BERNARD DUMAS*
JEFF FLEMING**
ROBERT E. WHALEY***

IMPLIED VOLATILITY FUNCTIONS: EMPIRICAL TESTS

ABSTRACT

Black and Scholes (1973) implied volatilities tend to be systematically related to the option's exercise price and time to expiration. Derman and Kani (1994), Dupire (1994), and Rubinstein (1994) attribute this behavior to the fact that the Black/Scholes constant volatility assumption is violated in practice. These authors hypothesize that the volatility of the underlying asset's return is a deterministic function of the asset price and time. Since the volatility function in their model has an arbitrary specification, the deterministic volatility (DV) option valuation model has the potential of fitting the observed cross-section of option prices exactly. Using a sample of S&P 500 index options during the period June 1988 and December 1993, we attempt to evaluate the economic significance of the implied volatility function by examining the predictive and hedging performance of the DV option valuation model.

Discussion draft: September 8, 1995

*Professor of Finance, HEC School of Management and Research Professor of Finance, Fuqua School of Business, Duke University, **Assistant Professor of Administrative Science, Jones Graduate School of Administration, Rice University and ***T. Austin Finch Foundation Professor of Business Administration, Fuqua School of Business, Duke University.

This research was supported by the Futures and Options Research Center at the Fuqua School of Business, Duke University. We gratefully acknowledge discussions with Jens Jackwerth and Mark Rubinstein and comments and suggestions by Peter Boessarts, Peter Carr, Jin-Chuan Duan, Denis Talay, and the participants at the Isaac Newton Institute, Cambridge University, in April 1995 and at the Chicago Board of Trade's Nineteenth Annual Spring Research Symposium, Chicago, Illinois, in May 1995.

IMPLIED VOLATILITY FUNCTIONS: EMPIRICAL TESTS

Claims that the Black and Scholes (1973) valuation formula no longer holds in financial markets have recently appeared with increasing frequency. When the Black/Scholes formula is inverted to imply volatilities from observed option prices, the volatility estimates differ across exercise prices and times to expiration.[1] For S&P 500 index option prices prior to the October 1987 market crash, for example, the implied volatilities form a "smile" pattern. Options that are deep in-the-money or out-of-the-money have higher implied volatilities than at-the-money options. After the crash, the pattern changed — implied volatilities of S&P 500 options decrease monotonically as the option exercise price rises relative to the current index level, with the rate of decrease increasing for options with shorter time to expiration.

The failure of the Black/Scholes model to describe the observed structure of option prices is thought to arise from its constant variance assumption.[2] Casual empiricism suggests that when stock prices go up volatility goes down, and vice versa. Accounting for nonconstant volatility within an option valuation framework, however, is no easy task. With stochastic volatility, option valuation generally requires a market price of risk parameter, which, among other things, is difficult to estimate. An exception to this valuation problem occurs where the volatility of the underlying asset's return is a deterministic function of asset price and/or time. In this case, it remains possible to value options based on the Black/Scholes partial differential equation although not by means of the Black/Scholes formula itself.

Derman and Kani (1994), Dupire (1994), and Rubinstein (1994) adopt variations of a deterministic volatility (DV) option valuation model. Rather than positing a structural form for the volatility function $\sigma(S,t)$, however, these authors suggest searching for a

[1] Rubinstein (1994) examines the S&P 500 index option market. Smile investigations have also been performed for the Philadelphia Exchange's foreign currency option market (e.g., Taylor and Xu (1993)), and for stock options traded at the LIFFE (e.g., Duque and Paxson (1993)) and the European Options Exchange (e.g., Heynen (1993)).

[2] Indeed, Black (1976, p. 177) succinctly states that "... if the volatility of a stock changes over time, the option formulas that assume a constant volatility are wrong."

binomial or trinomial lattice that achieves an *exact* cross-sectional fit of reported option prices. Rubinstein (1994), for example, uses an “implied binomial tree” approach. Current option prices are used to build a binomial tree with branches at each node that are designed (either by choice of up-and-down increment sizes or probabilities) to reflect the time-variation of volatility. With as many degrees of freedom as option prices, these models can obtain an exact fit of the observed structure of option prices at time *t*. Unfortunately, none of these studies examine the stability of the implied tree through time. In particular, if the implied binomial tree estimated at time *t* is correct, it should include the implied tree estimated at time *t*+1.

The purpose of this paper is to assess the stability of the implied volatility function $\sigma(S,t)$. A stable volatility function lends credibility to the DV approach. In this case, the DV framework should provide a better means of setting hedge ratios and valuing exotic options. On the other hand, if the function is not stable, it cannot be claimed that the true volatility function of the underlying asset price has been identified.

The paper is organized as follows. In Section I, we review option valuation under deterministic volatility and outline our implied volatility function estimation procedure. In Section II, we describe our sample of S&P 500 index option prices and document Black/Scholes implied volatility patterns that have appeared during the past six years. In Section III, we estimate cross-sectionally the implied volatility functions and describe the model’s goodness-of-fit. In Section IV, we assess how well the volatility function estimated at time *t* predicts option prices one week later. In Section V, we assess hedging performance. The study concludes in Section VI with a summary.

I. Option Valuation Under Deterministic Volatility

The valuation of options under the assumption that the volatility rate is a deterministic function of asset price and time is presented in this section. We begin by providing the Black/Scholes partial differential equation (PDE). Although the structure of the PDE remains the same under deterministic volatility, analytical formulas are generally not possible and approximation methods must be used. Since our objective is to infer the volatility function from a cross-section of option prices, we re-write the Black/Scholes backward equation as a forward equation to speed the estimation. Finally, we describe in detail our procedure for estimating the volatility function.

