

# **Recovering Stochastic Processes from Option Prices**

by

Jens Carsten Jackwerth

and

Mark Rubinstein\*

May 11, 2011

\* Jens Carsten Jackwerth is a professor of finance at the University of Konstanz ([jens.jackwerth@uni-konstanz.de](mailto:jens.jackwerth@uni-konstanz.de)). Mark Rubinstein is a professor of finance at the Haas School of Business, University of California at Berkeley ([rubinste@haas.berkeley.edu](mailto:rubinste@haas.berkeley.edu)).

## **Recovering Stochastic Processes from Option Prices**

### **Abstract**

How do stock prices evolve over time? The standard assumption of geometric Brownian motion, questionable as it has been right along, is even more doubtful in light of the stock market crash of 1987 and the subsequent prices of U.S. index options. With the development of rich and deep markets in these options, it is now possible to use options prices to make inferences about the risk-neutral stochastic process governing the underlying index. We compare the ability of models including Black-Scholes, naïve volatility smile predictions of traders, constant elasticity of variance, displaced diffusion, jump diffusion, stochastic volatility, and implied binomial trees to explain otherwise identical observed option prices that differ by strike prices, times-to-expiration, or times. The latter amounts to examining predictions of future implied volatilities.

Certain naïve predictive models used by traders seem to perform best, although some academic models are not far behind. We find that the better performing models all incorporate the negative correlation between index level and volatility. Further improvements to the models seem to require predicting the future at-the-money implied volatility. However, an “efficient markets result” makes these forecasts difficult, and improvements to the option pricing models might then be limited.

## Recovering Stochastic Processes from Option Prices

How do stock prices evolve over time?<sup>1</sup> Ever since Osborne (1959), the standard view has been that stock prices follow a geometric Brownian motion. Merton (1973) uses this assumption as the basis for an intertemporal model of market equilibrium, and Black and Scholes (1973) uses it as the basis for their option pricing model. Tests of options on stock in the early years of exchange-traded options more or less supported the implications of Brownian motion, see, for example, Rubinstein (1985). While it has long been well known that empirical return distributions exhibit fatter tails than implied by Brownian motion, evidence that something is not all right with this world is that S&P 500 index options since the crash of 1987 exhibit pronounced volatility smiles, see Jackwerth and Rubinstein (1996). A volatility smile describes implied volatilities that are largely convex and monotonically decreasing functions of strike prices.<sup>2</sup> Such volatilities contradict the assumptions of geometric Brownian motion and perfect markets, which would imply a flat line. Another way to describe this is that the implied risk-neutral probability densities are heavily skewed to the left and highly leptokurtic, unlike the lognormal assumption in Black-Scholes. Like the equity premium puzzle, this option pricing puzzle may ultimately lead us to a better understanding of the determinants of security prices.

There are three possibilities why option prices can spuriously exhibit volatility smiles: First, there are market imperfections, and observed option prices are always different from the true option prices at any time. The S&P 500 index option market is a rather deep and liquid market with rather unfettered access. Its daily notional volume is sizable, as reported in Table I for longer-term options. Even as the daily notional volume increased six-fold from \$1.5 billion in 1989 to \$8.5 billion in 1995, the volatility smile did not change. Most of our results are based on longer-term options, which account for about 4% of the total daily notional volume in all maturities. However, our results do not seem to be sensitive to our focus on the longer-term options.

---

<sup>1</sup> Li and Pearson (2008) repeated the here presented research very faithfully while updating some of the option pricing models.

<sup>2</sup> The implied volatility ( $\sigma^*$ ) causes the Black-Scholes formula to accurately price the option in the market:  $C = Sd^{-t}N(x) - Kr^{-t}N(x - \sigma^*\sqrt{t})$ , where  $S$  is the index level,  $d$  the dividend yield,  $t$  the time,  $N(\bullet)$  the cumulative normal distribution,  $x = \frac{\ln(Sd^{-t}/Kr^{-t})}{\sigma^*\sqrt{t}} + \frac{1}{2}\sigma^*\sqrt{t}$ , and  $K$  the strike price.

Table I about here

Since the S&P 500 index is rather high (370 dollars on average from 1986 through 1995), the value of an option is high compared to the bid/ask spread, which for at-the-money options is only some 42 cents, decreasing to 33 cents for out-of-the-money options. Moreover, we expect the true option price to be close to the mid-point quote for most of the time. Thus, market imperfections are not likely candidates to explain the volatility smile.

The second possibility is that option prices are measured correctly but that the implied probabilities are calculated incorrectly. For example, the wrong interpolation or extrapolation method is used to obtain a dense set of option prices across strike prices. Jackwerth and Rubinstein (1996) show however, that the choice of method does not really matter much because most methods back out virtually the same risk-neutral distribution, as long as there are a sufficient number of strike prices, say, about 15.<sup>3</sup>

The third possibility is that the observed option prices are systematically distorted, and that one can make money in the options market by exploiting such mispricing. Jackwerth (2000) takes this view to some extent.

We assume instead that we see correctly measured option prices that yield meaningful implied risk-neutral probability distributions. The volatility smile is then a way of describing the relation of option prices at the same time, with the same underlying asset and the same time-to-expiration, but with different strike prices. Option prices also provide three other types of comparisons that can be windows into an understanding of the stochastic process of the underlying assets:

- (1) Option prices at the same time, with the same underlying asset, and the same strike price, but with *different times-to-expiration*.
- (2) Option prices with the same underlying asset, the same expiration date, and the same strike price, but observed at *different times*.
- (3) Option prices at the same time, with the same time-to-expiration and with the same strike price, but with *different underlying assets*.

---

<sup>3</sup> The methods differ most in the tails, where they tend to agree on the *total* tail probability but distribute this probability differently. We avoid this difficulty by focusing on the center of the distribution and not using far-away-from-the-money option prices. Further evidence on the performance of different methods is surveyed in Jackwerth (1999).

Jackwerth and Rubinstein (1996) consider relationships among option prices at the same time and with the same underlying and time-to-expiration, but with different strike prices. The ultimate objective is to discover a single model that can explain all four relations simultaneously. For example, the post-crash smile of index options and the implied binomial tree model of Rubinstein (1994) strongly suggest that a key aspect of the “correct” model will be one that builds in a negative correlation between index level and at-the-money implied volatility. This can explain the relation in Jackwerth and Rubinstein (1996) and turns out in the post-crash period to be an empirical regularity of relation (2).

While we focus here on the smile in the S&P 500 data for the U.S., Tompkins (2001) documents that similar smiles, albeit not as steep as the U.S. smile, are seen in the UK, Japan, and Germany. In addition, Dennis and Mayhew (2000) show that individual option smiles in the U.S. are not as steep as the index smile, a finding that likely holds for the other markets as well but that has not been documented.

There are several rational economic reasons why the post-crash smile effect might obtain. First, corporate leverage effects imply that as stock prices fall, debt-equity ratios (in market values) rise, causing stock volatility to increase. Second, Kelly (1994) notes that equity prices have become more highly correlated in down markets, again causing an increase in volatility. Third, risk aversion effects can cause investors who are poorer after a downturn in the market to react more dramatically to news events. This would lead to increased volatility after a downturn. Fourth, the market could be more likely to jump down rather than up. Indeed, since the stock market crash period of 1987 until the end of 1998, the five greatest moves in the S&P 500 index have been down. Finally, as the volatility of the market increases, the required risk premium rises, too. A higher risk premium will in turn depress stock prices. We do not try to provide an economic explanation for observed smile patterns, but rather have the more limited objective of comparing alternative models that purport to explain relations (1) and (2). We leave to subsequent research an investigation of relation (3). A comparison of smile patterns for index options and individual stock options, as in Dennis and Mayhew (2000), provides a way to distinguish between leverage and wealth effects as explanations of the inverse correlation between at-the-money option implied volatilities and index levels. If leverage is the force behind the scenes, the downward slope of the smiles for index and stock options should be about the same. If the wealth

effect is predominant, the downward slope of the smile would be highest for index options and become less sloped the lower the ratio of a stock's systematic variance to its total variance.

To investigate the empirical problems, we suggest two main tests. Our first test investigates relation (1), using options prices at the same time and with the same underlying and strike price, but with different times-to-expiration. Here we find out how well different option pricing models are capable of simultaneously explaining option prices of different times-to-expiration. For this, we deduce shorter-term option prices from longer-term option prices. The volatility smile for the longer-term options is assumed known, and the volatility smile for the shorter-term options is unknown. The problem of relation (1) is to fit alternative option pricing models to the longer-term option prices. We can then compare the model values with the observed market prices for the shorter-term options and calculate pricing errors. To help understand the source of remaining errors, we also conduct a related experiment. We assume in addition that we also know the at-the-money implied volatility of the shorter-term options.

