Edgeworth Binomial Trees[†]

Mark Rubinstein
Paul Stephens Professor of Applied Investment Analysis
University of California, Berkeley

a version published in the *Journal of Derivatives* (Spring 1998)

Abstract

This paper develops a simple technique for valuing European and American derivatives with underlying asset risk-neutral returns that depart from lognormal in terms of prespecified non-zero skewness and greater-than-three kurtosis. Instead of specifying the entire risk-neutral distribution by the riskless return and volatility (as in the Black-Scholes case), this distribution is specified by its third and fourth central moments as well.

An Edgeworth expansion is used to transform a standard binomial density into a unimodal standardized discrete density – evaluated at equally-spaced points – with approximately the prespecified skewness and kurtosis. This density is in turn adjusted to have a mean equal to the riskless return (adjusted for the payout return, if any) and to a prespecified volatility.

European derivatives are then easily valued by using this risk-neutral density to weight their possible payoffs. European options with earlier maturities, American and exotic options can be valued in a consistent manner by using the method of implied binomial trees. These trees are particularly well-suited for this since they are generated from arbitrary discrete expiration-date risk-neutral probabilities -- precisely what is provided by the Edgeworth expansion.

The paper ends by translating several examples of alternative risk-neutral distributions into option prices and then into Black-Scholes implied volatility smiles. Implied trees are used to determine smiles for otherwise identical shorter-maturing options and future smiles for the original options conditional on knowing the future underlying asset price.

[†] The author thanks Jens Jackwerth for helpful comments.

Edgeworth Binomial Trees

The **risk-neutral valuation principle** as applied to European derivatives says that the present value of the derivative is its expected expiration-date payoff discounted to the present at the riskless return, where the expectation is calculated using **risk-neutral probabilities**. In addition, these same risk-neutral probabilities also apply to the underlying asset itself, so that its expiration-date expected value (based on these probabilities), after accounting for payouts, discounted by the riskless return is also its current price.

This principle has become the generator of modern derivatives pricing theory, with versions applying to American and exotic derivatives as well. While most derivatives pricing models share this principle in common, perhaps the most significant difference among them lies in their specification of the higher order moments (above the second) of the risk-neutral distribution. For example, the Black-Scholes formula assumes that the standardized risk-neutral distribution of the logarithm of underlying asset returns is normal with skewness 0 and kurtosis 3. This observation was not lost on Jarrow and Rudd[1] who developed a way of approximating the present value of a derivative by using an Edgeworth expansion allowing different values of skewness and kurtosis.

This paper simplifies their approach by applying an Edgeworth expansion directly to discretized risk-neutral probabilities (rather than to the derivatives pricing formula). In particular, the paper discretizes the risk-neutral distribution of the logarithm of underlying asset returns using equally-spaced points and provides a simple method of attaching risk-neutral probabilities to these points to match the first four central moments of a prespecified risk-neutral distribution. The paper then shows how the shape of the current **implied volatility smile** of European options depends on the assumed skewness and kurtosis of the underlying asset return.

This would be sufficient to value a European derivative. To value American-style derivatives, or several types of exotic derivatives, it is also necessary to know aspects of the risk-neutral stochastic process of the underlying asset return which explains how it can change from its current price to its price at the payoff of the derivative. This paper also takes the valuation problem this further step by using **implied binomial trees** as developed by Rubinstein.[2] This is a natural approach since these trees are expressly designed to accommodate arbitrary discretized expiration-date risk-neutral distributions. In particular, it compares the shapes of the implied volatility smile for European options of different maturities derived from the same implied tree.

I. Edgeworth Densities

We begin with a standardized binomial density b(x). This can be interpreted as a way of discretizing a continuous density function by concentrating it at equally-spaced discrete points. If n+1 is the number of points, then at each point j = 0, …, n the random variable x equals $[(2j) - n]/\sqrt{n}$ with associated probability b(x) equal to $[n!/j!(n-j)!]\,(½)^n$. For example, if n = 1, then x equals -1 or 1 with equal probability (½). If n = 4, then x equals -2, -1, 0, 1 or 2 with associated probabilities b(x) of 1/16, 1/4, 3/8, 1/4 and 1/16. Notice that this distribution has a mean of zero and a variance of one.

---

[1] R. Jarrow and A. Rudd, "Approximate Option Valuation for Arbitrary Stochastic Processes," *Journal of Financial Economics* 10, No. 3 (November 1982), pp. 347-369.

[2] M. Rubinstein, "Implied Binomial Trees," *Journal of Finance* 49, No. 3 (July 1994), pp. 771-818.

Given prespecified skewness ξ and kurtosis κ, Exhibit 1 describes how to transform this standardized binomial density b(x) into f(x), an approximately standardized density with approximately the desired skewness and kurtosis. Since this uses the Edgeworth expansion described by Johnson, Kotz and Balakrishnan,[3] we loosely refer to f(x) as a standardized "Edgeworth density."

