama–MacBeth Two-Pass Regressions of Implied Volatility Skew Measures on BKM Skewness and BKM Skewness Interacted with ATM Volatility and BKM Kurtosis

Dependent Variable	Skewness	(*t*-stat)	Volatility × Skewness	(*t*-stat)	Kurtosis × Skewness	(*t*-stat)	Adjusted R^2 (%)
90% vol – 110% vol	−2.24	(−4.97)					23.4
	−0.43	(−0.86)	−6.37	(−5.89)			31.4
	−2.11	(−2.14)	−4.45	(−4.08)	0.08	(1.76)	39.6
(90% vol – 110% vol)/100% vol	−0.11	(−12.08)					30.9
	−0.19	(−12.81)	0.33	(6.75)			43.4
	−0.27	(−13.45)	0.44	(7.25)	0.00	(5.51)	51.9
95% put price – 105% call price	−0.49	(−6.33)					19.6
	−0.47	(−3.61)	0.10	(0.19)			24.3
	−0.73	(−1.29)	0.36	(0.64)	0.01	(1.73)	28.8
25d put vol – 25d call vol	−2.02	(−3.90)					13.5
	3.28	(3.50)	−17.30	(−10.59)			48.6
	2.32	(1.66)	−16.41	(−9.37)	0.05	(1.02)	51.8
(25d put vol – 25d call vol)/50d vol	−0.09	(−13.48)					33.4
	−0.08	(−7.35)	−0.01	(−0.44)			35.2
	−0.14	(−8.71)	0.05	(1.60)	0.00	(4.85)	42.2

Notes: Values in parentheses are t-statistics computed from the time series of cross-sectional regression coefficients. Cells shaded gray have t-statistics with an absolute value greater than 2. The data are end-of-quarter snapshots of S&P 100 constituents from December 2004 to December 2009.

skew and the BKM skewness might tempt one to stick with the 90–110 skew given its simplicity and common usage. Connecting the theoretical results derived earlier with the real world data can help put that inclination to rest. According to the numerical simulations, the 90–110 measure (for a constant level of skewness) is not particularly sensitive to changes in the level of volatility if volatility is above the mid-twenties, but it is very sensitive at lower levels of volatility. In the S&P 100 constituent panel data used in these regressions, I find that the proportion of ATM volatilities less than 25% significantly exceeds one-half in every cross-section of data from December 2004 to June 2007. The 90–110 skew appears inappropriate when both low and high volatility stocks are being compared.

As a final robustness check, I ran the single stock regressions including the kurtosis variable in addition to the interaction terms. The conclusions from the previous regressions remain the same, and the average regression R^2 values rise by a couple of percentage points. In none of the cases is the additional kurtosis variable deemed statistically significant at conventional levels.

The regression results have practical implications. An active research area is to use panel data to explore linkages across markets. Cremers et al. [2008], for example, study the consistency of option pricing and credit derivative pricing for individual firms. The theoretical model they construct is based on skewness, but in their empirical work they proxy skewness with 92% strike put volatility minus ATM volatility (filling in data near these target values by assuming a linear in percentage moneyness relation). They find a significant effect on credit spreads for the level of ATM volatility, but a much more modest impact for skewness. In effect, their model suggests a regression such as *credit spread* $= \alpha + \beta$ (*volatility*) $+ \gamma$(*skewness*) $+ \varepsilon$, and they regress *credit spread* $= \alpha + \beta$(*volatility*) $+ \gamma$(*skew*) $+ \varepsilon$. The regression results in this article emphasize the importance of interactions among volatility, skewness, and kurtosis in explaining the level of implied volatility skew, suggesting that they should be included in regressions proxying skewness with implied volatility skew. Omitting the interaction variables might have significant impact on the statistical inference for such regressions.

CONCLUSIONS

The motivation for this article is to find out what practitioner measures of the implied volatility skew actually measure. The investigation is both theoretical and empirical. I begin with a model describing implied volatility as a function of option moneyness, and I work out the implications relating various measures of skew to volatility and skewness of the underlying risk-neutral distribution. The second step is to test the implications of the model against index and single stock option data.