The second test investigates relation (2), using option prices with the same underlying asset, expiration date, and strike price, but observed at different times. In this case, we use option valuation models to forecast future option prices conditional on the future underlying asset price. We calibrate alternative models on current longer-term option prices. Then, we wait 10 and 30 days, observe the underlying asset price, and assess the errors in our forecasts. A related test extends the forecasting procedure by incorporating information from both current longer-term and current shorter-term option prices. Again, to decompose the source of any remaining errors, we also assume in addition that we know in advance the future at-the-money option price.

For all tests, we evaluate five kinds of option valuation models (nine models altogether). We compare deterministic models and stochastic models and naïve trader rules. Related empirical work is in Dumas, Fleming, and Whaley (1998), Bates (2000), and Bakshi, Cao, and Chen (1997). The first paper investigates only different deterministic volatility models while the other two compare only different stochastic models.

The five categories of models are: first, mostly for reference, the Black-Scholes formula; second, two naïve smile-based predictions that use today's observed smile directly for prediction; third, two versions of Cox's (1996) constant elasticity of variance (CEV) formula; fourth, an implied binomial tree model; fifth, three parametric models that specify the stochastic process of the underlying, namely, displaced diffusion, jump diffusion, and stochastic volatility.

The naïve predictions do not rely on any solid theoretical basis, but we examine them because they are very simple and widely used by professionals. We show that they perform surprisingly well compared to the more rigorous academic models. We use the CEV model because it explicitly builds in an assumption that local volatility is negatively correlated with the underlying asset price and is therefore a natural candidate, given our observations.

Implied binomial trees, which are non-parametric, have been proposed by Rubinstein (1994), Derman and Kani (1994), and Dupire (1994). Work on implied risk-neutral distributions that is closely related has been conducted by Jackwerth and Rubinstein (1996) and Aït-Sahalia and Lo (1998).<sup>4</sup> We focus here on the implied binomial trees in Rubinstein (1994) and the generalizations in Jackwerth (1997). The generalizations allow us to incorporate information from times other than the end of the tree. We rely solely on the observed option prices in the market, and thus avoid having to specify a stochastic process a priori.

Next, we introduce the data. Then we conduct our two tests. Sections II and III are concerned with inferring shorter-term option prices from concurrent longer-term option prices (relation (1)), with an unknown and known term structure of volatilities respectively. Sections IV and V are concerned with forecasting future smiles using current longer-term option prices (relation (2)), with an unknown and known term structure of future volatilities respectively. We conclude with our surprising result that the naïve trader rules work as well as the more rigorous academic models.

## **Data**

The database includes minute-by-minute trades and quotes covering S&P 500 European index options, S&P 500 index futures, and S&P 500 index levels from April 2, 1986, through December 29, 1995.

All option models are parameterized to price the observed longer-term options best, those with times-to-expiration of between 135 and 225 days. In the first test, the models are then used to price shorter-term options with 45 to 134 days to expiration.

To obtain sets of option prices across several strike prices for the two times-to-expiration, we aggregate all daily quotes into two volatility smiles, one for the shorter- and one for the longer-term options. Throughout each day, we calculate the implied volatilities for all options with the

---

<sup>4</sup> A survey is in Jackwerth (1999).

same strike price and time-to-expiration. We compute the median implied volatilities for each strike price and treat this set as our representative daily volatility smile for a given time-to-expiration.

Interestingly, the number of available options quotes during the day does not influence the results very much. As in Jackwerth and Rubinstein (1996), we use only strike prices with strike price / index level ratios (moneyness) between 0.79 and 1.16 because of the lack of liquidity for the further away options.

The dividend yield is based on the actual payments throughout the life of the option. The interest rate is the average of the median implied borrowing and lending rates assuming put-call parity of all feasible pairs of options for a given time-to-expiration on a given day. The index level for our representative daily sets of option prices is the average of the daily high and low of a futures-based index. The futures-based index is obtained by deflating all futures quotes and trades by the median daily implied repo rate corresponding to the time-to-maturity of the future. For each minute, the median of all deflated quotes and trades is computed and used as the futures-based index for that minute.<sup>5</sup>

There are 2074 days in the almost ten years elapsed time for which we have a sufficient number of longer-term options. We specify two subperiods: a pre-crash period from April 2, 1986, through October 16, 1987, and a post-crash period from June 1, 1988, through December 29, 1995. We avoid the period right after the crash, which is often difficult to interpret empirically, as the market took about half a year to get settled again. The sample size is 1953 days: 386 days for the pre-crash period, and 1567 days for the post-crash period. For the empirical studies, where we need both shorter-term and longer-term options, there are 1639 days: 372 days for the pre-crash period, and 1267 days for the post-crash period. For the smile forecasts, we shorten the pre-crash period so that it ends September 16, 1987. This avoids forecasting across the crash with associated large errors.

To obtain the implied probability distributions as inputs for the implied binomial trees, we use the maximum smoothness method proposed in Jackwerth and Rubinstein (1996). As Figure 1 shows, even though the smoothness criterion does not rely on a lognormal prior distribution, the implied probability distribution calculated from option prices that are based on a lognormal distribution is very close to the underlying lognormal distribution.

---

<sup>5</sup> See Jackwerth and Rubinstein (1996) for a more detailed description of the data.



Figure 1 about here

Since we use the method of Jackwerth and Rubinstein (1996) for finding implied probability distributions, we sample the implied distributions on equal dollar-spaced asset values. Implied binomial trees are generally sensitive to the spacing at the end of the tree and do not work well with equal spacing. Thus, we have to resample the implied probability distributions onto equally *log-spaced* asset values that are given by a standard binomial tree with the same number of steps. To construct the standard binomial tree, we have to specify the volatility parameter, which we set equal to the implied volatility of the longer-term at-the-money option. For the resampling, we use piecewise-linear interpolation in the cumulative probabilities. The use of cubic splines improves performance only marginally but at a significantly higher computational cost.

The resulting probability distribution overprices options only slightly, with a median absolute error of about 3 cents for a test of 23 semiannual observations, even if the log spacing spans as few as 80 values. We could detect no pricing bias across strike prices.

Each model used to derive the shorter-term options is calibrated on the same set of longer-term options, although the different models imply different (risk-neutral) probability distributions for the longer-term options. Figure 2 depicts these distributions for the implied binomial tree, the Black-Scholes model, and the CEV model for a representative day after the crash. The implied binomial tree and the CEV models show rather similar implied distributions, but the implied binomial tree is more flexible.

Figure 2 about here

## **Looking Backward (Unknown Term Structure of Volatilities)**

The first test investigates otherwise similar options but with different times-to-expiration (relation (1)). We look backward in that we fit the models to concurrent longer-term option prices. Then, we price the shorter-term options using the fitted models.

### **Alternative Option Pricing Models**

The simplest model is Black-Scholes, which suggests that the volatility smile is truly flat (“flat smile” model). That is, we price shorter-term options according to a flat smile, which is

determined by the at-the-money implied volatility of the concurrent longer-term options. We repeat this calibration each day for all models. Note that the resulting parameters will vary from day to day. Such recalibrations are typical for how practitioners use the models and take account of the fact that the volatility smile moves from day to day.

Two other models, the displaced diffusion model of Rubinstein (1983) and the jump diffusion of Merton (1976), produce empirical results very similar to the Black-Scholes model in all tests, so we do not report these results separately. Neither of these models is flexible enough to generate volatility smile patterns that are sufficiently different from the flat smile of the Black-Scholes model.

The second set of predictions is often used by traders to accommodate observed non-flat smile patterns. These assume that the longer-term smile is the best guess for the shorter-term smile. There are two different versions depending on if the smile is defined across moneyness or strike price. One version is the “relative smile” prediction, which traders call the “sticky delta” method.<sup>6</sup> Here we assume that each shorter-term option with a given *moneyness* has an implied volatility equal to the corresponding longer-term option with the same moneyness. In a second version, we consider the “absolute smile” prediction, which traders call the “sticky vol” method. Here we assume that a shorter-term option with a given *strike price* has an implied volatility equal to the corresponding longer-term option with the same strike price.

For the relative and absolute smiles, we need to interpolate the longer-term implied volatilities to arrive at the corresponding implied volatilities for the shorter-term options. We use piecewise-linear interpolation since it gives very similar results to more sophisticated methods such as cubic spline or polynomial interpolation. We extrapolate for values (of moneyness or strike prices) outside the longer-term smile by assuming that all lower values have the same implied volatility as the lowest observable value and all higher values those of the highest observable value.<sup>7</sup>

Next, we examine two versions of the constant elasticity of variance model. The formula is based on a diffusion process specifying the instantaneous volatility  $\sigma(S,t)$  as:

---

<sup>6</sup> See Reiner (1998) who analyzes single-factor models that exhibit such “sticky delta” or “sticky vol” behavior. He finds that such processes would not be diffusion processes, which could cause problems in hedging such options.

