

BOUNDS OF PROBABILITY

Option prices can be used to determine the probability that future index values will lie in a given range, enabling the user to design better speculative strategies and calculate better hedge ratios.

David Shimko explains how

The now classic Black-Scholes (1973) option pricing model has been applied successfully to many option and derivative markets. The original model assumes that stock volatility is constant; if the model were precisely true, then implied volatility from European options of all strikes and maturities would be identical. Many traders claim to use the Black-Scholes model for pricing purposes. What they in fact use is a variant interpretation that allows the implied volatility of the underlying (the index) to vary according to the option's exercise price and maturity.

The volatility structure (the plot of implied volatility against the strike price) has graduated from a pricing convenience to a risk management tool. If the volatility structure shifts up or down in parallel fashion, for example, it is easy to determine the impact on the value of any option portfolio. The kappa (or vega) indicates how much an option's value changes for a small movement in implied volatility. Traders also use this information to trade volatility – for example, when they take a position that profits if their estimates of future implied volatility are correct.

Using the Black-Scholes formula, it is possible to calculate not only implied volatility, but any parameter of the model. For example, two market option prices can be used to determine the implied volatility and implied index value simultaneously. This procedure would be especially appropriate if the value of the index were unknown or unreliable. For example, in trading the S&P 100 option contract (OEX), the value of the underlying S&P 100 stock index is unreliable because not all stocks trade at the same time; the reported index changes lag real value changes by a small margin.¹

Therefore, if it is possible to calculate the

implied mean (from the implied index value) and the implied standard deviation of future index levels from two option prices; other characteristics of the distribution can be calculated with more than two option prices. Or, in reverse, the mean and standard deviation could be found by knowing the entire probability distribution² associated with different future index movements. The mean of this distribution is proportional to the expected ex-dividend index level, and the standard deviation of this distribution can be used to find the implied volatility. But the entire probability distribution gives more information than just the mean and standard deviation; it can provide a measure of the skewness of a distribution and its kurtosis.

For example, the market may place relatively greater probability on a downward price movement than an upward movement: this is known as a negative skew. When traders speak of "trading the skew", they are usually referring to predicting the slope of the implied volatility curve, and choosing an option position that profits if their forecast materialises. The use of the word "skew" is statistically correct, since a negatively sloped implied volatility curve gives rise to a negatively skewed probability distribution for future index levels. The skewness implied by the Black-Scholes model is uniformly small and positive.

Going beyond skewness, the market may believe that extreme upward and downward movements are more likely than allowed by the Black-Scholes assumptions: in this case, it is said that the implied market distribution is more *leptokurtic* than implied by Black-Scholes. This can be seen when the implied volatility curve "smiles" – is convex in the exercise price. In extreme cases, the smile creates a two-humped probability distribution. Using our methodology, we were able to observe several instances of two-humped distributions for the OEX in the late 1980s; they have since become more rare. The Black-Scholes model assumes a (one-humped) lognormal distribution for future index levels, usually associated with a kurtosis coefficient slightly higher than 3. To understand the importance of kurtosis, a trading example may be worthwhile: traders who consis-

The volatility structure has graduated from a pricing convenience to a risk management tool

tently sell strangles are implicitly betting that the market has overestimated the kurtosis of the true probability distribution.

This line of argument leads to two observations. First, it is desirable to know the entire probability distribution of future index values for both speculative and risk management purposes. Second, since many options trade simultaneously, it should be possible to generate these probability distributions from a large enough set of option prices. The fact that probability distributions can be recovered from option prices was first discovered by Breeden and Litzenberger.³ Also, the implied probability distribution has been calculated to test hypotheses regarding the behaviour of its mean.⁴ Other papers have used a parametric approach to determine the implied distribution. Our approach is to demonstrate how the original Breeden-Litzenberger approach can be implemented nonparametrically to recover implied probability distributions of future index values.

The proof of the following analysis can be found in the box on page 36. In the Breeden and Litzenberger approach, the second derivative of the call option price with respect to the exercise price of the option must be calculated to find the probability distribution. (This can be approximated by pricing butterfly spreads as well.) To calculate derivatives, smooth call option pricing functions are needed – for example, call option prices for an infinite number of exercise prices.