A. Black/Scholes valuation

Assuming the volatility rate of the underlying asset is a deterministic function of asset price and time, the Black/Scholes (1973) partial differential equation can be written

$$rc - rS\frac{\partial c}{\partial S} - \frac{1}{2}\sigma^2(S,t)S^2\frac{\partial^2 c}{\partial S^2} = \frac{\partial c}{\partial t} \qquad (1)$$

where r is the riskless rate of interest, S is the asset price, c is the option price, $\sigma^2(S,t)$ is the *local* volatility function, and t is time. The equation applies both to calls and puts and to European- and American-style options. What distinguishes the valuations are the boundary conditions. For a European-style call option, for example, the boundary condition, $c(S,T) \equiv \max(S - X, 0)$, is applied at the option's expiration. In the special case where the volatility rate is constant, $\sigma(S,t) \equiv \sigma$, the value of a European-style call can be obtained analytically, with the resulting formula being known as the "Black/Scholes formula." In general, however, valuation formulas cannot be obtained although option valuation remains possible through the use of numerical procedures. Rubinstein (1994) uses a binomial lattice approach, while Dupire (1994) recommends a trinomial approach. For reasons that will be discussed later, we use the Crank-Nicholson finite difference procedure.

B. Specifying the forward equation

The partial differential equation (1) is the backward equation of the Black/Scholes model. The call option value is a function of S and t for a fixed exercise price X and time to expiration T. At time t when S is known, however, we can equivalently consider the cross-section of option prices (with different exercise prices and times to expiration) to be functionally related to X and T. As Dupire (1994) shows, this means that the value of a European-style call option, $c(X,T)$, can be written as the forward partial differential equation,[3]

[3] The option price, c, and the underlying asset price, S, are taken to be forward prices (forward to the maturity date of the option). For that reason, equation (1) ignores interest and dividends, which are taken into account in the definition of forward prices.

$$\frac{1}{2}\sigma^2(F,T)X^2\frac{\partial^2 c}{\partial X^2}=\frac{\partial c}{\partial t} \quad (2)$$

with the associated initial condition, $c(X,0)\equiv\max(S-X,0)$. The advantage of the forward equation approach is that all option series with a common time to expiration can be valued simultaneously. Alternatively, we could have solved the Black/Scholes PDE as many times as we have options of different exercise prices and reached the same numerical results. Indeed, the backward equation approach would be necessary where one wanted to infer the volatility function from American-style option prices.

C. Estimation of the volatility function

We estimate the volatility function $\sigma(X,T)$ by fitting the option valuation model to the observed structure of option prices at time *t*. At this juncture, however, $\sigma(X,T)$ is an arbitrary function. We posit a number of different structural forms for $\sigma(X,T)$ including

$$\begin{aligned}
&\text{Model 0: } \sigma=a_0\\
&\text{Model 1: } \sigma=a_o+a_1X+a_2X^2\\
&\text{Model 2: } \sigma=a_o+a_1X+a_2X^2+a_3T+a_5XT\\
&\text{Model 3: } \sigma=a_o+a_1X+a_2X^2+a_3T+a_4T^2+a_5XT
\end{aligned} \quad (3)$$

Model 0 is the Black/Scholes model where the volatility rate is constant. Model 1 attempts to capture variation with the asset price, and Models 2 and 3 capture additional variation with respect to time.

In using parsimonious volatility function structures such as Models 1 through 3, our approach does not guarantee that the fitted, theoretical values match the observed option prices. As such, it can be distinguished from the approaches suggested by Derman and Kani (1994), Dupire (1994) and Rubinstein (1994). These authors, in various ways, suggest building a binomial or trinomial tree with variable volatility. Since these trees each contain as many degrees of freedom as there are option prices available to be fitted, an exact fit of quoted option prices can be obtained. The question is, however, whether or not the observed structure of option prices has been overfitted—a question that can be answered only by moving out of sample. If the volatility function is constant through time, as is assumed by the model, the implied tree at time *t* should contain the tree implied at time *t*+1. Alternatively, in terms of our approach, the volatility function estimated at time *t*+1 should have the same coefficients as the function estimated at time *t*. While the restricted functional form deteriorates the quality of the fit at time *t*, there is no reason to believe that it would deteriorate the quality of the prediction to time *t* plus one week, relative to the alternative tree-based procedure which is less parsimonious in terms of the number of parameters to be estimated.

Finally, we must qualify the use of our approach. Our overall procedure is not a testing procedure. The "null hypothesis" being investigated is that volatility is an exact function of asset price and time, so that options can be priced exactly by the no-arbitrage condition. Evidently, any deviation from such a strict theory, no matter how small, should cause a test statistic to reject it. No source of error has been postulated that makes a deviation from the theory tolerable.[4] If a source of error had been introduced, some restriction on the sampling distribution of the error could have been deduced. That restriction could have been the basis for a testing procedure.[5] Barring such an approach, for which no consensus exists in the literature, the empirical procedure employed here will have to remain a descriptive one. Whether prediction errors are large or not is a matter of judgment.

II. S&P 500 Option Prices and Implied Volatility Smiles

S&P 500 index options serve as the basis of our empirical analysis. In this section, we describe the data used in our analyses and document the commonly-observed pattern in Black/Scholes implied volatilities.

[4] The same difficulty arises in the empirical verification of an exact theory. See MacBeth and Merville (1979), Whaley (1982), and Rubinstein (1985).

[5] Jacquier and Jarrow (1995) introduce two kinds of errors in the Black/Scholes model: model error and market error, which they distinguish by assuming that market errors occur rarely. Other approaches to the problem include Lo (1986) who introduces parameter uncertainty, Clément, Gouriéroux and Montfort (1993) who randomize the martingale pricing measure, and Bossaerts and Hillion (1994) whose error is due to discreteness in hedging.