<sup>7</sup> The typical smile is u-shaped, so we will introduce some bias in the far-away-from-the-money options which will be undervalued. While the effect should not be too large as those options have low option premia, it stacks the tests somewhat against the naïve methods.

$$\sigma(S,t) = \sigma' S^{\rho-1} \quad (1)$$

where:  $\sigma'$  = volatility parameter;  
 $S$  = index level; and  
 $\rho$  = 1 + elasticity of variance with respect to index level.

Several well-known processes are nested in the CEV such as the absolute process ( $\rho = 0$ ), the square-root process ( $\rho = 0.5$ ), and the standard process underlying the Black-Scholes formula ( $\rho = 1$ ). If the parameter  $\rho$  is less than one, then the index level and the instantaneous volatility are inversely related, consistent with empirical findings. We measure a correlation between S&P 500 index level and at-the-money implied volatility of -0.72 for the 1880 observations from June 6, 1988, through December 28, 1995, so we impose an upper bound on  $\rho$  at 1. Options priced under this parameterization will exhibit a downward-sloping volatility smile.


We allow for two different lower bounds on  $\rho$ . In the restricted version, a lower bound of zero is chosen, corresponding to the absolute process. This limits the chance of bankruptcy to plausible levels. In a second version, we leave the  $\rho$  parameter unrestricted. Although this is mathematically legitimate, it may have undesirable economic implications (which we discuss later).

The CEV model has two additional parameters that need to be specified, the elasticity  $\rho - 1$  and the volatility  $\sigma'$ . We choose both parameters in an optimal fashion by minimizing the pricing error of the longer-term options. We then price the shorter-term options while using these two optimal parameters.

For the implied binomial tree model (Rubinstein, 1994), the natural starting point is the end of the tree, where we assume that all ending probabilities and asset nodal values are known from the method of Jackwerth and Rubinstein (1996).<sup>8</sup> In order to allow for subsequent generalizations, we use the notation of Jackwerth (1997). Let  $S_{i,j}$  be the nodal asset value at step  $i = 0, \dots, n$  and at level  $j = 0, \dots, i$ . Let  $P_{i,j}$  be the corresponding nodal probability. The recursive method is then, starting at the end of the tree:

---

<sup>8</sup> We prefer the implied binomial tree of Rubinstein (1994) to the Derman and Kani (1994) tree because it is less susceptible to numerical instabilities resulting from negative probabilities within the tree. Even the improvements by Barle and Cakici (1998) do not eliminate all the instabilities of the Derman-Kani tree.



Step One: 
$$P_{i-1,j} = (1 - w_{i,j})P_{i,j} + w_{i,j+1}P_{i,j+1} \quad (2)$$

Step Two: 
$$p_{i-1,j} = w_{i,j+1} \frac{P_{i,j+1}}{P_{i-1,j}} \quad (3)$$

Step Three: 
$$S_{i-1,j} = [(1 - p_{i-1,j})S_{i,j} + p_{i-1,j}S_{i,j+1}] / (r / \delta)^{1/n} \quad (4)$$

where:  $w_{i,j}$  =  $(j / i)$  is the transition probability weight of going down as we move to the prior step;  
 $r$  =  $1 +$  annualized interest rate;  
 $\delta$  =  $1 +$  annualized dividend yield;  
 $t$  = time-to-expiration in years; and  
 $n$  = number of steps.

To price shorter-term options, we find the step  $i^*$  of the implied binomial tree corresponding in time:

$$i^* = \text{integer}[0.5 + \frac{t_s}{t} n] \quad (5)$$

where:  $t_s$  = time-to-expiration of the shorter-term options; and  
 $t$  = time-to-expiration of the longer-term options.

We can now price all shorter-term options using the formula:

$$C(K_m) = \left( \sum_{j=0}^{i^*} P_{i^*,j} \max[0; S_{i^*,j} - K_m] \right) / r^{t_s} \quad (6)$$

where:  $K_m$  = strike price of the shorter-term option  $m$ .

We use implied binomial trees with 200 steps. Binomial trees are rather insensitive to the number of steps used. Even trees with as few as 80 steps yield pricing errors that are only 4 cents worse than the errors for 200 step trees which vary from 24-83 cents across different tests. Increasing the number of steps beyond 200 does not improve the fit of the trees significantly.

The last model is the stochastic volatility model of Heston (1993). Volatility is modeled as a mean-reverting square-root process. In his exact notation for the risk-neutral process:<sup>9</sup>

$$\begin{aligned} dS(t) &= rSdt + \sqrt{v(t)}Sdz_1(t) \\ dv(t) &= \kappa^* [\theta^* - v(t)]dt + \sigma\sqrt{v(t)}dz_2(t) \end{aligned} \tag{7}$$

where:  $\kappa^*$  = mean reversion;  
 $\theta^*$  = long-run variance;  
 $v(t)$  = current variance;  
 $\rho$  = correlation of  $z_1(t)$  and  $z_2(t)$ ;  
 $\sigma$  = volatility of volatility parameter;  
 $r$  = interest rate.

Again, we fit the first five parameters by minimizing the error of the longer-term options and then keep the parameters to price the shorter-term options.

More complex models incorporate stochastic interest rates and stochastic jumps in addition to stochastic volatility. Evidence in Bakshi et al. (1997), however, suggests that adding jumps only slightly improves on the stochastic volatility model for short-term options. Adding stochastic interest rates only slightly improves on the stochastic volatility model for long-term options. In both cases, there is no difference for hedging performance. Bates (2000) finds that adding jumps slightly improves performance over the stochastic volatility model, but that the implied probability distributions are inconsistent. Thus, we focus on the simpler stochastic volatility model since it seems to account for the lion's-share of performance improvement over the Black-Scholes model.

One model we do not include is the GARCH option pricing model of Duan (1995). Since some GARCH models converge to stochastic volatility models as the discrete time steps grow small (Heston and Nandi, 2000), we believe that we capture much of the effect by investigating the stochastic volatility model.

## **Empirical Results**

Our empirical results are stated in terms of pricing errors. To evaluate the performance of the alternative option models we use the root mean squared error in cents:

---

<sup>9</sup> Some of the variable names are used in other models, too, but we wanted to keep them as in the original model.

$$RMSE = \sqrt{\frac{1}{m} \sum_{i=1}^m [C_{\text{observed}}(K_i) - C_{\text{model}}(K_i)]^2} \quad (8)$$

where:  $C_{\text{observed}}$  = midpoint option price observed in the market;  
 $K_i$  = strike price for the observed option  $i$ ; and  
 $C_{\text{model}}$  = option price derived from an option valuation model.

Thus, the pricing error is an aggregate measure across all strike prices for options of a given time-to-expiration. Expressed in cents, the relative impact is less for a call option with a low strike price than for a call option with a high strike price. Choosing March 16, 1990, as an example, a typical day for the post-crash period, the at-the-money option (strike price = 34500 cents) is worth 975 cents. A low pricing error of 20 cents corresponds to 2%, and a high pricing error of 100 cents corresponds to 10%. The corresponding percentages for the 6%-in-the-money call (strike price = 32500) with price 2442 are 1% and 4%. The corresponding percentages for the 6%-out-of-the-money call (strike price = 36500) with price 148 are 14% and 68%.

There are two reasons we use cent errors. First, because we translate all puts into calls through put-call parity, we are concerned that expressing relative errors in terms of call prices might have adverse consequences for the corresponding put prices. The cent errors are unaffected however: if the put is over-priced by a cent, then the put-call-parity implied call is also over-priced by a cent. Second, the alternative of expressing errors in terms of implied volatilities means that for in-the-money calls small cent errors translate into large volatility errors.

Tables II and III exhibit the pricing errors for the alternative models.<sup>10</sup> Table II reports the pre-crash errors and Table III the post-crash pricing errors.

Table II about here

Table III about here

## Discussion

In the pre-crash period, assuming an unknown term structure of implied volatilities, the Black-Scholes formula yields median errors of 21 cents (Table II). These errors compare quite

---

<sup>10</sup> If any model cannot be evaluated on a given date, that date is eliminated from the sample. Thus, the number of observations for the two periods vary from table to table.

favorably with typical bid/ask spreads of 40 cents. Since pre-crash smiles are almost flat, there is no benefit to using the relative or absolute smile predictions.

The CEV model yields marginally lower errors because of the greater flexibility in choosing the  $\rho$  parameter, but it converges to the Black-Scholes formula as the parameter  $\rho$  takes on an optimal value of 1. This occurs for most days in the pre-crash period since the observed option prices imply probability distributions that are about lognormal. Moreover, the bound on  $\rho$  at zero for the restricted CEV model is almost never binding. This means that the two versions of the CEV model are virtually indistinguishable.