A main conclusion from the analysis is that many popular skew measures are strongly influenced by the levels of volatility and kurtosis implicit in the distribution. Using these measures to examine skew in isolation can be misleading if the analyst does not control for other changes in the distribution besides skewness. If the objective is to use a simple, robust measure of skewness that is not necessarily highly correlated with the level of volatility, the analysis in this article provides clear direction. The most descriptive and least redundant measure examined here is the (25 delta put–25 delta call)/50 delta volatility.

An objective defined in this article is to evaluate the volatility skew measures on whether they are valid skewness functionals, as defined in the statistics literature. A key takeaway from the statistics literature is that there is not a single measure of skewness that is unambiguously best for all purposes. Other evaluation techniques, for example, ones that identify specific trading opportunities, might be very informative about the usefulness of various volatility skew statistics. Furthermore the theoretical framework used in this article (linear volatility skew in delta) is a plausible and tractable reduced form model, but it lacks a deeper structural model as motivation. The model is agnostic about comovements among volatility, skewness, and kurtosis, as well as the univariate dynamics of these statistics. The approach taken in this study allows for sharp, interesting results, but it leaves obvious opportunities for further advancement.

ENDNOTES

This paper reflects the opinions of the author and does not necessarily reflect the opinions of Lyxor Asset Management Inc. The author thanks Bill Bane, Stephen Figlewski, Nikunj Kapadia, and Dan Kryzman for helpful comments on earlier drafts.

The appendix is available from the author on request.

[1]The regressions are OLS and exclude the deepest out-of-the-money options (ones with deltas less than 0.1 or greater than 0.9). Note that Bliss and Panagirtzoglu [2004] use a weighted least squares technique to downweight low-vega options, mitigating the convexity that appears in the raw data.

[2]A slightly better approximation appears to be a convex function such as $b_1(\sigma_{50\Delta})^{1.03}$, but both approximations are ad hoc and based on numerical experiments.

[3]Contract level data were obtained from iVolatility.com. End-of-day implied volatility, option delta, strike price, and soon are provided for each contract.

[4]All of the regressions in this and subsequent tables use Newey–West [1987] standard errors with 13 lags. Residuals from the regressions are highly serially correlated. For the most complicated regressions in Exhibit 7, for example, the Akaike Information Criterion selects from 3 to 18 lags when regressing residuals on lagged residuals. VARHAC covariance matrix estimation (den Haan and Levin [1998]) produces the same qualitative results.

REFERENCES

Bakshi, G., N. Kapadia, and D. Madan. "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options." *Review of Financial Studies*, Vol. 16, No. 1 (2003), pp. 101-143.

Backus, D., S. Foresi, and L. Wu. "Accounting for Biases in Black-Scholes." Unpublished paper, 2004.

Bank of England. "Recent Economic and Financial Developments." *Quarterly Bulletin*, Vol. 49, No. 1 (2009), pp. 5-26.

Bates, D. "Post-'87 Crash Fears in the S&P 500 Futures Option Market." *Journal of Econometrics*, 94 (2000), pp. 181-238.

Bliss, R., and N. Panigirtzoglu. "Option-Implied Risk Aversion Estimates." *Journal of Finance*, Vol. 59, No. 1 (2004), pp. 407-446.

Bowling, S., M. Khasawneh, S. Kaewkuekool, and B.R. Cho. "A Logistic Approximation to the Cumulative Normal Distribution." *International Journal of Industrial Engineering and Management*, Vol. 2, No. 1 (2009), pp. 114-127.

Carr, P., and L. Wu. "The Information Content of Straddles, Risk Reversals and Butterfly Spreads." Bloomberg Markets, 2005.

——. "Stochastic Skew in Currency Options." *Journal of Financial Economics*, Vol. 86, No. 1 (2007), pp. 213-247.

Cremers, M., J. Driessen, P. Maenhout, and D. Weinbaum. "Individual Stock-Option Prices and Credit Spreads." *Journal of Banking and Finance*, 32 (2008), pp. 2706-2715.

David, A., and P. Veronesi. "Macroeconomic Uncertainty and Fear Measures Extracted From Index Options." Unpublished paper, 2009.

Das, S., and R. Sundaram. "Of Smiles and Smirks: A Term-Structure Perspective." *Journal of Financial and Quantitative Analysis*, Vol. 34, No. 2 (1999), pp. 211-239.