Exhibit 1

Edgeworth Binomial Trees

C=f(S,t)

Edgeworth Expansion

b(x) = standardized binomial density

$$E(x) = 0, \quad E(x^2) = 1, \quad E(x^3) = 0, \quad E(x^4) = 3$$

f(x) = standardized "Edgeworth density"

$$E(x) \cong 0, \quad E(x^2) \cong 1, \quad E(x^3) \cong \xi, \quad E(x^4) \cong \kappa$$

$$f(x) = \left[1 + \frac{1}{6}\xi(x^3 - 3x) + \frac{1}{24}(\kappa - 3)(x^4 - 6x^2 + 3) + \frac{1}{72}\xi^2(x^6 - 15x^4 + 45x^2 - 15) \right] b(x)$$

The transformation works by multiplying the density b(x) by

$$1 + (1/6)\xi(x^3 - 3x) + (1/24)(\kappa - 3)(x^4 - 6x^2 + 3) + (1/72)(\xi^2)(x^6 - 15x^4 + 45x^2 - 15)$$

Notice that if skewness $\xi = 0$ and kurtosis $\kappa = 3$, then this expression equals 1 so that f(x) = b(x).

Unfortunately, this expansion is only an approximation. Generally, $\Sigma_j f(x_j) \neq 1$ and the moments are slightly in error. To correct for this, after the expansion, we rescale the probabilities so that they sum to 1 by replacing $f(x_j)$ with $f(x_j)/\Sigma_j f(x_j)$. Then, using this rescaled density, calculate its mean $M \equiv \Sigma_j f(x_j)x_j$ and its variance around that mean $V^2 \equiv \Sigma_j f(x_j)(x_j - M)^2$. Finally, replace the x_j with the standardized zero mean, standard deviation one random variable $(x_j - M)/V$. Notice that this continues to leave the variables x_j equally-spaced since $x_{j+1} - x_j = x_j - x_{j-1}$ for j = 1, …, n-1.

While this modification means that the probabilities sum to one and deliver the desired mean and standard deviation, the resulting skewness and kurtosis will only approximate the target levels (ξ ,

[3] N.L. Johnson, S. Kotz and N. Balakrishnan, *Continuous Univariate Distributions*, Volume 1, 2nd edition (John Wiley & Sons, 1994).

κ). Fortunately, in practice, these target levels are closely approximated, and the approximation tends to be improved by choosing n larger so that b(x) more closely resembles a normal distribution.

Another problem that can arise is that the original function f(x) may not be non-negative everywhere and so will not qualify as a density. However, there is a significant range of the target moments (ξ , κ) for which f(x) is positive. Even if $f(x) > 0$ for all x, it may not be unimodal. But again, the range of unimodal distributions is significant.[4] Exhibit 2 describes the locus of skewness, kurtosis pairs for which f(x) is positive and unimodal. These permit the use of density functions which are significantly leptokurtic and positively or negatively skewed.

Exhibit 2

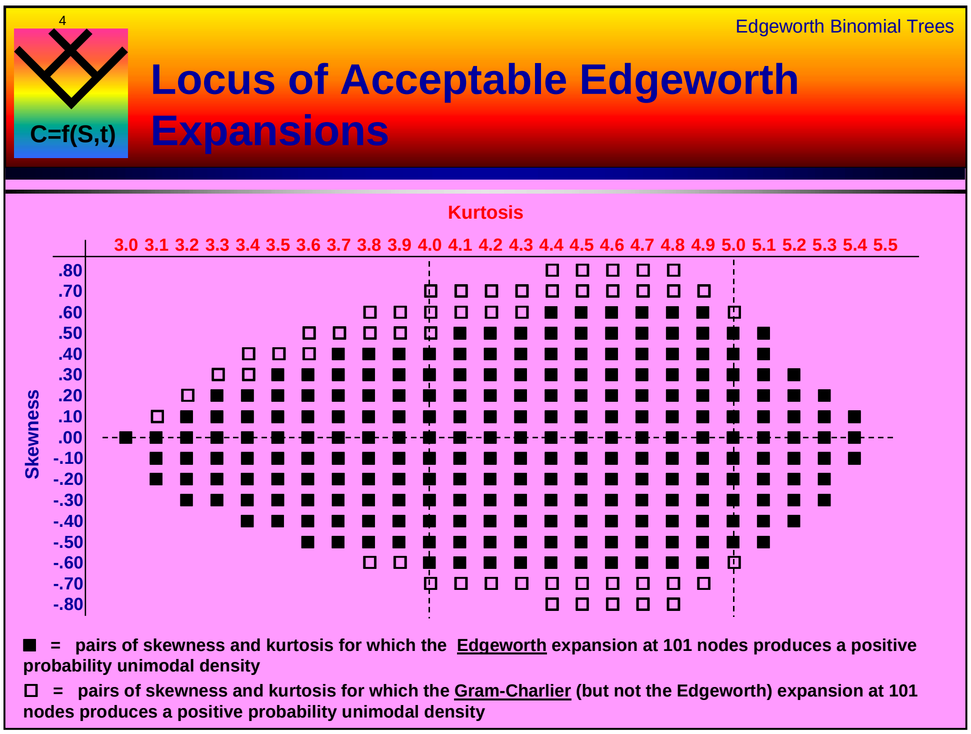

This locus can be enlarged as indicated above by using the Gram-Charlier expansion when the Edgeworth expansion fails to deliver these properties. This expansion is the same as the Edgeworth except the last term of that expansion is omitted. In other words, to use the Gram-Charlier expansion, multiply b(x) by:

$$1 + (1/6)\xi(x^3 - 3x) + (1/24)(\kappa - 3)(x^4 - 6x^2 + 3)$$

The locus of acceptable skewness, kurtosis pairs is also influenced by the number of discrete points n+1 at which the probabilities are concentrated. The more such points, the smaller the locus. For the purpose of valuing derivatives, n = 100 usually delivers more than the required accuracy, so the attached picture is based on this level.