A. Data Selection

Our sample contains observed prices of S&P 500 index options traded on the Chicago Board Options Exchange (CBOE) during the period June 1988 through December 1993.[6] S&P 500 index options are European-style and expire on the third Friday of the contract month. Originally, S&P 500 options traded at the CBOE expired only at the close of trading on the expiration day and were denoted by the ticker symbol SPX. When the Chicago Mercantile Exchange (CME) changed the expiration of their S&P 500 futures contract from the close to the open in June 1987, the CBOE introduced a second set of S&P 500 options with the ticker symbol NSX that expired at the open along with the futures. At the outset, the trading volume of the S&P 500 "open-expiry" option series was considerably lower than the "close-expiry" options. Over time, however, the trading volume grew and eventually exceeded that of the close-expiry options. On August 24, 1992, the CBOE reversed the ticker symbols of the two sets of options. Our sample contains SPX options throughout: close-expiry options until August 24, 1992 and open-expiry options afterward. During the first subperiod, the option's time to expiration is measured as the number of calendar days between the trade date and the expiration date; during the second, the number of days to expiration is the number of calendar days remaining less one.

In the last section, we argued that it is more efficient to estimate the implied volatility function using forward prices rather than spot prices. To carry out this estimation, therefore, we must transform observed spot prices for the cash index and option series into forward prices. For the index level, this means that we require both the term structure of interest rates as well as the daily cash dividends over the life of the option. To proxy for short-term riskless interest rates, we use the T-bill rates implied by the average of the bid and ask discounts. The history of these discounts was collected from the *Wall Street Journal*. The entire term structure was collected each day. The riskless rate corresponding to an option with time to expiration, T, is the rate obtained by interpolating the rates of the two T-bills whose maturities straddle the option expiration.

[6] The sample begins June 1988 since it was the first month that Standard and Poors began reporting daily cash dividends for the S&P 500 index portfolio. See Harvey and Whaley (1992b) regarding the importance of incorporating discrete daily cash dividends in index option valuation.

The daily cash dividends for the S&P 500 index portfolio were collected from the *S&P 500 Information Bulletin*. To compute the present value of the dividends paid during the option's life, *PVD*, the daily dividends are discounted at the rates corresponding to the ex-dividend dates and summed over the life of the option, that is,

$$PVD = \sum_{i=1}^{n} D_i e^{-r_i t_i} \qquad (4)$$

where D_i is the *i*-th cash dividend payment, t_i is the time to ex-dividend from the current date, r_i is the interest rate corresponding to the time to ex-dividend (interpolated from the current term structure of interest rates), and *n* is the number of dividend payments during the option's life.[7] The implied forward price of the S&P 500 index is therefore

$$F = (S - PVD)e^{rT}, \qquad (5)$$

where *T* is the time to expiration of the option. The forward price of an option is simply the current price carried forward to the option's expiration at the appropriate riskless interest rate (e.g., ce^{rT}).

The reported cash index level, *S*, in (5) is only an expositional device and is never used for valuation purposes in our analysis. Indeed, fearing imperfect synchronization between the option market and other markets,[8] we use neither the reported S&P index nor the S&P 500 futures price from the CME.[9] Instead, we infer the cash index level simultaneously with the parameters of the volatility function from the cross-section of option prices. In this way, our empirical procedure relies only on observations from a single market and not others. Consequently, no auxiliary assumption of market integration is necessary.[10]

In conducting our investigation, we estimate each model of the volatility function once each week during the sample period June 1988 through December 1993. Wednesdays are used because fewer holidays fall on Wednesday than any other trading

[7] The convention introduces an inconsistency, with small consequences, between option prices of different maturities. In constructing the forward version of the S&P 500 index level, one assumes that the dividends to be paid during the option's life are certain.
[8] See Fleming, Ostdiek and Whaley (1995).
[9] For a detailed description of the problems of using a reported index level in computing implied volatility, see Whaley (1994, Appendix).
[10] This is not quite true since we use Treasury bill rates in computing forward prices.

day. Where a Wednesday was a holiday during the sample period, the trading day immediately preceding Wednesday was used.

We estimate the volatility functions by minimizing the sum of squared errors of the observed option prices from the options' theoretical values.[11] To proxy for observed option prices, we use bid/ask price midpoints. Trade prices, by their very nature, are executed at the bid or at the ask, depending on the motivation of the last trade. For the cross-sectional estimation of the volatility function, this means that trade prices induce noise, making the estimation of the function less precise. Bid/ask midpoints, on the other hand, contain no such error.

Three exclusionary criteria were applied to the data. First, we eliminated options with less than six and with more than one hundred days to expiration. Options with less than six days to expiration have relatively small time premia, hence the estimation of volatility is extremely sensitive to nonsynchronous option prices and other possible measurement errors. Options with more than a hundred days to expiration, on the other hand, are unnecessary since our objective is only to determine whether the volatility function remains valid over a span of one week. Including longer-term options would only serve to deteriorate the cross-sectional fit.

The second exclusionary criterion filtered out deep in- and out-of-the-money options. Like in the case of extremely short-term options, deep in- and out-of-the-money options have little time premia and hence contain little information about the volatility function. In addition, these options have little trading activity, hence price quotes are generally not supported by actual trades. To operationalize this criterion, we eliminate options whose absolute "moneyness," that is, $\left|100\left(\frac{F}{X}-1\right)\right|$, is greater than ten percent.[12]

Finally, only options with bid/ask price quotes during the last half hour of trading (i.e., 2:45 to 3:15 PM (CST)) are used. Restricting ourselves to such a tight window of time is necessary because, as we explained earlier, we also imply the cash index level from

---

[11] The algorithm used for the minimization is "AMOEBA" from Press, Teukolsky, Vetterling, and Flannery (1992). The routine is based on the downhill simplex method of Nelder and Mead (1965).