Demeterfi, K., E. Derman, M. Kamal, and J. Zou. "A Guide to Volatility and Variance Swaps." *The Journal of Derivatives*, Vol. 6, No. 4 (1999), pp. 9-32.

Dennis, P., and S. Mayhew. "Risk Neutral Skewness: Evidence from Stock Options." *Journal of Financial and Quantitative Analysis*, Vol. 37, No. 3 (2002), pp. 471-493.

den Haan, W., and A. Levin. "Vector Autoregressive Covariance Matrix Estimation." Unpublished paper, 1998.

Doksum, K.A. "Measures of Location and Asymmetry." *Scandinavian Journal of Statistics*, Vol. 2, No. 1 (1975), pp. 11-22.

Fama, E., and J. MacBeth. "Risk, Return, and Equilibrium: Empirical Tests." *Journal of Political Economy*, Vol. 81, No. 3 (1973), pp. 607-636.

Gemmill, G. "Did Option Traders Anticipate the Crash? Evidence From Volatility Smiles in the U.K. with U.S. Comparisons." *Journal of Futures Markets*, Vol. 16, No. 8 (1996), pp. 881-897.

Groeneveld, R., and G. Meeden. "Measuring Skewness and Kurtosis." *The Statistician*, Vol. 33, No. 4 (1984), pp. 391-399.

Hull, J., I. Nelken, and A. White. "Merton's Model, Credit Risk, and Volatility Skews." *Journal of Credit Risk*, Vol. 1, No. 1 (2004), pp. 1-27.

Oja, H. "On Location, Scale, Skewness and Kurtosis of Univariate Distributions." *Scandinavian Journal of Statistics*, 8 (1981), pp. 154-168.

Toft, K., and B. Prucyk. "Options on Leveraged Equity: Theory and Empirical Tests." *Journal of Finance*, 52 (1997), pp. 1151-1180.

Mixon, S. "Option Markets and Implied Volatility: Past versus Present." *Journal of Financial Economics*, 94 (2009), pp. 171-191.

Natenberg, S. *Option Volatility and Pricing*. New York: McGraw-Hill, 1994.

Newey, W., and West, K. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55 (1987), pp. 703-708.

van Zwet, W.R. *Convex Transformation of Random Variables*. Mathematics Centre Tract 7. Amsterdam, Mathematisch Centrum, 1964.

Xing, Y., X. Zhang, and R. Zhao. "What Does the Individual Option Volatility Smirk Tell Us About Future Equity Returns?" *Journal of Financial and Quantitative Analysis*, Vol. 45, No. 3 (2010), pp. 641-662.

Appendix to
“What Does Implied Volatility Skew Measure?”
[NOT FOR PUBLICATION]

Scott Mixon

May 2010

This appendix provides supporting analysis regarding the linear-in-delta theoretical model for implied volatility.

- **Section A1** provides context by displaying several probability density functions (and associated statistics) generated with the model. [Page 2]
- **Section A2** elaborates on numerical experiments providing evidence on a theoretical property that a skewness functional should satisfy: location and scale invariance. [Page 3]
- **Section A3** elaborates on another theoretical property of valid skewness orderings: preservation of van Zwet's c-ordering of distributions. [Page 5]

A1. EXAMPLES

Chart A1 displays four probability density functions, each with 20% annualized at-the-money implied volatility. In these cases, and in the others in this appendix, the interest rate and dividend yield are set at 0.0%, and the time to maturity of the options is set at three months. The four functions differ in the level of skewness and, therefore, in the associated implied volatility skew statistics. Table A1 displays implied volatility skew statistics associated with these distributions.

Chart A1.