To get smooth functions (since options

¹ Option traders can overcome this problem by pricing the cash options off the index futures. This solves the asynchronous trading problem of stocks in the index, but the futures prices have to be adjusted for time value. If the discount rate used for this adjustment is incorrect (it is not generally equal to the Treasury bill rate), then the implied underlying cash price is also incorrect.

² This article considers only *certainty-equivalent distributions*: the index probabilities associated with the risk-neutral pricing model. In this model, either all investors are risk-neutral, or perfect arbitrage is always possible.

³ Breeden, Douglas and Robert Litzenberger, 1978, *Prices of state-contingent claims implicit in option prices*, *Journal of Business* 51, pages 621–652.

⁴ See Longstaff, Francis, 1993, *Martingale restriction tests of option pricing models*, University of California, Los Angeles, working paper.

do not trade at *all* exercise prices), one must interpolate. The results of the analysis are very sensitive to the choice of interpolation method; the one presented here provided the smoothest-appearing results among those tried. Call option prices can be interpolated directly. However, the call option prices can be made smoother in general by interpolating implied volatility.⁵ For example, if five options give five different implied volatility figures, the implied volatility could be linearly interpolated to get an option price for every single striking price. We elected an even smoother "interpolation" by allowing implied volatility to be represented as a best-fit least-squares parabola. That is, we allowed the implied volatility function to be a parabolic function of the exercise price, $A_0 + A_1X + A_2X^2$, for every exercise price within the range of traded exercise prices. Outside this range, we assume implied volatility is constant. To make the analysis concrete, the following example is offered.

Consider European options written on the S&P 500 index. For example, on October 21, 1991, *The Wall Street Journal* reported these prices for calls and puts with a December expiration:

Exercise	Call	Put
325	66.500	0.3125
345	46.000	0.8750
360	33.000	2.0000
365	27.750	2.6250
375	20.125	4.2500
385	13.500	7.1250
390	9.625	8.7500
395	7.250	11.0000
400	5.375	13.7500
405	3.375	17.0000
410	1.875	19.7500
425	0.250	34.0000

Only option pairs for which both call and put prices are listed are used.

First, we wish to find the implied index level and risk-free interest rate. Theoretically, these data are available from the marketplace. However, we have already mentioned that index and futures values are unreliable estimators of the true index; we also argue that the T-bill rate may not be the appropriate risk-free dis-

count rate.⁶ We therefore chose to find the implied index level and interest rate. This is quite simple for European options. Put-call parity requires that the following equation holds:

$$C - P = SD - XB$$

where D is a dividend discount factor, and B is a risk-free discount factor. A simple linear regression of (call - put) on the exercise price yields the following: the intercept is 387.9812 and the slope is -0.991646, with $R^2 = 99.97\%$. The intercept is an estimate of the ex-dividend index value (SD), and the slope is an estimate of minus the T-bill price per \$1 face value (-B). The T-bill matures on the same day as the option. The closing index value was reported at 390.02, implying a continuous proportional dividend yield over the two-month period of 3.14% (annualised). The dividend yield falls within the historical range. The T-bill price estimate gives an implied risk-free rate of 5.03% (continuously compounded) over the two-month period. On October 21, the quoted T-bill for December 19 maturity has an ask yield of 5.04%.⁷

At this point, like many option traders, we create a chart of implied volatilities across strike prices. The implied SD and B values are used to calculate implied volatilities for the options listed above. The generalised Black-Scholes equation (A2) is then used together with market call prices to find the implied volatilities (v).⁸ The implied volatilities can be written as $v(X)$, since they usually differ for every X. We smooth the implied volatilities according to a quadratic relation (equation A3) with the exercise price. The regression estimates for our example imply that

$$v(X) =$$

$0.89167 - 0.0038715X + 0.0000044458X^2$ with $R^2 = 75.20\%$. The rough and smooth implied volatilities are shown in figure 1.

In this case, the implied volatility profile is negatively sloped, implying a negative skew to the distribution of future index values: the market is placing relatively greater weight on a fall in the index than on a rise. The slight convexity indicates that the market is also placing greater weight on large movements (positive or negative) than would be implied by the Black-Scholes model.