[12] The moneyness variable may also be written in terms of spot prices, that is, $\left|100\left(\frac{S-PVD}{Xe^{-rT}}-1\right)\right|$.

the observed option prices. The option prices used for that purpose must, therefore, be reasonably synchronous. The cost of this criterion is, of course, that we include only a reduced number of option quotes. The cost is not too onerous, however, since we find quotes for an average of 44 call/put option series during the last half-hour each Wednesday.[13] Seventeen of the 292 Wednesday cross-sections had only one contract expiration; 141 had two; 129 had three; and five had four.

B. Volatility smiles

To illustrate the Black/Scholes "implied volatility smile", we used bid and ask price quotes for call options[14] during the 2:45-3:15 PM window on April 1, 1992 (a typical day) and computed implied volatilities[15] based on the Black/Scholes call option valuation formula,

$$c = (S - PVD)\,N(d_1) - Xe^{-rT} N(d_2)\,, \tag{6}$$

where $d_1 = [\ln((S - PVD)/Xe^{-rT}) + .5\sigma^2 T]/\sigma\sqrt{T},$ $d_2 = d_1 - \sigma\sqrt{T},$ and *N(.)* is the cumulative normal density function. The pattern of implied volatilities is displayed in Figure 1. Note that these are the Black/Scholes implied volatilities and not a graph of the volatility function $\sigma(S,t)$. The fact that they do not fall on an horizontal line is, of course, evidence that the Black/Scholes formula does not hold.

Several features in Figure 1 deserve comment. First, observe that implied volatilities corresponding to bid and ask quoted prices are closest together for options that are approximately at the money (where percent moneyness is zero). Bid and ask implied volatilities are further apart as moneyness moves away from 0, particularly to the right of

[13] To assess the reasonableness of using the 2:45-3:15 PM window for estimation, we computed the mean absolute return and the standard deviation of return of the nearby S&P 500 futures (with at least six days to expiration) by fifteen-minute interval throughout the trading day across the days of the sample period. The results indicated that the lowest mean absolute return and standard deviation of return occur just prior to noon. The end-of-day window is only slightly higher, while the beginning-of-day window is nearly double. We chose to stay with the end-of-day window for ease in interpretation of the results, although our plans are to replicate the steps of the study using the 11:30-12 noon window each day.

[14] For this exercise only, we use the reported stock index level in the estimation of volatility. Since the reported index is always stale, we use only call options. While a stale index causes the implied volatilities of the calls to be biased downward or upward depending on whether the reported index is above or below its true level, the bias for all calls will be in the same direction. With puts, the bias is opposite.

[15] For this illustration only, we did not enforce the moneyness criterion.

the figure where the call options are deep in the money. The reason is that the market for these options is less active so market makers require a larger spread.[16]

Second, the so-called "smile" is not a smile at all. The label arose prior to the October 1987 market crash when, in general, the Black/Scholes implied volatilities were symmetric around zero moneyness, with in-the-money and out-of-the-money options having higher implied volatilities than at-the-money options. The pattern displayed in Figure 1, however, is indicative of the pattern that has generally appeared since the crash, that is, call (put) option implied volatilities increasing monotonically as the call (put) goes deeper in the money (out-of-the-money).

Third, the smile tends to attenuate as time to expiration increases. For the calls with 17 days to expiration, the range of implied volatilities is from slightly more than 10 percent to nearly 30 percent. For the 45-day and 80-day calls, implied volatilities are not higher than 22 percent while having approximately the same lower bound.

III. Estimation Results

Using the S&P 500 index option data described in the previous section, we estimate the four different volatility function specifications given in (3). As was noted earlier, Model 0 is the Black/Scholes constant volatility model. Model 1 allows the volatility rate to vary with asset price but not with time. Models 2 and 3 attempt to capture time variation. A fifth model, called the "Switching Model," uses the volatility functions given by Models 1, 2 and 3, depending on whether the number of different option expirations on a given day is one, two, or three. This model is introduced due to the fact that in some of our cross-sections, there is little or no time-to-expiration variation, undermining our ability to precisely estimate the relation between the volatility rate and time. We estimate each model by minimizing the sum of squared errors between the observed option prices and their theoretical values based on the DV option valuation model. For all volatility function specifications, we truncate the volatility rate at one

[16] Spreads are generally competitively determined. The CBOE has rules governing the maximum spread allowed for options with different degrees of moneyness.

percent. Finally, to avoid possible problems with index level staleness, the cash index level is simultaneously estimated along with the volatility function's parameters.

A. Goodness-of-fit

Table 1 contains summary statistics of the estimation results across all 292 days in the sample period June 1988 through December 1993. Average root mean squared valuation error (RMSVE), average valuation error outside the bid/ask spread (AVERR), and the frequency with which the specified model has a lower RMSVE than the switching model (FreqSW) are reported.

The average RMSVE results reveal that there is a strong relation between volatility and asset price. When the volatility rate is a quadratic function asset price (Model 1), the average RMSVE of the DV option valuation model is less than half of that of the Black/Scholes constant volatility model (Model 0), .3010 vs. .6497. Time also appears to have an important effect. In moving from Model 1 to Model 2, the average RMSVE is reduced further (i.e., from .3010 to .2300), albeit not quite as dramatically. The addition of the time variable to the volatility function appears to be important, although most of the incremental explanatory power appears to come from the cross-product term, *XT*. Adding a quadratic time to expiration term (Model 3) reduces the average RMSVE to its lowest level of the assumed specifications, .2264. The switching model's RMSVE is virtually the same.