Probability density functions generated by the linear-in-delta skew model

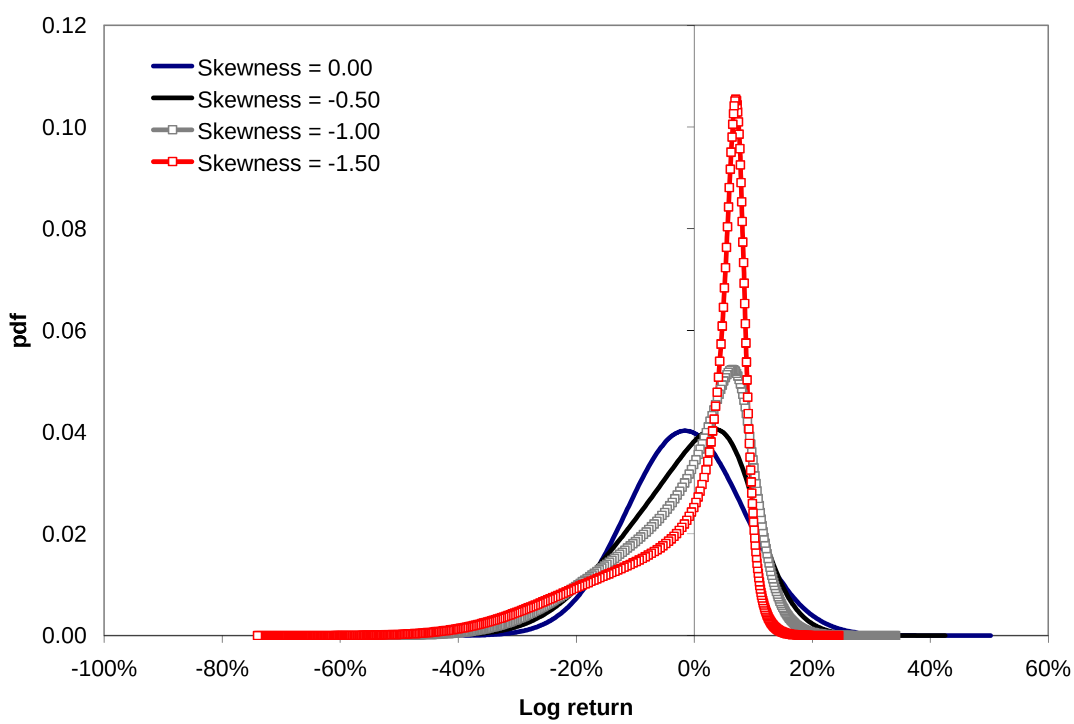

Table A1.

Parameters and statistics for examples in Chart A1

PARAMETERS		OUTPUTS					
ATM Volatility	b (Volatility slope in delta space)	Volatility of risk neutral density	Skewness of risk neutral density	90% Vol – 110% Vol	(90 – 110) / 100	25 delta put vol – 25 delta call vol	25 delta put vol – 25 delta call vol) / 50 delta vol
% per year		% per year		Volatility points	%	Volatility points	%
20.0%	0.000	20.0%	0.00	0.0	0.00	0.0	0.00
20.0%	0.061	20.4%	-0.50	4.2	0.21	3.1	0.15
20.0%	0.127	21.1%	-1.00	9.1	0.45	6.4	0.32
20.0%	0.204	22.0%	-1.50	15.7	0.77	10.2	0.51

A2. LOCATION AND SCALE INVARIANCE

Chart A2(a).

Robustness tests for Exhibit 5 (scaled)