[4] Note that the rescaling and standardizing changes in f(x) have no effect on the sign of f(x) or its modality.

Exhibit 3 gives us some idea of the flexibility of the Edgeworth approach in designing density functions.

Exhibit 3

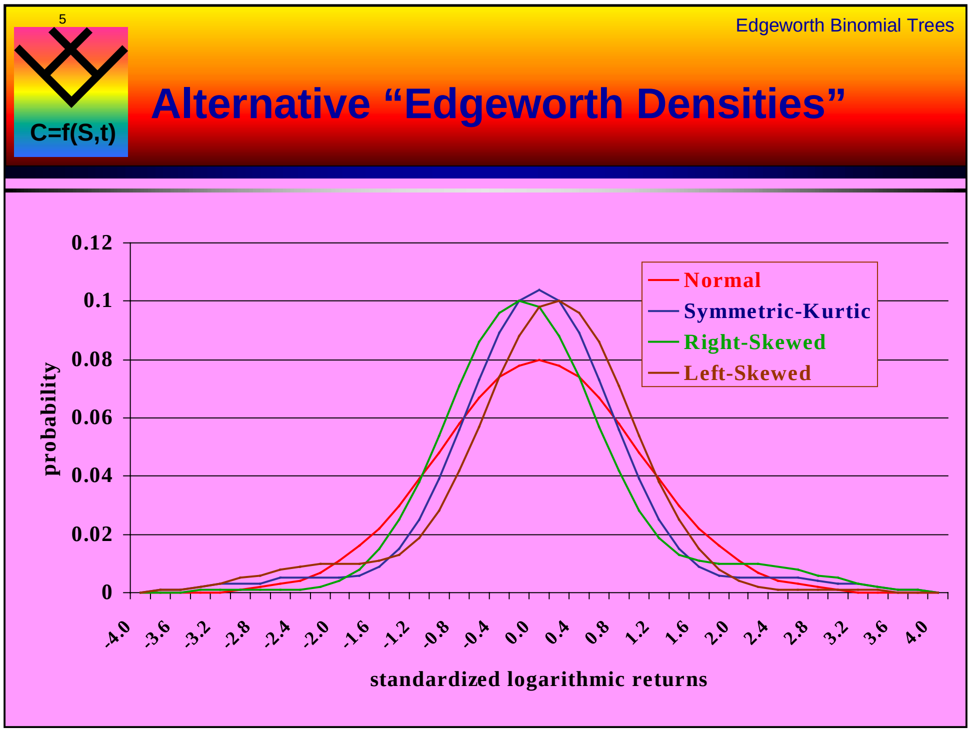

Listed below are the resulting moments in each of the four examples:

	E(x)	E(x^2)	E(x^3)	E(x^4)
Normal (ξ = 0, κ = 3)	0	1	0	2.98
Symmetric-Kurtic (ξ = 0, κ = 5.4)	0	1	0	5.31
Right-Skewed (ξ = .8, κ = 4.8)	0	1	.79	4.73
Left-Skewed (ξ = -.8, κ = 4.8)	0	1	-.79	4.73

The symmetric-kurtic example (ξ = 0, κ = 5.4) is chosen to approximate the risk-neutral distributions that can be inferred from the prices of many foreign currency options. The leptokurtotic, left-skewed example (ξ = -.8, κ = 4.8) is chosen to approximate the risk-neutral distribution now commonly inferred from post-1987 crash prices of S&P 500 European options, as well as index options priced on German and Japanese equity indexes. Notice also in this case, the near second modality in the lower left-hand tail corresponds remarkably well the the slight second modality frequently observed in these implied distributions.

II. Risk-Neutral "Edgeworth Densities" and European Option Values

The application of "Edgeworth densities" to the risk-neutral distribution for European derivatives valuation requires that the the standardized distribution be transformed to the have the risk-neutral mean and standard deviation. In particular, denote $P_j \equiv f(x_j)$ as the risk-neutral probability associated with expiration-date price S_j of the asset underlying the derivative. S_j itself is constructed from x_j as follows:

$$S_j = Se^{\mu t + \sigma\sqrt{t}\, x_j} \quad \text{and} \quad (r/d)^t = \Sigma_j P_j (S_j/S)$$

where $S \equiv$ current underlying asset price
$r \equiv$ annualized riskless return
$d \equiv$ annualized payout return of the underlying asset
$t \equiv$ time-to-expiration in years
$\mu \equiv$ annualized risk-neutral expectation of the <u>logarithm</u> of S_j/S
$\sigma \equiv$ annualized risk-neutral volatility of the <u>logarithm</u> of S_j/S