The frequency with which competing models have lower RMSVE's than the switching model supports the switching model as being the "best" of the available alternatives. Only Model 3 comes close, having a lower RMSVE in 37.3 percent of the cross-sections examined. Model 2 is next with 14.7 percent. The constant volatility model never has a lower RMSVE.

The average bid/ask spread of the option series used in our estimations is approximately 47 cents over the sample period. With such a wide trading cost band, few of the theoretical option values lie outside the range of observed bid and ask prices. The average absolute valuation error outside the bid/ask spread (i.e., the difference between the fitted value and the ask price if the fitted value exceeds the ask, the difference between the bid price and the fitted value if the fitted value is below the bid, or zero if the fitted value lies between the bid and ask prices), denoted AVERR, is less than 5 cents for the switching model. So, even a volatility function with as few as six parameters provides a cross-sectional fit that is largely within trading cost bands. Naturally, adding more parameters would eventually ensure a perfect fit.

Figures 2 and 3 show the actual and fitted prices and valuation errors by option series on that day. Figure 2 contains the actual bid/ask price midpoints and the fitted values of the options on April 1, 1992. Given the wide range of option exercise prices, the deviations from the model values appear quite small. The solid fitted value line appears to fall on the observed prices across all option series. Figure 3, on the other hand, is more informative. For the 17-day call option, for example, the valuation error is positive and largest for deep in-the-money calls and falls virtually monotonically as the calls go less and less in the money. The valuation errors for the 17-day puts, however, appear much more random. Deep out-of-the-money puts have large positive valuation errors, and the valuation errors fall as the puts become more in the money. Then, the pattern reverses, with in the money puts having increasing positive valuation errors. For the longer time to expiration options, the valuation errors are less systematic.

B. Parameter estimates

The average parameters estimated for each of the volatility function are also interesting. Model 0 is, of course, the constant volatility model of Black/Scholes. When this model was fitted each week during our 292-week sample period, the mean estimated coefficient $\hat{a}_0$ was 15.72 percent. The distribution of implied volatilities was slightly skewed to the right in the sense that the median estimated coefficient was 15.17 percent. The minimum estimate was 9.43 percent on December 29, 1993, and the maximum was 27.16 percent on January 16, 1991.

Model 3 is the least parsimonious volatility function that we consider. Figure 4 contains an illustration of the estimated volatility surface generated by Model 3 on Wednesday, April 1, 1992. On that day, there were 73 option series across three maturities that satisfied our exclusionary criteria. The estimated coefficients were:

$$\hat{a}_0 = 113.620 \qquad \hat{a}_1 = -.242512$$
$$\hat{a}_2 = -.0000148231 \qquad \hat{a}_3 = .347430$$
$$\hat{a}_4 = -.0000016749 \qquad \hat{a}_5 = .0512376$$

with an implied cash index level of 404.61. Figure 4 shows that, for a given time to expiration, the local volatility displays a smile-like pattern. As the index level falls from its current level of 404.61, volatility increases at an increasing rate. As the index level rises, volatility continues to fall, however, it eventually begins to rise once the index level rises above approximately 500. Looking at the other horizontal axis, volatility does not appear to have an extraordinary term structure shape. Indeed, the surface appears quite flat as the time to expiration of the option increases.

The coefficients estimated using the observed option prices on April 1, 1992 can also be used to deduce the shape of the risk-neutral probability distribution at the end of the options' lives. Figure 3 shows the implied distributions for the April, May, and June 1992 option maturities. Note that all of the distributions are skewed to the left. This result is opposite the right-skewness that is implicit in the Black/Scholes assumption of lognormally-distributed asset prices. The wider variance for the May and then June expirations merely reflects the fact that the longer is the option's time to expiration, the greater is the probability of large asset price moves. It is interesting to note that our implied distribution does not exhibit the bimodality of the distribution implied in Rubinstein (1994). One possible explanation for this is that our analysis assumes a more parsimonious volatility function.

IV. Prediction Results

The estimation results reported in the last section indicate that the more terms that are added to the volatility function, the better the DV option valuation model does at explaining the observed structure of option prices. A critical assumption of the model, however, is that the volatility function is deterministic and remains stable through the option's life. In this section, we evaluate how well the volatility function estimated one week values the same options one week later.

Figures 6 and 7 show the actual and fitted prices and valuation errors by option series on a typical day during our sample period, April 8, 1992. Table 2 contains summary statistics of the prediction results across all 291 days remaining in the sample. Turning first to the figures, it is interesting to note that the theoretical option values (i.e., the solid line) lie frequently outside the bid/ask price spread (i.e., the dashed lines). Figure 6 shows that this is particularly true for longer term options. Figure 7 further illustrates the time to expiration bias. In this figure, the bid and ask prices are normalized by the bid/ask price midpoint. Hence, the dashed lines are symmetric around zero. The theoretical values are also normalized by the bid/ask price midpoint. Consequently, the figure shows clearly where, and by how much, the fitted value deviated from the quotes. While the valuation errors are largest for the 80-day options, they are still quite large for the 45-day and even some of the 17-day options.

Table 2 provides the summary statistics of the prediction errors across all days in the sample. The table shows that the errors are generally quite large relative to the Table 1 results. The average RMSVE across the days in the sample is about 57 cents for all DV models except Model 0. Recall that the in-sample errors were about 23 cents. The magnitude of the error should not be surprising, however, in the sense that new market information has disseminated over the week, presumably causing the level of market volatility to be revised. Indeed, the average mean absolute valuation error (AVERR) outside the bid/ask spread is nearly 30 cents!