The smoothed volatilities can be used to find the smoothed call prices (via Black-Scholes), and the smoothed call prices can be differentiated to find the values of the density function and cumulative distribution function for each possible value of X. The distribution is only calculated between the endpoints from the smoothed volatility values. Equations A7 and A8 provide analytic expressions for the density and cumulative distribution values under the parabolic volatility structure assumption.

This leaves the technical problem: what should we do with probabilities beyond the range of traded strike prices? We chose to assume the tail distributions were log-normal.

⁵ Although we use Black-Scholes implied volatility for the purposes of interpolation, we do not assume that the Black-Scholes model holds


⁶ See Longstaff (1993)

⁷ These numbers do not always conform to market data so closely

⁸ For the example, we smoothed values of $v = \sigma\sqrt{t}$. This simplification obviates the need to know the time to maturity of the options

1. Implied volatility profile

December maturity, 21 October 1991



RECOVERING INDEX PROBABILITY DISTRIBUTION FROM OPTION PRICES

Use prices of European calls and puts for a set of different exercise prices (X), and the same time to maturity (τ). Put-call parity can be established for any X and τ :

$$C(X, \tau) - P(X, \tau) = SD(\tau) - XB(\tau) \quad (A1)$$

The current level of the stock price (or futures price) is S , and $D(\tau) \leq 1$ represents a time-dependent dividend adjustment factor. The product of S and $D(\tau)$ represents the present value of the expected future ex-dividend price of the stock. $B(\tau) \leq 1$ represents the value of a zero-coupon bond per dollar of face value; the bond matures on the same date as the options. For every maturity τ observed, regress the option price difference ($C(X, \tau) - P(X, \tau)$) on a constant and the exercise price. The constant is an estimate of $SD(\tau)$, and the negative of the slope is an estimate of $B(\tau)$. These estimates are termed S' and B' .

We ultimately wish to interpolate $C(X, \tau)$ for unobserved X and fixed τ . Begin by calculating the Black-Scholes implied volatility for each call option. To calculate implied volatilities, we find $\sigma(X)$ for each X_1 to satisfy the modified Black-Scholes equation below:

$$C(X, \tau) = S'N(d_1) - XB'N(d_2) \quad (A2)$$

$$d_1 = \left[\ln\left\{ \frac{S'}{XB'} \right\} + \frac{1}{2}\sigma^2\tau \right] / \left[\sigma\sqrt{\tau} \right]$$

$$d_2 = d_1 - \sigma\sqrt{\tau}$$

$N(\bullet)$ is the cumulative normal distribution function. The known values of S' and B' (from the previous regression) are used to calculate the implied volatility function, $\hat{\sigma}(X, \tau)$, hereafter called the volatility structure.

Implied volatilities are then "smoothed".¹ The volatility structure is represented as a parabola of best least-squares fit:

$$\hat{\sigma}(X, \tau) = A_0(\tau) + A_1(\tau)X + A_2(\tau)X^2 \quad (A3)$$

The smoothed volatility structure gives a value of σ for every X .² These values of σ are used to generate smooth call option prices, using the Black-Scholes equation.

Following Cox, Ross and Rubinstein,³ we write the call option price for fixed τ and arbitrary X as:

$$C(X, \tau) = B(\tau) \int_X^{\infty} (S^* - X) f(S^*, \tau) dS^* \quad (A4)$$

S^* is the (random) value of the certainty-equivalent ex-dividend stock price; $f(S^*, \tau)$ is the probability density, and $F(S^*, \tau)$ is the cumulative probability density. The assumed pricing relationship is more general than Black-Scholes.

Breeden and Litzenberger demonstrated that the partial derivatives with respect to the exercise prices of the options are related to the distribution function $F(S^*)$ and the density function $f(S^*)$ in the following manner:

$$C_x(X, \tau) = -B(\tau) [1 - F(S^*, \tau | S^* = X)] \quad (A5)$$

$$C_{xx}(X, \tau) = B(\tau) f(S^*, \tau | S^* = X) \quad (A6)$$

The results can be verified with Leibniz's rule for differentiation under an integral. Under the parabolic implied volatility structure assumption, we can calculate the appropriate derivatives of the call pricing function as a function of $\sigma(X)$, together with (A5) and (A6) to find the frequency and cumulative frequency values between the endpoints. This procedure gives an analytic expression for the density functions under the parabolic implied volatility assumption. Other interpolation techniques may be used; in all cases, analytic expressions can be calculated for the implied distributions. For the parabolic volatility structure, the distribution is described as follows:⁴