Then:

$$(r/d)^t = \Sigma_j P_j e^{\mu t + \sigma\sqrt{t}\, x_j} = (\Sigma_j P_j e^{\sigma\sqrt{t}\, x_j})e^{\mu t}$$

$$\log[(r/d)^t] = \log(\Sigma_j P_j e^{\sigma\sqrt{t}\, x_j}) + \mu t$$

$$\mu = [\log(r/d)] - [\log(\Sigma_j P_j e^{\sigma\sqrt{t}\, x_j})]/t$$

Note that this formula for μ is similar to the formula commonly used if S_j/S conforms to a risk-neutral lognormal distribution. In that case, $\mu = [\log(r/d)] - ½\sigma^2$. But since for our case, S_j/S may not be lognormal (indeed that is the purpose of this exercise), the correction to (r/d) – the annualized risk-neutral mean of S_j/S – to obtain the annualized risk-neutral mean of the <u>logarithm</u> of S_j/S must be more general. Exhibit 4 summarizes these transformations.

Exhibit 4

6

Edgeworth Binomial Trees

C=f(S,t)

Risk-Neutral "Edgeworth Densities"

Define $P_j \equiv f(x_j)$.

Given r, d, t and σ, adjust S_j by choosing μ such that:

$$S_j = Se^{\mu t + \sigma\sqrt{t}\, x_j} \qquad (r/d)^t = \Sigma_j P_j (S_j/S)$$

Then:

$$(r/d)^t = \Sigma_j P_j e^{\mu t + \sigma\sqrt{t}\, x_j} = (\Sigma_j P_j e^{\sigma\sqrt{t}\, x_j})e^{\mu t}$$

$$\log[(r/d)^t] = \log(\Sigma_j P_j e^{\sigma\sqrt{t}\, x_j}) + \mu t$$

$$\mu = [\log(r/d)] - [\log(\Sigma_j P_j e^{\sigma\sqrt{t}\, x_j})]/t$$

Note that the logarithms of the resulting S_j will be equally spaced:

$$\log S_{j+1} - \log S_j = \log S_j - \log S_{j-1} \quad (j = 1, \ldots, n-1)$$

The discrete density P_j for $j = 0, \ldots, n$, defined on the points S_j now has the desired mean, standard deviation, and approximately the desired skewness and kurtosis. To value European calls with striking prices K_i, $i = 1, \ldots, m$, we simply calculate:

$$C(K_i) = (\Sigma_j P_j \max[0, S_j - K_i])/r^t$$

For four different risk-neutral densities (corresponding to our Exhibit 3), the attached picture converts these call values into their Black-Scholes implied volatilities to obtain alternative option smiles.

Of course, the smile consistent with discretized lognormal probabilities ("Normal") is flat as required for the validity of the Black-Scholes formula. In the other cases, the Black-Scholes formula, although still useful as a device for translating option prices into implied volatilities, does not hold. For example, a symmetric risk-neutral Edgeworth density with kurtosis of 5.4 translates into a more or less symmetric smile pattern around-the-money, valuing options less than their Black-Scholes values near-the-money, but higher than their Black-Scholes values sufficiently away-from-money. In Exhibit 5, the 0.0 point along the horizontal axis corresponds to a striking price of $K(r/d)^{-t}$. Other points along the axis refer to the number of standard deviations of changes in the natural logarithm of the underlying asset price, where a single standard deviation equals $\sigma\sqrt{t}$.

The smiles of skewed densities, not surprisingly, are skewed. For example, the left-skewed density has a skewness of -.8 and a kurtosis of 4.8. This translates into a smile for which a call with a striking price two standard deviations in-the-money has an 8% higher implied volatility (26%) than the implied volatility of the corresponding call with a striking price two standard deviations out-of-the-money (18%)[5].