The average RMSVE for Model 1 is lower than the more elaborate models out of sample. One interpretation of this finding is that the more complex volatility function specifications overfit the observed structure of option prices. Out of sample, the more parsimonious models tend to have smaller errors.

The prediction errors are also categorized by moneyness and time to expiration, and summary statistics for each category are presented. The table shows that prediction errors increase with the time to expiration, consistent with our conclusions based on examining Figures 6 and 7. But, what is perhaps more interesting is that the at-the-money options have the largest prediction errors for all times to expiration. This arises because at-the-money options are the most sensitive to volatility (where time premium is highest).

When options are deep in- or out-of-the-money, volatility mismeasurement has less impact on option valuation.

What is most troubling about the analysis thus far is that, although the RMSVEs must be considered large for all practical purposes, we have no real means for evaluating what size of valuation error should truly be considered "large." One benchmark that comes to mind is the valuation error that would have been achieved by means of the Black/Scholes formula applied to a fitted volatility smile such as the ones of Figure 1. Quite evidently, applying the Black/Scholes formula to a nonconstant volatility is internally inconsistent since the Black/Scholes formula is based on an assumption of constant volatility. Nonetheless, the procedure could conceivably be used as a practical way of predicting option prices.[17] We would expect the DV option valuation model, which is based on an internally consistent specification, to represent an improvement on the Black/Scholes approach.

To verify whether that is the case, we implement on the Black/Scholes formula a two-step procedure which is similar to the one we have so far applied to the volatility-function specification. On day t, we fit the same switching volatility specification to the Black/Scholes implicit volatility smile, and then, on day t+7, we apply the Black/Scholes formula to the same smile but given the new cash index level. The valuation errors that are achieved in this fashion are also summarized in Table 2. Note that the Black/Scholes errors are almost uniformly smaller than those of the deterministic volatility approach. The average RMSVE across the entire sample period is 48 cents for the ad hoc Black/Scholes procedure, where it is nearly 56 cents for the DV (Model 1) option valuation model. In viewing the various option categories, the greatest pricing improvement appears to be for at-the-money options, whose average RMSVEs are reduced by 10 cents or more. Put simply, the deterministic volatility approach does not appear to be an improvement over the traditional, albeit inconsistent, Black/Scholes formula with changing volatility.

[17] The Black/Scholes procedure could not serve to predict American or exotic option prices from European option prices, which is the major benefit claimed for the implied volatility tree approach.

V. Hedging Results

A key motivation for developing the DV option valuation model is to provide better risk management. If volatility is a deterministic function of asset price and time, setting hedge ratios based on the DV option valuation model should present an improvement over the constant volatility model. In this section, we evaluate hedging performance. Our methodology assumes that the hedge portfolio is continuously rebalanced through time. The hedge portfolio is formed on day *t* and then is unwound one week later. Under this scheme, the hedging error is defined as

$$\varepsilon_t = \Delta c_{\text{observed},t} - \Delta c_{\text{theoretical},t}, \tag{7}$$

where $\Delta c_{\text{observed},t}$ is the change in the observed option price from day *t* until day *t*+7 and $\Delta c_{\text{theoretical},t}$ is the change in the model's theoretical value.

Table 3 contains a summary of the hedging error results. Across the overall sample period, Model 0—the Black-Scholes constant volatility model—performs best of all the deterministic volatility function specifications! Its average root mean squared hedging error (RMSHE) is .4547, compared with .4892, .5078, .5084, and .5075 for Models 1 through 3 and the switching model, respectively. The results indicate that the more parsimonious is the volatility function, the better is the hedging performance.

The ad hoc Black/Scholes procedure described in the last section also performs well from a hedging standpoint. The average RMSHE is only .4406, and it outperforms each of the DV specifications in more than 50 percent of the days examined. Consistent with the prediction results reported in Table 2, the DV option valuation model does not appear to be an improvement. Better risk management results can obtained using an ad hoc procedure.

VI. Summary and Conclusions

Claims that the Black and Scholes (1973) valuation formula no longer holds in financial markets have appeared with increasing frequency recently. When the Black/Scholes formula is used to imply volatilities from observed prices of options, the volatility estimates vary systematically across exercise prices and times to expiration. Derman and

Kani (1994), Dupire (1994), and Rubinstein (1994) argue this systematic behavior is driven by the fact that the volatility rate of asset return varies with the level of asset price and time. They go on to propose that volatility is a deterministic function of asset price and volatility and develop appropriate binomial or trinomial option valuation procedures.

In this paper, we apply the deterministic volatility option valuation approach to S&P 500 index option prices during the period June 1988 through December 1993 and find a number of interesting results. First, because of the limitless flexibility of the volatility function's specification, the DV model always does better in-sample than does the constant volatility model. Second, when the fitted volatility function is used to value options one week later, the model's predictions deteriorates with the complexity of the assumed volatility specification. Third, hedge ratios determined by the Black/Scholes model appear more reliable than those obtained from the DV option valuation model.

References

Black, F., 1976, Studies of stock price volatility changes, *Proceedings of the 1976 Meetings of the American Statistical Association,* Business and Economic Statistic Section, 177-181.

Black, F. and M.S. Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637-659.

Bossaerts, P. and P. Hillion, 1994.

Clément, C. Gouriéroux and A. Montfort, 1993.

Derman, E. and I. Kani, 1994, The implied volatility smile and its implied tree, Goldman Sachs Quantitative Strategies Research Notes, January 1994.

Derman, E. and I. Kani, 1994, Riding on the smile, *Risk* 7, 32-39.

Dupire, B., 1993, Pricing and hedging with smiles, Working paper, Paribus Capital Markets, London.

Dupire, B., 1994, Pricing with a smile, *Risk* 7, 18-20.