$$f(S|S=X) = n(d_2) [d_{2x} - (A_1 + 2A_2X)(1 - d_{2x}) - 2A_2X] \quad (A7)$$

$$F(S|S=X) = 1 + Xn(d_2)(A_1 + 2A_2X) - N(d_2) \quad (A8)$$

$$d_{1x} = -1/(Xv) + (1 - d_1/v)(A_1 + 2A_2X)$$

$$d_{2x} = d_{1x} - (A_1 + 2A_2X)$$

$$v = \sigma\sqrt{\tau}$$

The implied distribution above can be used to calculate numerically the implied moments of the distribution. For example, the mean of the CEQ distribution (p_1) should equal $SD(\tau)/B(\tau)$.⁵ The implied bond price can be inverted to find the implied marginal cost of riskless capital (r) in the option market. The variance of the CEQ distribution (p_2), translated to return form,⁶ gives a unique measure of instantaneous implied volatility, σ' . This calculation obviates the need to calculate weighted average implied volatilities to determine a single working volatility figure.

We also calculate the skewness (p_3) and kurtosis (p_4) of the implied CEQ distribution. The summary of these parameter calculations is shown below:

$$p_1 = E[S^*] \quad (A9)$$

$$p_2 = E[(S^* - p_1)^2]$$

$$p_3 = E[(S^* - p_1)^3] / p_2^{3/2}$$

$$p_4 = E[(S^* - p_1)^4] / p_2^2$$

Let

$$q = \sqrt{p_2} / p_1,$$

the coefficient of variation for the CEQ distribution. The unambiguous implied return volatility for the CEQ distribution is given by

$$\sigma' = \sqrt{\ln(q^2 + 1) / \tau} \quad (A10)$$

If the current true index level (S) is known, then the implied dividend discount factor can be determined by:

$$D' = p_1 B' / S \quad (A11)$$

The implied continuous proportional dividend yield of the index, σ , over the life of the option is given by:

$$\delta' = -\ln(D') / \tau \quad (A12)$$

A benchmark lognormal distribution with the same mean and variance of the implied CEQ distribution has the following higher moments:⁷

$$I_3 = 3q + q^3 \quad (A13)$$

$$I_4 = 3 + 16q^2 + 15q^4 + 6q^6 + q^8$$

I_3 represents the skewness of the lognormal distribution, and I_4 its kurtosis. These higher moments can be used to compare the implied distribution to the lognormal distribution with the same mean and variance. ■

¹ At first, we tried interpolating volatility in a piecewise-linear fashion. This procedure leads to choppy implied CEQ distributions. However, it retains the desirable property that when call prices are recalculated, market prices of observed calls are recovered.

² Note that Black-Scholes pricing holds if $A_1(\tau) = A_1$ for all τ , and $A_2 = 0$. These restrictions could be used to test the Black-Scholes model.

³ Cox, J., S. Ross and M. Rubinstein, 1979, *Option pricing: A simplified approach*, *Journal of Financial Economics* 7, pages 229-264.

⁴ For any interpolation technique, the derivatives may be calculated numerically as well. In this case, the simple functional form of the volatility profile allows us to calculate the appropriate derivatives analytically.

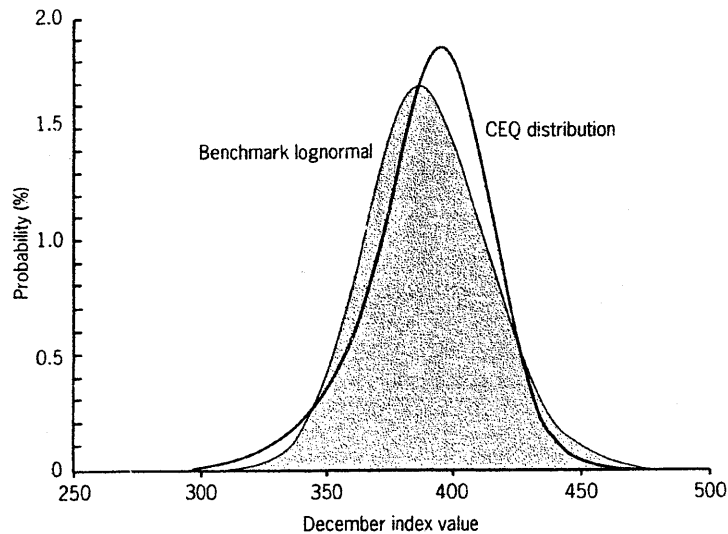
⁵ The equality of SD/B to the mean of the CEQ distribution has been questioned on empirical grounds by Longstaff (1993). If certainty equivalent pricing does not hold, we must use a discount rate different from the risk-free rate; our techniques are flexible enough to handle this possibility.