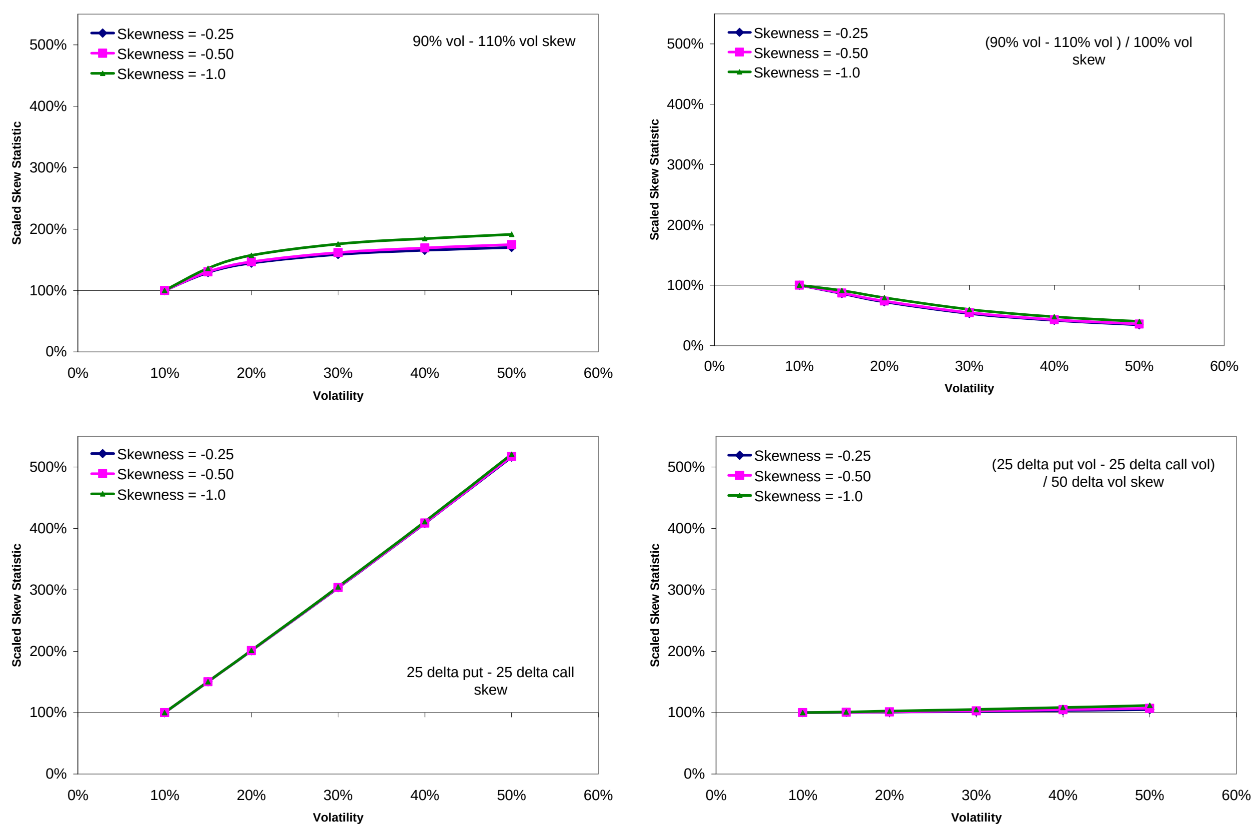

Charts A2(a) and A2(b) elaborate on the robustness tests for Exhibit 5 in the paper. The paper shows implied volatility skew statistics as a function of volatility for skewness = -0.50. Figures in Chart A2(a) display the same scaled skew statistics for skewness = -0.25, -0.50, and -1.00; figures in Chart A2(b) show unscaled values.

Chart A2(b).

Robustness tests for Exhibit 5 (not scaled)