Duque, J. and D. Paxson, 1993, Implied volatility and dynamic hedging, *Review of Futures Markets* 13, 381-421.

Fleming, J., 1994, The quality of market volatility forecasts implied by S&P 100 index option prices, working paper, Fuqua School of Business, Duke University.

Fleming, J., B. Ostdiek and R. E. Whaley, 1993, Trading costs and the relative rates of price discovery in the stock, futures, and option markets, Working paper, Fuqua School of Business, Duke University.

Fleming, J., and R.E. Whaley, 1993, The value of wildcard options," *Journal of Finance* 49 (March 1994), 215-236.

Harvey, C.R., and R.E. Whaley, 1992a, Market volatility prediction and the efficiency of the S&P 100 index option market, *Journal of Financial Economics* 30, 43-73.

Harvey, C.R., and R.E. Whaley, 1992b, Dividends and S&P 100 index option valuation, *Journal of Futures Markets* 12, 123-137.

Heynen, R., 1993, An empirical investigation of observed smile patterns, *Review of Futures Markets* 13, 317-353.

Jacquier, E. and R. Jarrow, 1995,

Lo, A., 1986, Statistical tests of contingent-claims asset-pricing models: A new methodology, *Journal of Financial Economics* 17.

MacBeth, J. D. and L. J. Merville, 1980, Tests of the Black/Scholes and Cox call option valuation models, *Journal of Finance* 35, 285-301.

Nelder, J.A. and R. Mead, 1965, *Computer Journal* 7, 308-313.

Press, W.H., S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, 1992, *Numerical Recipes in FORTRAN: The Art of Scientific Computing*, Second edition, Cambridge University Press.

Rubinstein, M., 1994, Implied binomial trees, *Journal of Finance* 49, 771-818.

Shimko, D., 1993, Bounds of probability, *Risk* 6, 33-37.

Taylor, S.J. and X. Xu, 1993, The magnitude of implied volatility smiles: Theory and empirical evidence for exchange rates, *Review of Futures Markets* 13, 355-380.

Whaley, R.E., 1982, Valuation of American call options on dividend-paying stocks: Empirical tests, *Journal of Financial Economics* 10, 29-58.

Whaley, R.E., 1993, Derivatives on market volatility: Hedging tools long overdue, *Journal of Derivatives* 1, 71-84.

Table 1: Daily Average Root Mean Squared Valuation Errors for Estimation Models. RMSVE is the root mean squared valuation error, computed each day, and then averaged across all days in the sample period. AVERR is the average absolute valuation error outside the observed bid/ask quotes. FreqSW is the frequency of days, expressed as a ratio of the total number of days, on which a particular model has a lower daily RMSVE than the Switching Model.

Overall Sample:

	RMSVE	AVERR	FreqSW
Model 0	.6497	.3484	.000
Model 1	.3010	.0946	.010
Model 2	.2300	.0517	.147
Model 3	.2264	.0496	.373
Switching	.2268	.0499	—

Subcategories:

Moneyness (%) Lower	Moneyness (%) Upper		Days to Expiration: Less than 40 RMSVE	Less than 40 FreqSW	40 or more but less than 70 RMSVE	40 or more but less than 70 FreqSW	70 or more but less than 100 RMSVE	70 or more but less than 100 FreqSW
-10	-5	Model 0	.4333	.190	.5407	.239	.7916	.087
		Model 1	.2412	.550	.2777	.396	.3565	.115
		Model 2	.2403	.238	.2770	.226	.3075	.337
		Model 3	.2394	.247	.2771	.239	.3033	.058
		Switching	.2397	—	.2776	—	.3036	—
-5	0	Model 0	.5350	.021	.6550	.004	.6873	.029
		Model 1	.3646	.080	.2453	.262	.3050	.184
		Model 2	.1998	.226	.1978	.249	.2316	.404
		Model 3	.1930	.302	.1920	.253	.2292	.088
		Switching	.1933	—	.1925	—	.2292	—
0	5	Model 0	.4469	.031	.5711	.026	.7680	.015
		Model 1	.3395	.139	.2263	.389	.2465	.415
		Model 2	.1950	.236	.1861	.231	.2072	.370
		Model 3	.1888	.260	.1819	.231	.2045	.074
		Switching	.1893	—	.1819	—	.2045	—
5	10	Model 0	.4425	.188	.8514	.005	1.2568	.008
		Model 1	.2266	.672	.2462	.470	.3320	.205
		Model 2	.2541	.251	.2246	.219	.2212	.311
		Model 3	.2560	.237	.2207	.219	.2065	.068
		Switching	.2561	—	.2229	—	.2065	—

Table 2: Daily Average Root Mean Squared Valuation Errors for Prediction Models. RMSVE is the root mean squared valuation error, computed each day, and then averaged across all days in the sample period. AVERR is the average absolute valuation error outside the observed bid/ask quotes. FreqSW (FreqBS) is the frequency of days, expressed as a ratio of the total number of days, on which a particular model has a lower daily RMSVE than the Switching Model (Black/Scholes Model).