Exhibit 5

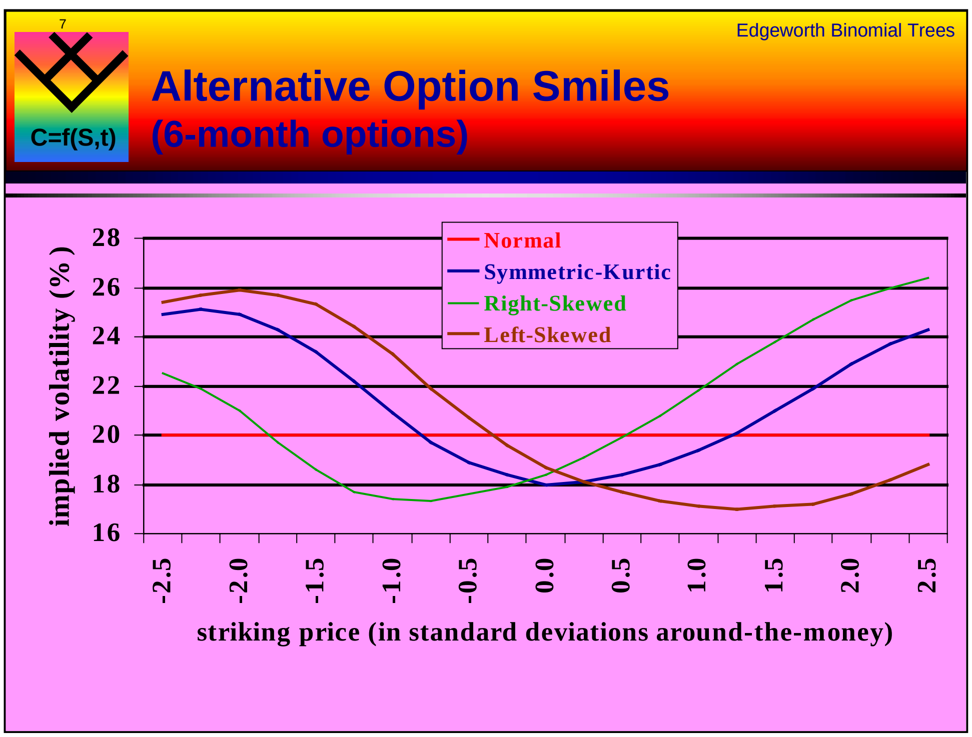

[5] The negatively sloped portion of the left-skewed and symmetric-kurtic smiles on the far left at around -2.0 standard deviations is probably caused by the almost second modality in that occurs at about -2.0 standard deviations in the corresponding Edgeworth densities.

III. Risk-Neutral Stochastic Processes

The risk-neutral probability distribution at the expiration date is only part of the story. We also want to know the stochastic process that leads to this distribution. In a discrete version of the Black-Scholes model, this can be described by a recombining binomial tree with constant multiplicative up and down moves, and constant riskless and payout returns. After of sequence of these moves, the probabilities at the end of the tree can be made to approximate closely a risk-neutral lognormal distribution with a prespecified volatility and mean.[6] However, if the target risk-neutral distribution departs significantly from lognormal, this simple binomial stochastic process must perforce be inconsistent with this.

So one might ask, is there a way to modify the binomial model which leaves its major advantages in tact – its intuitive simplicity and numerical tractability – but at the same time is consistent with the target risk-neutral distribution? It turns out this can be done even while retaining the main attractive features of the the binomial approach:[7]

- binomial price moves,
- recombining nodes,
- ending nodal values organized from lowest to highest,
- constant riskless and payout returns, and
- all paths leading to the same ending node having the same risk-neutral probability.

This last feature means that if you stand at a node at the end of the tree and look backwards, you will see many paths from the beginning of the tree that lead to that node. Each of these paths has the same probability. This does not mean that all paths in the tree have the same probability, but that *conditional on ending up at a particular terminal node*, the paths have the same probability. However, in an important way the modified binomial tree differs from the standard tree: *it does not require constant move sizes.* It allows the local volatility of the underlying asset return to vary with changes in the underlying asset price and time. In addition, it can be shown that given the ending risk-neutral distribution, the riskless and payout returns, and with the above assumptions, there exists a unique consistent binomial tree, which moreover, preserves the property that there are no arbitrage opportunities in the interior of the implied tree (all risk-neutral move probabilities, although they may be different at each node, are non-negative).

As shown in exhibit 6, deriving this generalized tree from the expiration-date risk-neutral density is quite easy.

Step Zero: Start at the end of the tree where there is a node corresponding to each S_j for $j = 1, . . ., n$. Corresponding to each of these is a risk-neutral nodal probability P_j. Calculate the corresponding risk-neutral probability of a single path to that node $P = P_j/[n!/j!(n-j)!]$. Consider two adjacent ending nodes, with path probability, nodal value pairs (P^+, S^+) and (P^-,S^-). Let (P, S) be this pair for the prior node in the tree that has these two possible outcomes.

[6] See J. Cox, S. Ross and M. Rubinstein, “Option Pricing: A Simplified Approach,” *Journal of Financial Economics* 7, No. 3 (September 1979), pp. 229-263,

[7] See M. Rubinstein, “Implied Binomial Trees,” *Journal of Finance* 49, No. 3 (July 1994), 771-818.

Step One: The path probability for the prior node P must equal the sum $P^- + P^+$, because, having arrived at the prior node, either a down move ending up at S^- with probability P^- or an up move ending up at S^+ with probability P^+ must happen.

Step Two: By the laws of conditional probability, conditional on being at the prior node S, the probability of then moving up to S^+ must be P^+/P, so that the up "move probability" $p = P^+/P$; similarly, the down move probability must be $1 - p = P^-/P$.

Step Three: Since p is a risk-neutral probability and the move probability of ending up at S^+ is p and $1 - p$ is the move probability of ending down at S^-, then S must be the risk-neutral expectation of S^- and S^+ discounted at the riskless return over the period with a correction for payouts. Under the assumption of constant riskless and payout returns,[8] then the single-move riskless return $r \equiv r^{t/n}$ and the single-move payout return $\delta \equiv d^{t/n}$.