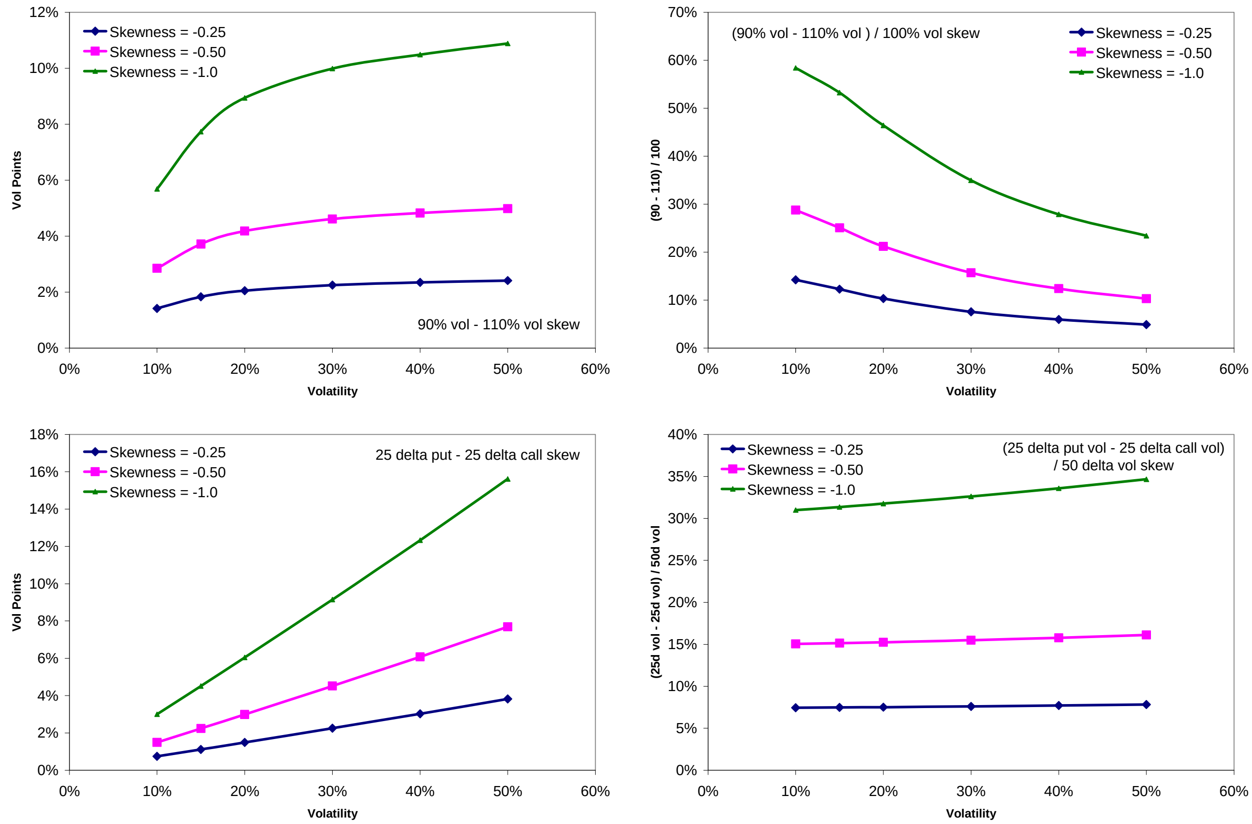

A3. PRESERVATION OF VAN ZWET’S CONVEXITY ORDERING

Table A3 displays the input parameters and volatility skew statistics for an example case in which some standard implied volatility skew measures fail to preserve van Zwet’s convexity ordering. More specifically, I compute implied volatility skew statistics for a low volatility, highly left-skewed distribution and for a high volatility, less left skewed distribution. Two of the implied volatility skew statistics (the 90 - 110 skew and the 25 delta put volatility minus 25 delta call volatility) are higher for the less left-skewed distribution, thereby sending precisely the wrong message as measures of skewness. Their significant correlation to volatility overpowers their connection to skewness in these cases.

The two panels of Chart A3 illustrate the distributions used to compute the above counterexample. The top panel shows the standardized CDFs of the two random variables. The distributions are c-comparable, with $F <_c G$, according to Oja's [1981] result regarding the crossing points of the functions. The crossing points are marked on the chart. The bottom panel shows the PDFs of the two random variables.

TABLE A3.

Example: Two measures of skew fail to preserve van Zwet's c-ordering

PARAMETERS		OUTPUTS					
ATM Volatility	b (Volatility slope in delta space)	Volatility of risk neutral density	Skewness of risk neutral density	90% Vol – 110% Vol	(90 – 110) / 100	25 delta put vol – 25 delta call vol	25 delta put vol – 25 delta call vol) / 50 delta vol
% per year		% per year		Volatility points	%	Volatility points	%
45.0%	0.312	49.9%	-1.00	10.9	0.23	15.6	0.35
12.5%	0.100	13.2%	-1.25	8.9	0.71	5.0	0.40

Table A4 provides results for another counterexample regarding van Zwet's convexity ordering. The table shows that the (90 – 110)/100 measure is lower for a highly skewed distribution than it is for a more symmetric distribution. As before, the measure's correlation to volatility overpowered its correlation to skewness. The result is a measure that can, for certain parameter configurations, give precisely the wrong suggestion regarding the level of skewness in the market.

The panels of Chart A4 mimic those in Chart A3. The crossing points of the two CDFs are marked in order to show consistency with Oja's [1981] results showing c-precedence. The relation $F <_c G$ holds for the two distributions, but the measure (90 – 110)/100 does not respect that relation.

TABLE A4.