Implied binomial trees do slightly worse, with a median error of 22 cents. This is surprising, because the implied binomial tree nests the Black-Scholes formula if the ending probability distribution is lognormal. The implied probability distributions are not perfectly lognormal, however. The method used to derive them fits the implied probability distribution to match midpoint option prices that do not necessarily observe a flat smile, even in the pre-crash period. Although a flat smile usually prices all options within their bid/ask spreads, Jackwerth and Rubinstein (1996) argue in favor of using midpoints for numerical reasons in deriving the implied distributions.<sup>11</sup> Finally, the implied distribution has to be resampled onto a log-spacing that introduces a small error in the order of a few cents.<sup>12</sup>

Similarly, the stochastic volatility model also converges to the Black-Scholes model in the pre-crash period.

The post-crash period produces more interesting results. Black-Scholes does the worst, with a median error of 125 cents. The relative and absolute smile predictions are best, with median errors of 51 and 54 cents, respectively. The implied binomial tree holds the middle ground with a median error of 78 cents.

The performance of the CEV depends strongly on restriction in the choice of the  $\rho$  parameter, with median errors of 110 (restricted) and 57 cents (unrestricted). Limiting  $\rho$  to values between 0 and 1 causes the CEV to come in second-to-last, while the unrestricted version is a

---

<sup>11</sup> Since we are fitting midpoints, we have to be concerned about overfitting. Therefore, we also use implied probability distributions with a lower penalty parameter that produces smoother distributions with less fit. Finally, we use the (very smooth and unimodal) probability distribution implied in the CEV model after fitting the CEV model to observed option prices. The errors do not change significantly with the degree of smoothness in the probability distributions, so we use the original implied probability distributions throughout.

<sup>12</sup> One further concern is that the number of known options could influence the error since each incremental option embodies some additional information that could be used in implying option prices. Empirically, we find that the error is virtually uncorrelated with the number of known options.

relatively good model apart from its high standard deviation. The unrestricted version of the CEV model is questionable on economic grounds, however.

Three lines of argument support a restriction on the parameter  $\rho$ . First, several empirical studies using S&P 500 index returns find that estimates of  $\rho$  are mostly confined to the interval between 0 and 1.<sup>13</sup> We note that such estimates stem from the actual return process and not the risk-neutral process, which we use in our option pricing model.<sup>14</sup>

Second, Bates (1996) argues that for stock market indices it is inconceivable that there is a significant probability of bankruptcy; with  $\rho$  sufficiently negative, however, this is possible. Let us consider  $\rho = -4$ , which is the closest integer to the mean optimal unrestricted CEV coefficient in the post-crash period. For three scenarios of interest rates, dividend yields, volatilities, and times-to-expiration that are typical for that period, the CEV model suggests a probability of bankruptcy for the S&P 500 (index = 0) over the next 4 to 7 month's of between 2 and 304 out of 10,000. Typical values tend to lie in the center of this range and seem rather high, given the implication that the whole U.S. economy would be destroyed.

Third, He and Leland (1993) derive a condition that any stochastic diffusion process of asset prices has to meet for the economy to be in equilibrium. They assume a pure exchange economy with dynamically complete asset markets where investors maximize their expected utility of future consumption and have nonrandom exogenous income. This condition will hold if  $\rho = 1$ , which corresponds to the Black-Scholes process. The condition will also hold if, unrealistically, the expected return of the market equals plus or minus the risk-free return. If we accept this as a viable approximation of the U.S. economy, we must conclude that the CEV process is not compatible with that economy in general.

### **Looking Backward (Known Term Structure of At-the-Money Volatilities)**

Continuing the test of relation (1), we still investigate otherwise similar options but with different times-to-expiration. We again look backward in that we price shorter-term options using several methods based on concurrent longer-term option prices. In this case, we want to incorporate a known term structure of at-the-money volatilities. We can then assess how much

---

<sup>13</sup> Beckers (1980), Gibbons and Jacklin (1988), and Christie (1982).

<sup>14</sup> See Jackwerth (2000) for more empirical evidence that the risk-neutral process can be quite different from the actual process.



failure of the models to incorporate the term structure of implied volatilities contributes to overall pricing errors. To do this, we need generalized versions of the models.

### **Alternative Option Pricing Models**

We generalize the Black-Scholes model by assuming a flat volatility smile that is determined solely by the at-the-money implied volatility of the shorter-term options, which we assume is known. For the relative smile prediction, we again assume that the implied volatility for identical levels of moneyness remains the same as for the longer-term options, but we also shift the smile vertically by a constant amount so that we match the implied volatility of the shorter-term at-the-money option. For the absolute smile prediction, we shift the smile vertically by a constant amount in a similar way.

To generalize use of the CEV model, we keep the shape parameter  $\rho$  constant, as determined by the best fit of the longer-term options, but allow the volatility parameter  $\sigma'$  to be chosen optimally so as to fit the shorter-term at-the-money option. Then we price the other shorter-term options using the adjusted parameters.

In order to incorporate a known term structure of volatilities into an implied binomial tree, we use the generalized binomial tree as described in Jackwerth (1997). We can use the same three-step algorithm used for the implied binomial tree, but a generalized binomial tree allows for arbitrary weight functions (instead of the linear weight function associated with a standard implied binomial tree) while preserving no-arbitrage in the interior of the tree. Matching of the volatility term structure is achieved by choosing a weight  $w(0.5)$  between 0 and 1. We know that  $w(0) = 0$  and  $w(1) = 1$ . We then use a piecewise linear function to fit through the three points. This is illustrated for  $w(0.5) = 0.75$  in Figure 3.

Figure 3 about here

We can now select a weight  $w(0.5)$  that causes the at-the-money shorter-term option value to equal its observed market price. Then, all other shorter-term options are priced using this optimal  $w(0.5)$ .

Finally, for the stochastic volatility model, we calibrate the parameters to the longer-term options, and then reset the parameter for the current volatility of the stock price  $v(t)$  to fit the shorter-term at-the-money option.

## Empirical Results and Discussion

Tables IV and V show the pricing errors for the alternative models: Table IV the pre-crash pricing errors and Table V the post-crash pricing errors.

Table IV about here

Table V about here

Assuming a known term structure of at-the-money volatilities now means that the at-the-money shorter-term option is properly priced. The median error for the Black-Scholes formula improves to 13 cents. The relative and absolute smile predictions perform nearly as well, with median errors of 15 and 16 cents, respectively. For the CEV model, the  $\rho$  parameter is 1 for most days, again. Both the restricted and the unrestricted CEV model therefore converge to the Black-Scholes formula.

The generalized binomial tree produces a negligible improvement in the median error compared to the unknown term structure of volatilities. The reason for this disappointing performance in the pre-crash period is closely related to the problem of fitting the longer-term options too well, and thus implying distributions that are not perfectly lognormal. Adjusting the trees according to these distributions to match the shorter-term at-the-money option can result in excessive adjustments that improve the fit for some of the shorter-term options (particularly the at-the-money and near-the-money options) and harm the fit of other options. Thus, the overall effect is either no improvement or even worse performance compared to the assumption of an unknown term structure of volatilities.

The stochastic volatility model also converges to the Black-Scholes model and performs almost as well. In the pre-crash period overall, where the implied probability distributions are about lognormal and the volatility smiles almost flat, using the Black-Scholes formula remains the superior method.

In the post-crash period, assuming a known term structure of at-the-money volatilities, the Black-Scholes formula is again the worst-performing model, with a median error of 121 cents. The relative and absolute smile predictions again do well, with median errors of 42 cents in both cases.

Again, the restricted CEV does almost as poorly as the Black-Scholes model with 101 cents. The unrestricted CEV does much better, with a median error of 40 cents, but the objections to a negative parameter  $\rho$  still apply. Generalized binomial trees do very well, with median error 38 cents, and so does the stochastic volatility model with 43 cents.

Note that the better models now perform very similarly to one another. This is an indication that, after adjusting for the level of the shorter-term smile, there is little advantage in one model over the others.

We conclude that the observed term structure of volatilities is more complex than the models allow. From the present research, it is not clear if a better model of the term structure of volatility can remain deterministic, or if extension to a stochastic volatility structure are needed.

### **Looking Forward (Unknown Term Structure of Volatilities)**

The second empirical test investigates the same options but at different times (relation (2)). Here, we evaluate forecasts of smiles 10 or 30 days in the future. We focus first on incorporating information from the current longer-term option prices. The results from incorporating the shorter-term options turn out to be very similar.

#### **Alternative Option Pricing Models**

For the purpose of forecasting, we use the flat, relative, and absolute smiles as first defined. All smile predictions are based on the current longer-term options, and after waiting for 10 and 30 days we then measure the prediction errors. For the CEV models, we use the first approach again. After estimating the parameters  $\rho$  and  $\sigma'$  from the longer-term options, and potentially restricting  $\rho$ , we use these estimates for our smile forecasts.

We build the implied binomial trees based on the longer-term options. We then wait 10 or 30 days, and we assess which node  $(i, j)$  in the tree is closest to the index level on that future day. If we can now determine the risk-neutral probability distribution, conditional on being at node  $(i, j)$ , then we can price options expiring at the end of the remaining tree.