Exhibit 6

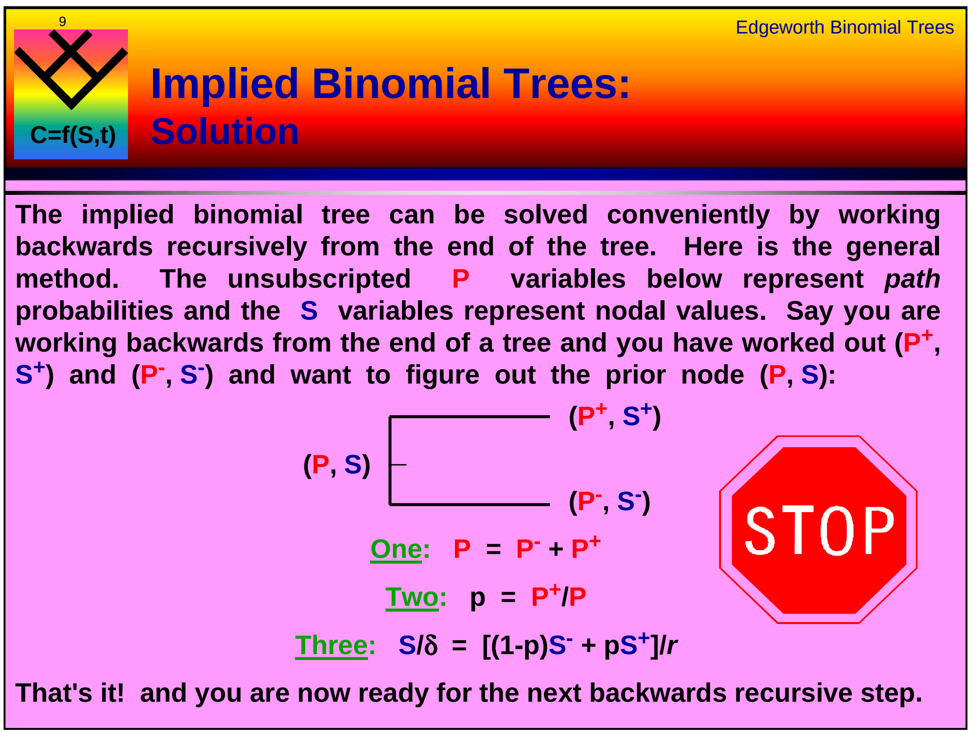

Having shown above how to fill in the binomial tree one-move prior to expiration, the same procedure can then be followed recursively working backwards to fill in the entire tree.

IV. Implied Tree Option Valuation

Having constructed the binomial tree of underlying asset prices, it can be used to value derivatives. As we have seen, the value of a European call expiring at the end of the tree is simply its risk-neutral expected payoff discounted back to the present at the riskless return. No tree is needed for this, only the ending nodal values and associated risk-neutral probabilities.

---

[8] The ratio r/d cannot be chosen independently of the current underlying asset price and its expiration-date risk-neutral distribution. Indeed, $(r/d)^t = \Sigma_j P_j(S_j/S)$.

However, a tree is needed to infer the current value of options with earlier maturities – as well as American or exotic options – in a way that is consistent with the values of the longer maturity options. Exhibit 7 shows what happens when the tree that is constructed by the working backwards procedure from six-month options is used to value options with shorter maturity. For example, the current values of three-month options can be calculated from the underlying asset nodal values and probabilities half-way through the six-month option tree. The option payoff is evaluated using these three-month underlying asset values weighted by their associated risk-neutral probabilities, and discounted back to the present using the riskless return over three months.

In terms of risk-neutral probabilities, the total volatility after three months is less than after six months. Therefore, to compare option smiles, each constructed from otherwise identical options but with different maturities, it is useful to scale the horizontal axis in units of standard deviations of the underlying asset price over the life of the options, as we have done in Exhibit 7.