Example: Another measure of skew failing to preserve van Zwet's c-ordering

PARAMETERS		OUTPUTS					
ATM Volatility	b (Volatility slope in delta space)	Volatility of risk neutral density	Skewness of risk neutral density	90% Vol – 110% Vol	(90 – 110) / 100	25 delta put vol – 25 delta call vol	25 delta put vol – 25 delta call vol) / 50 delta vol
% per year		% per year		Volatility points	%	Volatility points	%
15.0%	0.094	15.6%	-1.00	7.9	0.52	4.7	0.31
30.0%	0.256	33.1%	-1.25	13.6	0.44	12.8	0.43

Chart A3. Two standardized distributions, *F(x)* more skewed to left than *G(x)*

Panel A: CDFs

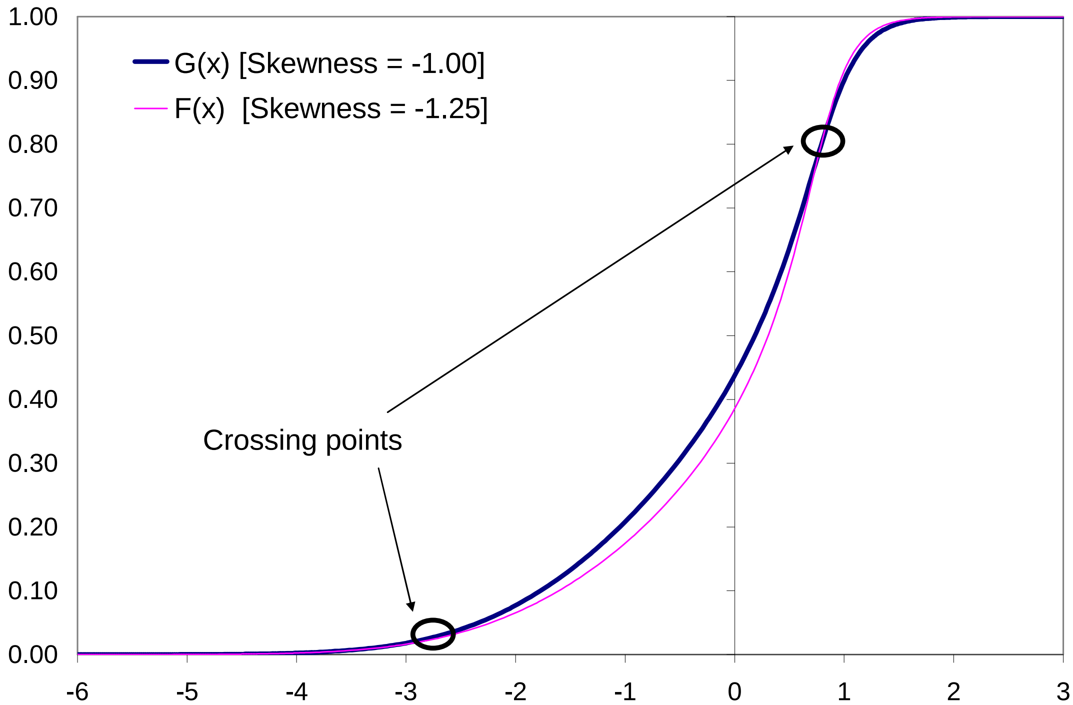

Panel B: PDFs

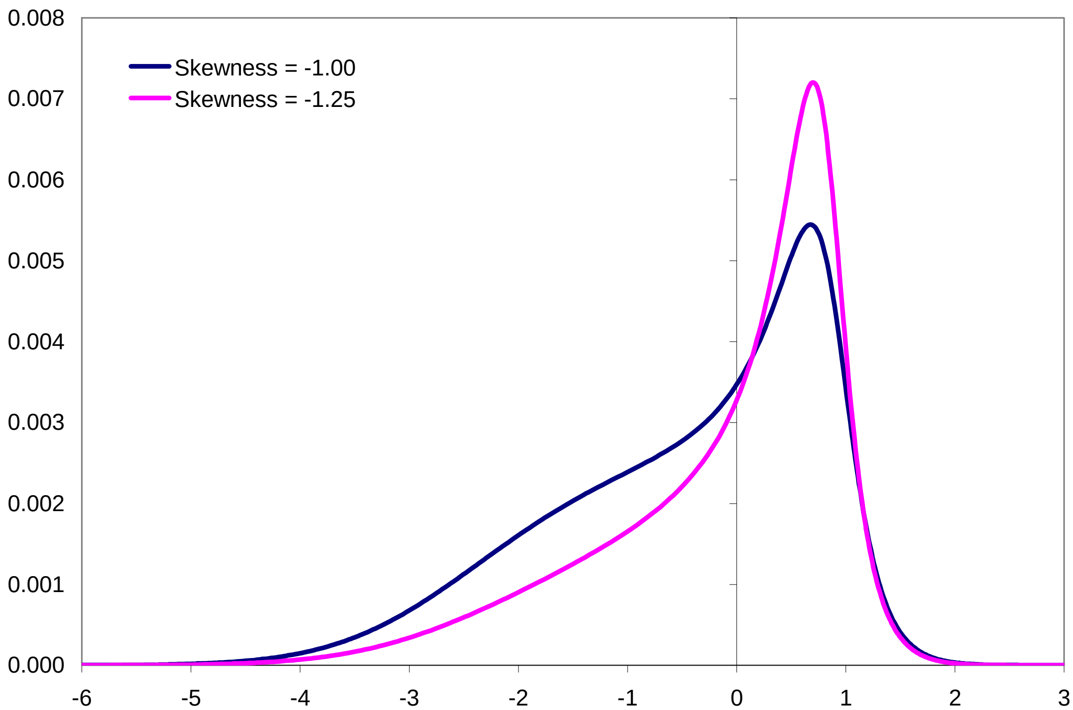

Chart A4. Two standardized distributions, *F(x)* more skewed to left than *G(x)*

Panel A: CDFs

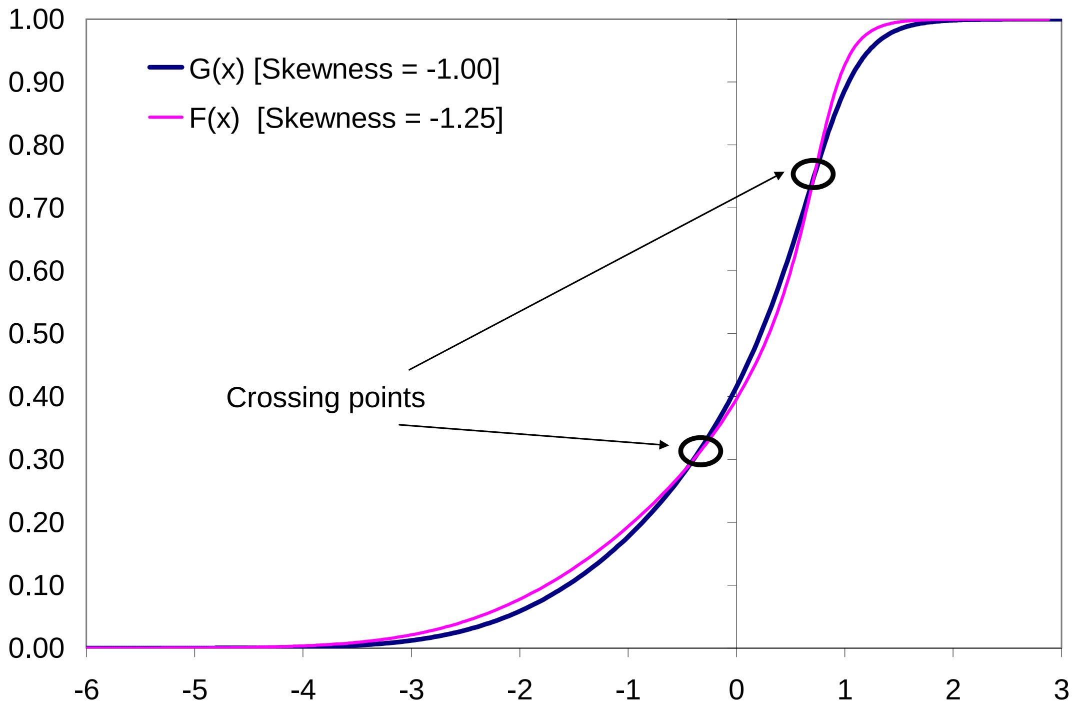

Panel B: PDFs