There is an easy way to obtain the conditional probabilities. First, we go out for, say, 10 days into the future, and find the closest underlying asset value on the step of the tree closest in time to the observed asset value. The corresponding node  $(i, j)$  is the future node. Due to coarseness of the tree, the time and asset value at node  $(i, j)$  are likely to be a bit off the observed values. To

correct for this, we update the observed option prices by assuming that the implied volatilities remain constant for small changes in the asset value and time. Then we choose the asset value to be exactly the one at node (i, j) in the tree and the time to be exactly the one of the selected step i. In addition, we use the interest rate and dividend yield that we use for the original tree. Given all this information and the implied volatilities of the observed options, we now calculate hypothetical option prices (with slightly different time-to-expiration and stock price), which replace the observed options.

The final step is to evaluate the options according to the tree. Here we need conditional probabilities ( $P'$ ) which we obtain by setting  $P'_{i,j} = 1$  since we know with probability one that in 10 days we arrived at that node (i, j). We can easily calculate the probability of an up- and down-move from node (i, j) as:

$$p_{i,j} = w_{i+1,j+1} \frac{P_{i+1,j+1}}{P_{i,j}} \quad (9)$$

where:  $p_{i,j}$  = probability of moving up from node (i, j); and  
 $(1 - p_{i,j})$  = probability of moving down from node (i, j).

We can then recursively develop a tree of the conditional probabilities  $P'$  as we go forward from node (i, j):

$$P'_{k,l} = P'_{k-1,l} (1 - p_{k-1,l}) \quad \text{for } k = 0 \quad (10)$$

$$P'_{k,l} = P'_{k-1,l} (1 - p_{k-1,l}) + P'_{k-1,l-1} p_{k-1,l-1} \quad \text{for } k = 1, \dots, l-1 \quad (11)$$

$$P'_{k,l} = P'_{k-1,l-1} p_{k-1,l-1} \quad \text{for } k = l \quad (12)$$

The conditional probabilities at the end of the tree ( $= P'_{n,l}$ ) are then used to price the set of observed options through:

$$C(K_m) = \left( \sum_{l=0}^n P_{n,l}' \max[0; S_{n,l} - K_m] \right) / r^{t(n-i)/n} \quad (13)$$

where:  $K_m$  = strike price of the observed option m;  
 $r$  =  $1 +$  interest rate;  
 $t$  = time-to-expiration of the original tree; and  
 $i$  = step corresponding to the forecasting node (i, j).

Note that this procedure can adjust properly for small deviations in time, especially since a 200-step tree corresponds to about one day per step. The change in the index from one level to the next, even in a 200-step tree, however, can be coarse enough to affect the results when we update the observed smile by setting the index to its closest match observed in the tree.

There is an elegant way around this problem. Rather than building a tree of conditional probabilities starting at the closest index level, we start the tree of conditional probabilities at two nodes simultaneously. We choose the nodes just above and below the observed index level. Next we assign conditional probabilities  $P'_{i,j+1}$  and  $P'_{i,j}$  to the upper and lower node (summing to one), which causes the expected index level to coincide with the observed index level. At this point we unravel the conditional probability tree to the end. This method proves to be much more stable than the original version.

For the stochastic volatility model, we again calibrate the parameters  $\kappa^*$ ,  $\theta^*$ ,  $v(t)$ ,  $\rho$ , and  $\sigma$  to the longer-term options. We then wait for 10 or 30 days and use those parameters to forecast the option prices at those times.

### **Empirical Results and Discussion**

Sometimes it is not possible to wait for 10 or 30 days because there are no data available. In this case, for the shorter wait we search for the closest wait to 10 days between 7 and 13 days; for the longer wait we search for the closest wait to 30 days between 20 and 40 days. Tables VI and VII exhibit the pre-crash and post-crash pricing errors for the various models, based on knowing current longer-term option prices only.

Table VI about here

Table VII about here

Table VI confirms the result that in the pre-crash world with about flat smiles, the lognormal Black-Scholes model performs well, with a 38 cent error for 10 day predictions and with a 72 cent error for 30 day predictions. All other models collapse into the nested special case of the Black-Scholes model.

In the post-crash period in Table VII, the Black-Scholes model and the restricted CEV model perform worst, with 172 and 136 cent errors for 10 day predictions and 172 and 145 cent errors for 30 day predictions. The implied trees trail the relative smile, the unrestricted CEV, and the stochastic volatility models, which perform equally well. The best model is the absolute smile, with a 44 cent error for 10 day predictions and a 63 error for 30 day predictions. To investigate this rather surprising result, we graph results for four models in Figure 4.

Figure 4 about here

We can see that the absolute smile model incorporates the fact that index level and implied volatility are negatively correlated. Thus, as the index goes up over the next 30 days, the absolute smile updates the future at-the-money volatility with the current volatility of an option with the same strike price as the future stock price. Such an option is currently an out-of-the-money call option. Its implied volatility tends to be lower than the current at-the-money volatility because of the negative slope of the typical post-crash smile.

We can tie the CEV model nicely into this framework. We have established already that the restricted CEV model yields (and predicts) smiles similar to the Black-Scholes model. In addition, regressing the error on the unrestricted CEV forecasts against the relative smile model yields a correlation of 0.997. Thus, those two models behave almost identically.

### **Incorporating Known Current Shorter-Term Option Prices**

The initial notion was that we could improve our forecasts of option prices by extending the models to incorporate both current shorter-term and longer-term option prices. For the Black-Scholes formula and the relative and absolute smile models, we combine the shorter-term and longer-term implied volatilities in order to find the forward term structure of implied volatilities, and we base our forecasts on this term structure. For the CEV and the stochastic volatility models, the parameters are chosen so as to minimize the pricing error of the shorter-term and longer-term options simultaneously. These optimal parameters are then used for the forecasts.

We again use the generalized binomial tree, where we match the shorter-term options by minimizing the pricing error through choosing a weight  $w(0.5)$  between 0 and 1.<sup>15</sup> We can then use the conditional probability tree to forecast option prices and measure the pricing error.

It turns out that adding the information incorporated in shorter-term options does not help us in forecasting future option prices. The results are 1 to 4 cents worse than the pricing errors reported in tables VI and VII, where we use only the longer-term options. The only model that uses the additional information in a meaningful way is the generalized binomial tree in the post-crash period, which drops from an error of 67 cents for 10-day predictions to an error of 59 cents.

### **Looking Forward (Known Term Structure of At-the-Money Volatilities)**

To continue the test of relation (2), we investigate again the same options but at different times. We then incorporate a known term structure of at-the-money future volatilities. This allows us to untangle the pricing errors related to the models from the impact of the shift in future at-the-money implied volatility. We use a simple vertical shift to each model's forecasted volatility smile to fit the future at-the-money implied volatility correctly.

Tables VIII and IX present the pricing errors for the alternative models: Table VIII the pre-crash pricing errors, and Table IX the post-crash pricing errors.

Table VIII about here

Table IX about here

For the pre-crash period reported in Table VIII, assuming a known term structure of future at-the-money volatilities, all methods perform very similarly. Pricing errors are around 21 cents for 10- and 30-day predictions.

More interesting is the post-crash period reported in Table IX. The Black-Scholes and the restricted CEV model are predictably the worst-performing models. All other models, except for the generalized binomial tree, perform similarly well, with 22-27 cent errors for the 10-day predictions and 30-33 cent errors for the 30-day predictions.

---

<sup>15</sup> An extension to this method is to choose two points on the weight function, say,  $w(0.33)$  and  $w(0.66)$ , and again get the best fit with the shorter-term options. This additional complexity of the weight function does not improve the performance of the binomial trees, so we use the simple adjustment  $w(0.5)$  throughout.

We also investigate the median pricing error outside the bid/ask spread, where we still assume that we know the future at-the-money implied volatility. We do not report the details, as all these errors are rather small: 7 cents pre-crash and 12 cents post-crash for the better-performing models. This is again an indication that, after incorporating the level of the future smile, there is little advantage in one model over another. It seems that most models can generate the general shape of the smiles equally well, except for the rigid Black-Scholes and the restricted CEV models. All models seem to struggle with getting the level of the smile right, however.

This work suggests that we could improve the performance of all option pricing models if we were able to develop good forecasts of future implied volatility. Some preliminary studies reveal several problems.

Implied volatilities are very persistent, so we investigate models in levels and in differences of implied volatilities. The most important effect is the negative relation to log returns, but this effect is already captured in the better-performing models; they all build in the negative correlation between the stock price and volatility.

We also investigate the explanatory power of various macroeconomic factors such as default risk (credit spread), change in the term structure of interest rates, inflation, industrial production, and consumption growth. None of these factors has explanatory power. Nor do NYSE trading volume or interest rates.