Exhibit 7

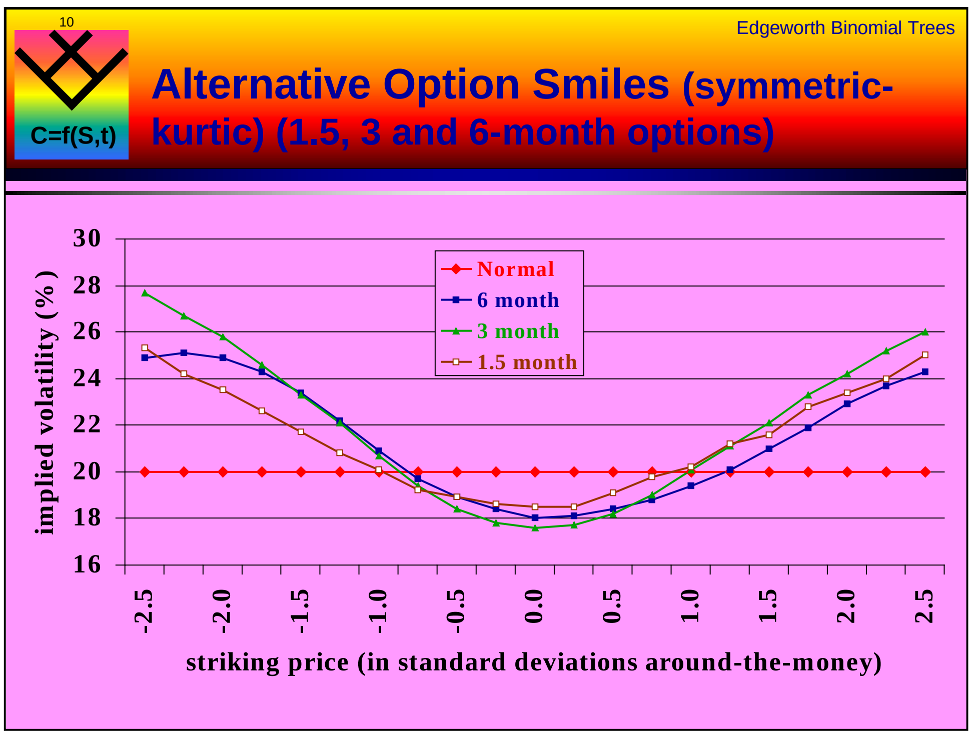

As we can see for the symmetric-kurtic case, the shapes of the smiles for options of different maturities are roughly similar. Exhibits 8 and 9 compare these smile shapes for options based on right-skewed and left-skewed six-month risk-neutral density functions. Other things equal, the smile shapes for different maturity options are roughly similar, with a somewhat more extreme smile the shorter the maturity.

The implied binomial tree also predicts future smiles conditional on the underlying asset price and the remaining time-to-expiration. For example, consider our six month options. We can think of the future evolving along one of the paths in the tree. In three months, we will be standing at some node down the tree. Emanating from this node is the remaining tree which determines the prices of options at that node. Again knowing these prices, we can use the Black-

Scholes formula to convert them into an implied volatility smile for the original options which now have three months to expiration.