⁶ That is, the annualised volatility of $\ln(S^*)$.

⁷ See Jarrow, R. and A. Rudd, 1982, *Approximate option valuation for arbitrary stochastic processes*, *Journal of Financial Economics* 10, pages 347-369.

2. Implied CEQ index distribution

December maturity, 21 October 1991



We matched the frequency and cumulative frequency of the distribution with a lognormal distribution in each tail. At the minimum X value (325) and the maximum X value (425), we calculated the density and distribution values, and searched for lognormal distributions that had the same values.

The resulting distribution is pieced together from its three parts. The distribution is shown, along with the benchmark lognormal distribution, in figure 2.

The lognormal distribution is the shorter in maximum height. The negative skew in the actual distribution is apparent from figure 2. Traders may be satisfied with examining the diagram, and making speculative plays based on guesses about changes in probabilities. Or, the distribution can be used to calculate numerically the implied moments of the future index distribution. For example, the mean is equal to the sum of the products of the possible index values multiplied by the respective probabilities. The moments are tabulated as follows:

Moment no	Description	CEQ moment	Lognormal moment
0	Area under curve	1.046	1.000
1	Mean of distribution	388.915	388.915
2	Variance	573.090	573.090
3	Skewness	-0.511	0.185
4	Kurtosis	3.465	3.061

The higher moments (above 0) are corrected by normalising the area to one. We matched the actual distribution to a lognormal distribution with the same mean and variance. The variance of the index level can also be converted to a return volatility (see page 36); the resulting

implied annual return volatility is 15.06% or 6.15% for the two-month period. This figure is in the range of the Black-Scholes implied volatilities. Notably, this calculation yields a single volatility figure – researchers have in the past used at-the-money implied volatility or averages of implied volatilities as proxies for true volatility.

If we examine a benchmark lognormal distribution with the same mean and variance as the implied index distribution, we find a lognormal skewness coefficient of 0.185, and kurtosis of 3.061. The implied index distribution is negatively skewed and more leptokurtic than the benchmark lognormal distribution. This numerically confirms our earlier observation from the implied volatility curve regarding the skewness and kurtosis.

The remaining question is: what is the value of this extra knowledge of index probabilities? First, we resolve the trader's paradox that the volatility of the index implied by different options varies; with a unique picture of the probability distribution, there is a unique measure of volatility. There is also a unique measure of the mean of the distribution which does not depend on the cash or futures index prices. Second,

The method makes it possible to analyse apparently rational and irrational probability movements and compare them across markets

this distribution can be used to test the suitability of Black-Scholes for pricing in a given option market. Our data indicate that, while Black-Scholes performs relatively well for index options, it does not perform as well for, say, crude oil futures options. Third, we suggest that these moments (mean, standard deviation, skewness, and kurtosis) may be tracked through time to create better forecasts and risk management capabilities. For example, the means and standard deviations of these distributions tend to be negatively correlated through time. This will cause Black-Scholes call deltas to overestimate true market deltas, and lead to hedging underperformance for those who use Black-Scholes deltas to hedge. Fourth, we argue that this method will complete the trader's knowledge set by making it possible to analyse apparently rational and irrational probability movements and compare them across markets. This technique does not provide a trading system, but it provides a new way to examine the information implicit in option prices.

In short, if we all agree that implied volatility is an important informational, speculative and hedging tool, should we not say the same of implied skewness, kurtosis, modality ... indeed, the probability that the index will fall in any given range? ■

David Shimko is assistant professor of finance at the department of finance, University of Southern California, Los Angeles. The author thanks Nikko-LOR for financial support and Mark Rubinstein for comments on the academic draft of this paper