Finally, we use time-to-expiration, past index levels (1, 30, and 60 days), past implied volatility levels (1, 30, and 60 days), past log returns (30 and 60 days), and day-of-the-week dummies. Again, none of these variables has explanatory power. Neither does it help to take logarithms of times-to-expiration, index levels, or implied volatilities.

This suggests an “efficient markets results” for implied volatilities. Today’s implied volatility seems to incorporate all the information publicly available at that time. Such result, while making it difficult to improve on option pricing models, does make sense economically; implied volatilities are nothing but scaled option prices, and option prices should not be predictable in efficient markets.

## **Conclusion**

We have asked what we can learn from option prices about the risk-neutral stochastic process of the S&P 500 index. This question has been partially addressed by others, but we believe we are



the first to compare implied tree models with parametric models and naïve trader rules and for pre- and post-crash subperiods.

In our first test, we choose the parameters of nine different option valuation models to best fit the observed market prices of longer-term options. We then price shorter-term options using those parameters, assuming the term structure of at-the-money volatilities is unknown. If the term structure of volatilities is known, we allow one parameter to be chosen so that the at-the-money shorter-term option is priced properly. In a second test, we calibrate our models to the longer-term options and predict 10- and 30-day future volatility smiles. We rerun the test where we assume that we know the future at-the-money implied volatility.

First, in the pre-crash period, all models match the performance of the Black-Scholes model. The reason is that the volatility smiles were almost flat pre-crash, and the additional flexibility of the more complex models is not needed.

Second, in the post-crash period, the naïve trader rules perform best. In the backward-looking test the relative smile model has a median error of 51 cents, and the absolute smile model an error of 54 cents. In the forward-looking test, the absolute smile model has a median error of 44 cents. As the standard deviation of these errors is rather large, it is difficult to distinguish the different models economically. This could well be why options traders use the naïve methods instead of the more complicated academic models.

Third, all models, except the Black-Scholes and the restricted CEV model, perform very similarly. The incorporation of the inverse relation between index level and volatility achieves this superior performance. Beyond the models however, there is considerable information incorporated in the level of the smile, as all models improve significantly when we add information about the at-the-money implied volatility of the shorter-term or the future option. It would seem fruitful to devote more research to incorporation of future at-the-money implied volatility into option pricing models. Our attempts at modeling future at-the-money implied volatility suggest an “efficient markets result” which could make it impossible to forecast these future volatilities. Such result fundamentally limits the possibilities to improve on option pricing models.

## **Acknowledgements**

The authors gratefully acknowledge a research grant from the Q-Group. For helpful comments the authors would like to thank an anonymous referee, David Brown, Jim Hodder, David Modest and seminar participants at the AFA meetings 1996, Berkeley Program in Finance 1998, and at Erasmus, Konstanz, Warwick, Dartmouth, Iowa, UBC, Madison, and Oxford.

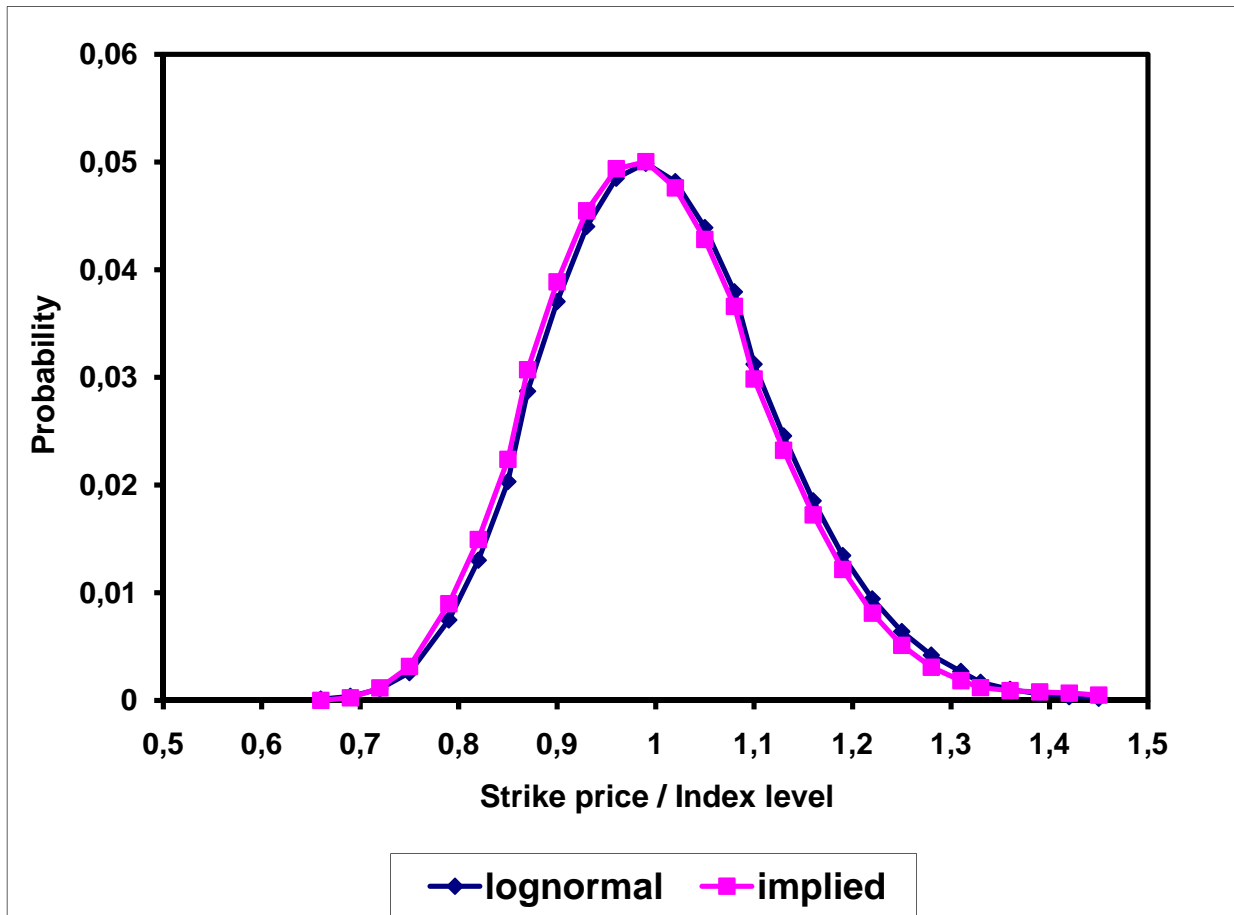
## References

- Ait-Sahalia, Y., Lo, A. W. (1998). Nonparametric estimation of state-price densities implied in financial asset prices. *Journal of Finance* 53 (2), 499-547.
- Bakshi, G., Cao, C., Chen, Z. (1997). Empirical performance of alternative option pricing models. *Journal of Finance* 52 (5), 2003-2049.
- Barle, S., Cakici, N. (1998). How to grow a smiling tree. *Journal of Financial Engineering* 7 (2), 127-146.
- Bates, D. S. (1996). Testing option pricing models, in G. S. Maddala and C. R. Rao, eds: *Handbook of Statistics, Vol. 15: Statistical Methods in Finance* (North Holland, Amsterdam), 567-611.
- Bates, D. S. (2000). Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics* 94 (1-2), 181-238.
- Beckers, S. (1980). The constant elasticity of variance model and its implications for option pricing. *Journal of Finance* 35 (3), 661-673.
- Black, F., Scholes M. (1973). Pricing of options and corporate liabilities. *Journal of Political Economy* 81 (3), 637-659.
- Christie, A. A. (1982). The stochastic behavior of common stock variances. *Journal of Financial Economics* 10 (4), 407-432.
- Cox, J. C. (1996). The constant elasticity of variance option pricing model. *Journal of Portfolio Management* 23, special issue, 15-17.
- Dennis, P., Mayhew, S. (2000). Implied volatility smiles: Evidence from options on individual equities. Working paper, University of Virginia.
- Derman, E., Kani, I. (1994). Riding on a smile. *Risk* 7 (2), 32-39.
- Duan, J.-C. (1995). The GARCH option pricing model. *Mathematical Finance* 5 (1), 13-32.
- Dumas, B., Fleming J., Whaley R. E. (1998). Implied volatility functions: Empirical tests. *Journal of Finance* 53 (6), 2059-2106.
- Dupire, B. (1994). Pricing with a smile. *Risk* 7 (1), 18-20.
- Gibbons, M., Jacklin, C. (1988). CEV diffusion estimation. Working paper, Stanford University.
- He, H., Leland, H. (1993). On equilibrium asset price processes. *The Review of Financial Studies* 6 (3), 593-619.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6 (2), 327-343.
- Heston, S. L., Nandi, S. (2000). A closed-form GARCH option valuation model. *Review of Financial Studies* 13 (3), 585-625.
- Jackwerth, J. C. (1997). Generalized binomial trees. *Journal of Derivatives* 5 (2), 7-17.
- Jackwerth, J. C. (1999). Option implied risk-neutral distributions and implied binomial trees: A literature review. *Journal of Derivatives* 7 (2), 66-82.