Exhibit 8

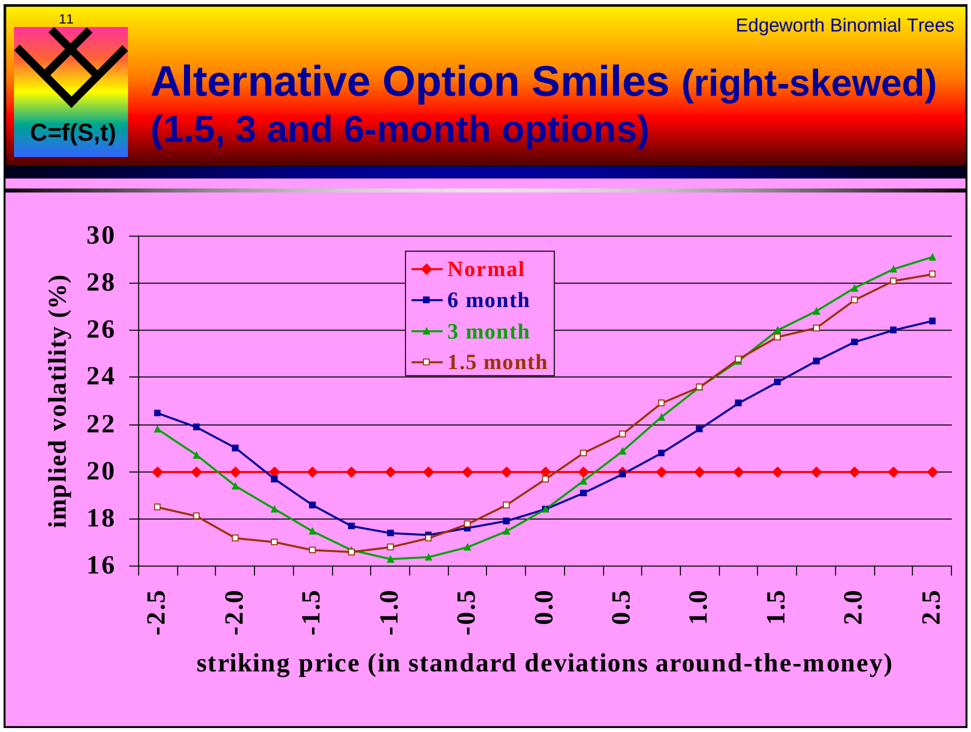

Exhibit 9

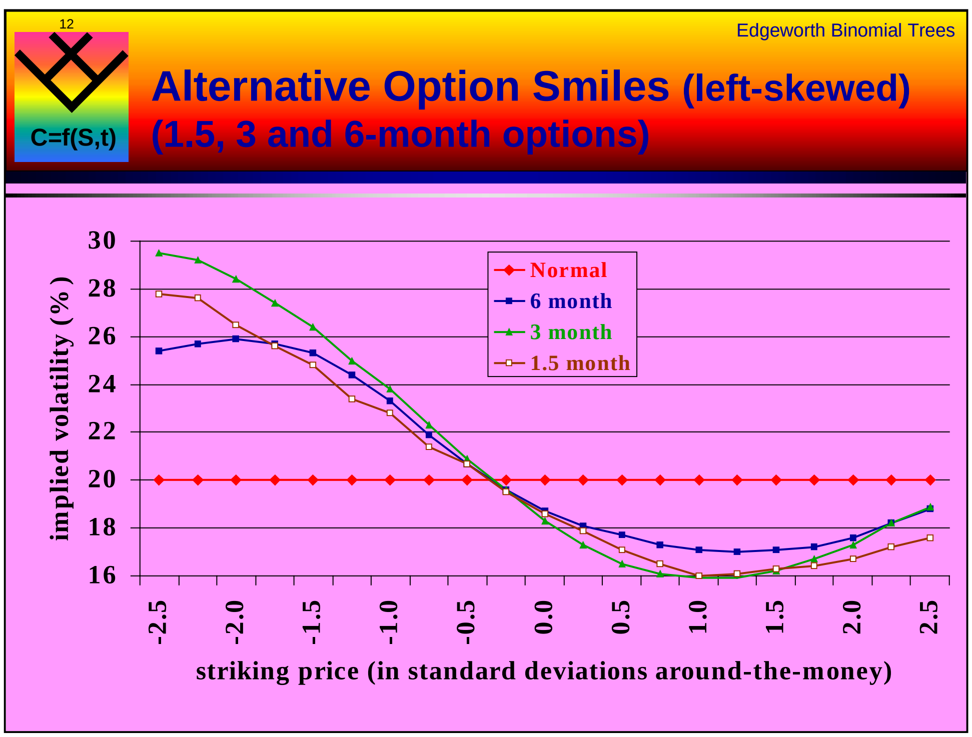

For the symmetric-kurtic case, Exhibit 10 shows the future three month smiles conditional on the underlying asset price remaining about the same (100), decreasing by about one standard deviation (90) or increasing by about one standard deviation (110). The spacing along the horizontal axis continues to be in units of standard deviation remaining to expiration, but now the total standard deviation is 10% (.2√.25) instead of 14.14% (.2√.5). When the asset price is unchanged at about 100, the smile retains it original shape but reflects lower implied volatility everywhere. When the asset price falls by 10%, the smile is negatively sloped over most of its relevant range, which is the sort of shape that translates into a left-skewed risk-neutral distribution. The intuition for this may be that with higher probability than the lognormal in its lower tail, if the asset price moves down then the conditional distribution becomes left-skewed. This reasoning could also explain why the smile becomes positively sloped if the asset price rises by 10%.

Exhibit 10

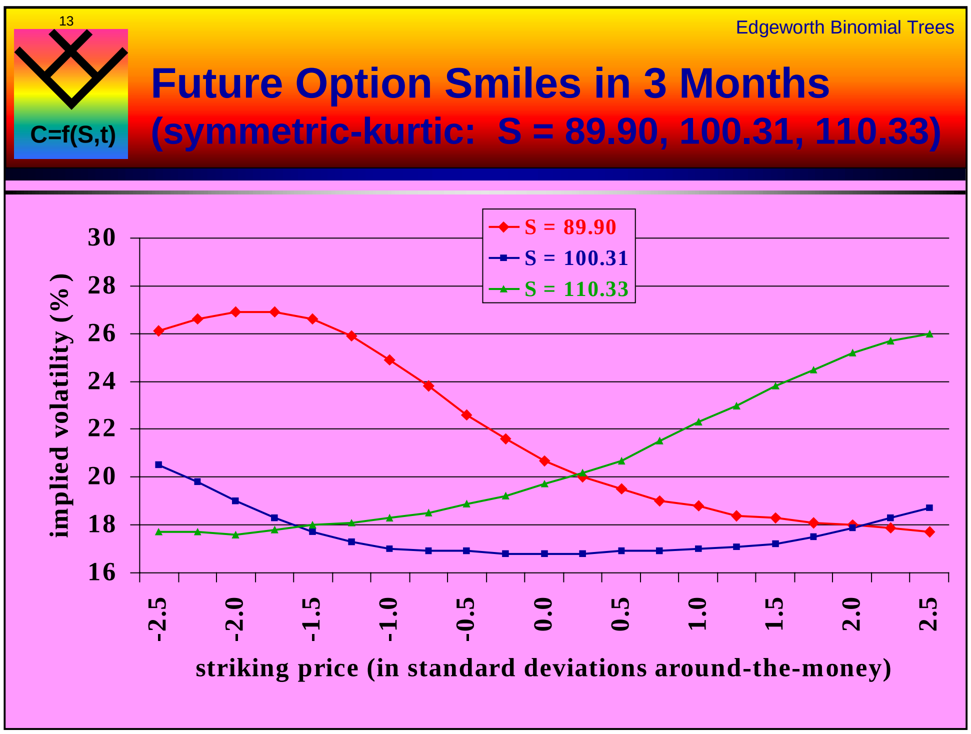

Exhibits 11 and 12 look at what happens to the original smiles for the right- and left-skewed cases. Again we see that if the asset price remains about the same (100), then the shapes of the original smiles three months later are more or less retained, but reflect lower implied volatilities everywhere. Most interesting is that for the originally right-skewed (left-skewed) smile, if the underlying asset price goes down (up) by 10%, the slope of the remaining smile reverses and becomes left-skewed (right-skewed).

Exhibit 11

Exhibit 12