- Jackwerth, J. C. (2000). Recovering risk aversion from option prices and realized returns. *Review of Financial Studies* 13 (2), 433-451.
- Jackwerth, J. C., Rubinstein, M. (1996). Recovering probability distributions from option prices. *Journal of Finance* 51 (5), 1611-1631.
- Kelly, M. A. (1994). The U.S. equity market has become more diversified. J. P. Morgan Equity Derivative Strategies, New York.
- Li, M., Pearson, N. D. (2008). A horse race among competing option pricing models using S&P 500 index options. Working Paper, Georgia Institute of Technology and University of Illinois at Urbana-Champaign.
- Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica* 41 (5), 867-887.
- Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3 (1-2), 125-144.
- Osborne, M. F. M. (1959). Brownian motion in the stock market. *Operations Research* 7 (2), 145-173.
- Reiner, E. (1998). Volatility rules and implied processes, presented at the RISK conference of New Advances in Derivatives Modelling and Analysis. Warburg Dillon Read LLC, UBC, Equity Structured Products.
- Rubinstein, M. (1983). Displaced diffusion option pricing. *Journal of Finance* 38 (1), 213-217.
- Rubinstein, M. (1985). Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active CBOE option classes from August 23, 1976 through August 31, 1978. *Journal of Finance* 40 (2), 455-480.
- Rubinstein, M. (1994). Implied binomial trees. *Journal of Finance* 49 (3), 771-818.
- Tompkins, R. G. (2001). Implied volatility surfaces: Uncovering the regularities for options on financial futures. *The European Journal of Finance* 7 (3), 198 - 230.

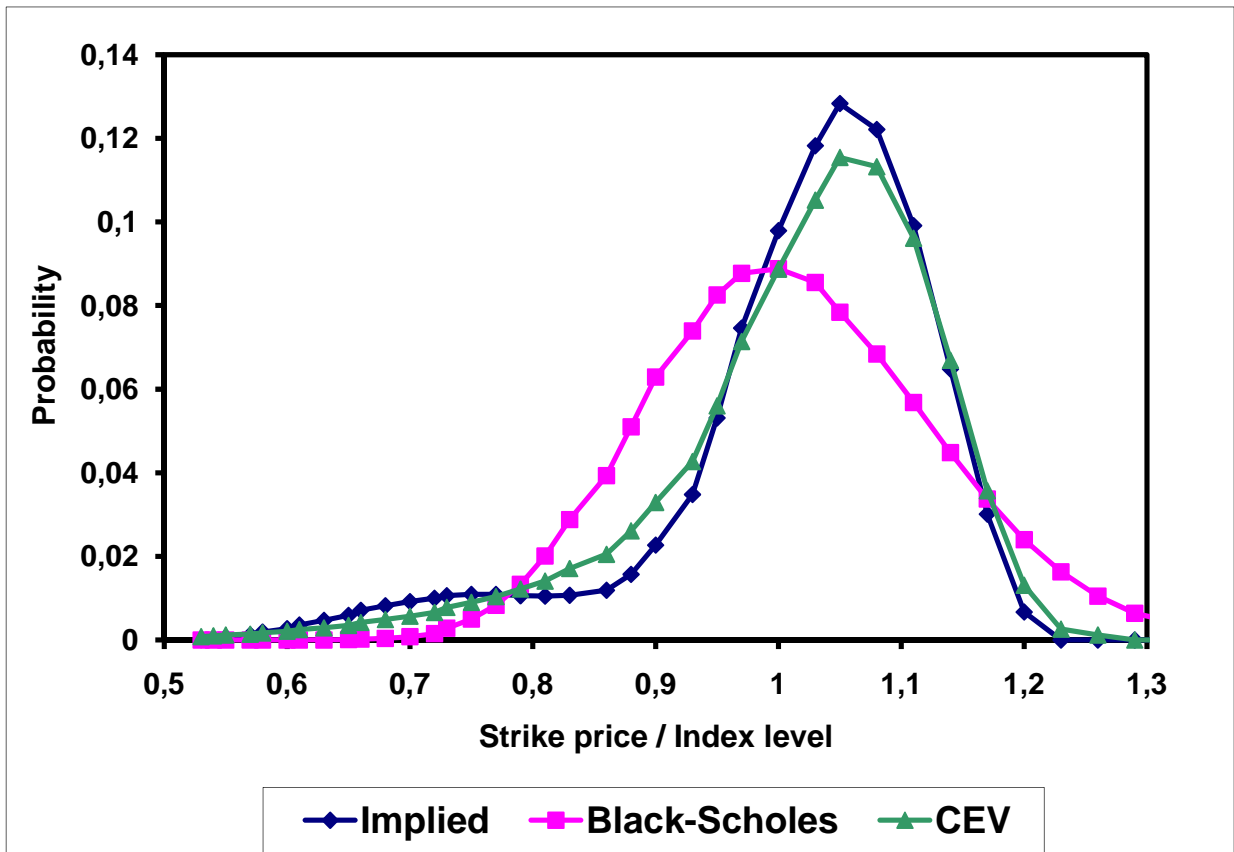
**Figure 1**  
**Lognormal and Implied Probability Distributions**

In this figure, we graph a lognormal distribution with an annualized volatility of 0.1602, which perfectly prices a set of 15 hypothetical options with strike prices in the range from 0.86 to 1.08. Given those option prices, we can derive the corresponding implied probability distribution according to the maximum smoothness criterion, graphed as well. The interest rate and dividend yield are chosen to match the observed options with 189 days-to-expiration on March 16, 1990.



**Figure 2**  
**Implied, Lognormal, and CEV Probability Distributions**


For the options with time-to-expiration of 189 days on March 16, 1990, we graph the implied probability distribution, the lognormal distribution based on the volatility of the at-the-money option, and the CEV model implied distribution where the (unrestricted) CEV parameters are chosen to fit the observed option prices best.



**Figure 3**

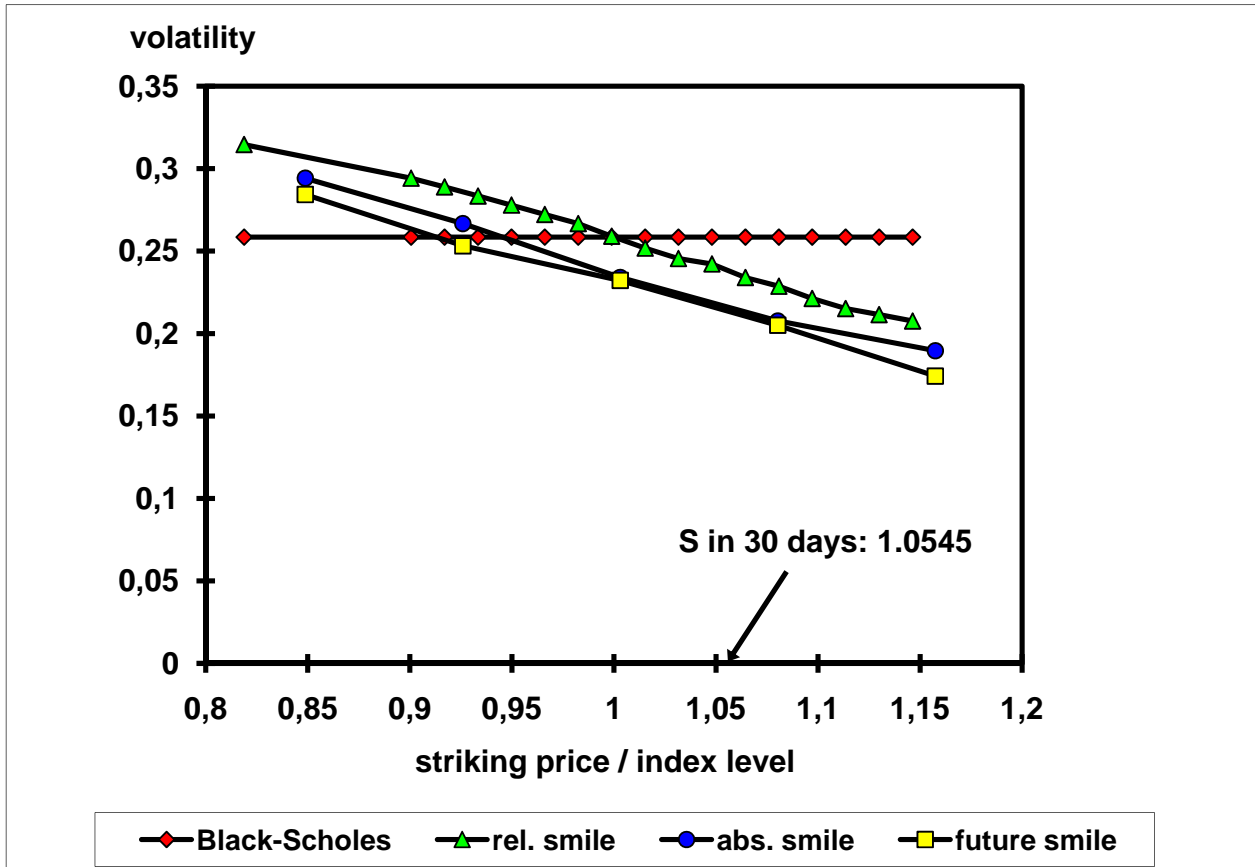
**Piecewise Linear Weight Function for a Generalized Binomial Tree**

In a standard binomial tree, the weight function is the diagonal of the unit square. A generalized binomial is parameterized by any continuous function across the corners. Here, we use a piecewise linear specification that is determined by the weight at the halfway point.



**Figure 4**  
**Observed Future Smile and Predictions**  
**Black-Scholes, Relative Smile, and Absolute Smile Models**

For options with time-to-expiration of 150 days on October 16, 1990, we graph the predicted future smiles on November 15, 1990, 30 days into the future. The predictions are based on the Black-Scholes, the relative smile, and the absolute smile model. We also graph the future smile as we observe it 30 days later.





**Table I**  
**Daily Notional Volume**

Daily notional volume (in million \$) across moneyness and time, based on 252 trading days per year. Time-to-expiration is 135-225 days.

Year	----- Moneyness -----							
	less than 0.91		0.91 – 0.97		0.97 – 1.03		more than 1.03	
	Calls	Puts	Calls	Puts	Calls	Puts	Calls	Puts
1986-87	84	247	53	152	78	76	93	44
1988-89	16	433	29	231	118	296	284	85
1990-91	32	942	41	877	637	700	829	345
1992-93	20	1534	73	1113	495	629	999	103
1994-95	53	2242	109	1937	613	1954	1401	233

**Table II**  
**Pricing Errors for Pre-Crash Period**  
**(Unknown Term Structure of Volatilities)**

This table uses seven methods to price shorter-term options, given concurrent longer-term option prices. The sample covers the pre-crash period from April 2, 1986, through October 16, 1987, where there are 360 observations.

Pricing method	----- Cent Error -----		
	Mean	Median	Std. Dev.
Black-Scholes formula	27	21	21
Relative smile prediction	27	21	21
Absolute smile prediction	28	22	21
CEV restricted formula	26	20	21
CEV unrestricted formula	26	20	21
Implied binomial tree	29	22	21
Stochastic volatility	27	21	22

**Table III**  
**Pricing Errors for Post-Crash Period**  
**(Unknown Term Structure of Volatilities)**

This table uses seven methods to price shorter-term options, given concurrent longer-term option prices. The sample covers the post-crash period from June 1, 1988, through December 31, 1994, where there are 1278 observations.

Pricing method	----- Cent Error -----		
	Mean	Median	Std. Dev.
Black-Scholes formula	126	125	40
Relative smile prediction	54	51	24
Absolute smile prediction	57	54	24
CEV restricted formula	110	110	41
CEV unrestricted formula	75	57	57
Implied binomial tree	82	78	37
Stochastic volatility	78	73	61

**Table IV**  
**Pricing Errors for Pre-Crash Period**  
**(Known Term Structure of At-the-Money Volatilities)**

This table uses seven methods to price shorter-term options, given concurrent longer-term option prices. The sample covers the pre-crash period from April 2, 1986, through October 16, 1987, where there are 360 observations.

Pricing method	----- Cent Error -----		
	Mean	Median	Std. Dev.
Black-Scholes formula	16	13	12
Relative smile prediction	17	15	12
Absolute smile prediction	18	16	12
CEV restricted formula	15	12	11
CEV unrestricted formula	15	12	11
Generalized binomial tree	28	21	21
Stochastic volatility	17	15	14

**Table V**  
**Pricing Errors for Post-Crash Period**  
**(Known Term Structure of At-the-Money Volatilities)**

This table uses seven methods to price shorter-term options, given concurrent longer-term option prices. The sample covers the post-crash period from June 1, 1988, through December 31, 1994, where there are 1278 observations.

Pricing method	----- Cent Error -----		
	Mean	Median	Std. Dev.
Black-Scholes formula	124	121	40
Relative smile prediction	43	42	17
Absolute smile prediction	43	42	17
CEV restricted formula	102	101	38
CEV unrestricted formula	43	40	26
Generalized binomial tree	42	38	27
Stochastic volatility	51	43	58

**Table VI**  
**Pricing Errors for Pre-Crash Period**  
**(Unknown Term Structure of Volatilities)**

This table uses seven methods to predict future option prices, given current longer-term option prices and the future underlying asset price. The cent errors correspond to evaluations of options closest to 10 (+/-3) days into the future, the cent errors in parentheses to evaluations of options closest to 30 (+/- 10) days into the future. The sample period covers the pre-crash period from April 2, 1986, through September 16, 1987, where there are 363 (363) observations.

Forecasting Method	----- Cent Error -----					
	Mean		Median		Std. Dev.	
Black-Scholes formula	50	(81)	38	(72)	36	(54)
Relative smile prediction	51	(82)	42	(71)	34	(52)
Absolute smile prediction	52	(84)	42	(73)	35	(54)
CEV restricted formula	49	(80)	40	(71)	35	(53)
CEV unrestricted formula	50	(80)	40	(72)	35	(53)
Implied binomial tree	54	(87)	44	(69)	40	(64)
Stochastic volatility	50	(81)	40	(71)	34	(54)

**Table VII**  
**Pricing Errors for Post-Crash Period**  
**(Unknown Term Structure of Volatilities)**

This table uses seven methods to predict future option prices, given the current longer-term option prices and the future underlying asset price. The cent errors correspond to evaluations of options closest to 10 (+/-3) days into the future, the cent errors in parentheses to evaluations of options closest to 30 (+/- 10) days into the future. The sample period covers the post-crash period from June 1, 1988, through November 30, 1994, where there are 1553 (1553) observations.

Forecasting Method	----- Cent Error -----					
	Mean		Median		Std. Dev.	
Black-Scholes formula	175	(181)	172	(172)	58	( 71)
Relative smile prediction	73	(102)	56	( 78)	61	( 83)
Absolute smile prediction	56	( 76)	44	( 63)	43	( 55)
CEV restricted formula	139	(154)	136	(145)	54	( 68)
CEV unrestricted formula	74	(102)	56	( 77)	60	( 83)
Implied binomial tree	83	(116)	67	( 97)	62	( 79)
Stochastic volatility	75	(111)	57	( 83)	66	(132)

**Table VIII**  
**Pricing Errors for Pre-Crash Period**

**(Known Term Structure of Future At-the-Money Volatilities)**

This table uses seven methods to predict future option prices, given current longer-term option prices, the future at-the-money volatility, and the future underlying asset price. The cent errors correspond to evaluations of options closest to 10 (+/-3) days into the future, the cent errors in parentheses to evaluations of options closest to 30 (+/- 10) days into the future. The sample period covers the pre-crash period from April 2, 1986, through September 16, 1987, where there are 353 (353) observations.

Forecasting Method	----- Cent Error -----					
	Mean		Median		Std. Dev.	
Black-Scholes formula	21	(20)	19	(18)	11	(10)
Relative smile prediction	26	(26)	22	(23)	14	(14)
Absolute smile prediction	27	(28)	23	(24)	15	(16)
CEV restricted formula	22	(21)	19	(19)	12	(11)
CEV unrestricted formula	22	(21)	19	(19)	13	(11)
Generalized binomial tree	24	(24)	21	(21)	13	(13)
Stochastic volatility	24	(23)	20	(20)	13	(13)



**Table IX**  
**Pricing Errors for Post-Crash Period**  
**(Known Term Structure of Future At-the-Money Volatilities)**

This table uses seven methods to predict future option prices, given current longer-term option prices, the future at-the-money volatility, and the future underlying asset price. The cent errors correspond to evaluations of options closest to 10 (+/-3) days into the future, the cent errors in parentheses to evaluations of options closest to 30 (+/- 10) days into the future. The sample period covers the post-crash period from June 1, 1988, through November 30, 1994, where there are 1267 (1267) observations.

Forecasting Method	----- Cent Error -----					
	Mean		Median		Std. Dev.	
Black-Scholes formula	163	(156)	167	(159)	45	(43)
Relative smile prediction	30	( 37)	23	( 31)	24	(23)
Absolute smile prediction	30	( 37)	23	( 31)	24	(23)
CEV restricted formula	125	(122)	131	(128)	41	(39)
CEV unrestricted formula	32	( 37)	27	( 33)	21	(19)
Generalized binomial tree	34	( 53)	28	( 51)	24	(27)
Stochastic volatility	32	( 38)	22	( 30)	31	(49)