Overall Sample:

	RMSVE	AVERR	FreqSW	FreqBS
Model 0	.7840	.4493	.172	.089
Model 1	.5572	.2854	.460	.326
Model 2	.5665	.2997	.265	.405
Model 3	.5626	.2975	.302	.395
Switching	.5621	.2969	—	.395
Black/Scholes	.4829	.2266	.605	—

Subcategories:

Moneyness (%) Lower	Moneyness (%) Upper		Days to Expiration: Less than 40 RMSVE	FreqSW	FreqBS	40 or more but less than 70 RMSVE	FreqSW	FreqBS	70 or more but less than 100 RMSVE	FreqSW	FreqBS
-10	-5	Model 0	.4612	.222	.300	.6398	.277	.296	.7645	.240	.279
		Model 1	.2722	.517	.574	.3818	.535	.478	.5087	.365	.471
		Model 2	.2698	.209	.543	.3959	.239	.478	.5220	.231	.490
		Model 3	.2688	.296	.539	.3924	.270	.465	.5116	.221	.500
		Switching	.2690	—	.548	.3929	—	.465	.5114	—	.490
		Black/Scholes	.3111	.452	—	.4067	.535	—	.4976	.501	—
-5	0	Model 0	.6422	.282	.199	.8515	.294	.189	.8342	.360	.221
		Model 1	.4892	.380	.373	.5842	.482	.360	.7128	.500	.397
		Model 2	.4517	.233	.429	.6033	.254	.377	.8104	.316	.346
		Model 3	.4473	.314	.446	.6015	.298	.368	.7799	.191	.375
		Switching	.4483	—	.443	.6039	—	.368	.7847	—	.368
		Black/Scholes	.4093	.557	—	.4673	.632	—	.5591	.632	—
0	5	Model 0	.5676	.293	.265	.7372	.294	.215	.9341	.274	.156
		Model 1	.4694	.411	.429	.5515	.461	.382	.6793	.556	.370
		Model 2	.4464	.310	.505	.5707	.320	.390	.7743	.356	.363
		Model 3	.4477	.293	.509	.5741	.276	.382	.7657	.185	.378
		Switching	.4462	—	.509	.5740	—	.382	.7663	—	.378
		Black/Scholes	.4164	.491	—	.4528	.618	—	.5123	.622	—
5	10	Model 0	.4600	.266	.259	.8558	.156	.133	1.3147	.091	.038
		Model 1	.3047	.608	.580	.4560	.408	.440	.6188	.402	.394
		Model 2	.3139	.283	.542	.4577	.312	.436	.6184	.402	.439
		Model 3	.3180	.252	.531	.4568	.252	.436	.6152	.189	.432
		Switching	.3151	—	.545	.4564	—	.445	.6145	—	.432
		Black/Scholes	.3325	.455	—	.4172	.555	—	.5070	.568	—

Table 3: Daily Average Root Mean Squared Hedging Errors for Prediction Models. RMSHE is the root mean hedging error, computed each day, and then averaged across all days in the sample period. FreqSW (FreqBS) is the frequency of days, expressed as a ratio of the total number of days, on which a particular model has a lower daily RMSHE than the Switching Model (Black/Scholes Model).

Overall Sample:

	RMSHE	FreqSW	FreqBS
Model 0	.4547	.577	.467
Model 1	.4892	.557	.416
Model 2	.5078	.265	.395
Model 3	.5084	.275	.405
Switching	.5075	—	.402
Black/Scholes	.4406	.598	—

Subcategories:

Moneyness (%) Lower	Moneyness (%) Upper		Days to Expiration: Less than 40 RMSHE	FreqSW	FreqBS	40 or more but less than 70 RMSHE	FreqSW	FreqBS	70 or more but less than 100 RMSHE	FreqSW	FreqBS
-10	-5	Model 0	.3399	.415	.470	.4148	.510	.476	.4412	.456	.389
		Model 1	.2636	.585	.620	.3796	.594	.490	.4432	.567	.467
		Model 2	.2800	.250	.540	.3990	.231	.406	.4672	.344	.456
		Model 3	.2776	.300	.565	.3946	.336	.441	.4526	.111	.478
		Switching	.2790	—	.575	.3966	—	.427	.4532	—	.478
		Black/Scholes	.3252	.425	—	.4064	.573	—	.3974	.522	—
-5	0	Model 0	.4081	.458	.427	.4329	.618	.536	.4733	.536	.440
		Model 1	.3762	.477	.427	.4796	.600	.423	.5296	.632	.464
		Model 2	.3694	.238	.465	.5056	.227	.400	.5634	.368	.408
		Model 3	.3675	.250	.485	.5043	.264	.395	.5576	.088	.408
		Switching	.3674	—	.488	.5046	—	.400	.5584	—	.408
		Black/Scholes	.3551	.512	—	.4199	.600	—	.4524	.592	—
0	5	Model 0	.3728	.561	.496	.4142	.692	.522	.4472	.625	.523
		Model 1	.4300	.424	.355	.5272	.536	.326	.5645	.625	.398
		Model 2	.4119	.267	.443	.5494	.259	.362	.6099	.477	.383
		Model 3	.4121	.248	.443	.5531	.250	.348	.6145	.070	.367
		Switching	.4109	—	.439	.5520	—	.348	.6140	—	.367
		Black/Scholes	.3670	.561	—	.4120	.652	—	.4584	.633	—
5	10	Model 0	.3466	.426	.418	.3478	.577	.498	.4413	.570	.405
		Model 1	.3140	.496	.500	.4005	.547	.398	.4565	.653	.413
		Model 2	.3148	.287	.512	.4250	.274	.388	.4969	.504	.364
		Model 3	.3163	.270	.504	.4277	.244	.383	.5069	.058	.347
		Switching	.3143	—	.512	.4253	—	.383	.5060	—	.347
		Black/Scholes	.3191	.488	—	.3420	.617	—	.4079	.653	—

Figure 1: Black-Scholes implied voaltility smile patterns on April 1, 1992. Implied volatilities corresponding to call options of different times to expiration and different exercise prices are obtained from quoted option prices by inverting the Black-Scholes formula. The two curves correspond to bid and ask quoted option prices. Moneyness is defined as $100[(S-PVD)/Xe^{-rT}-1]$, and volatility is expressed as an annual percentage.

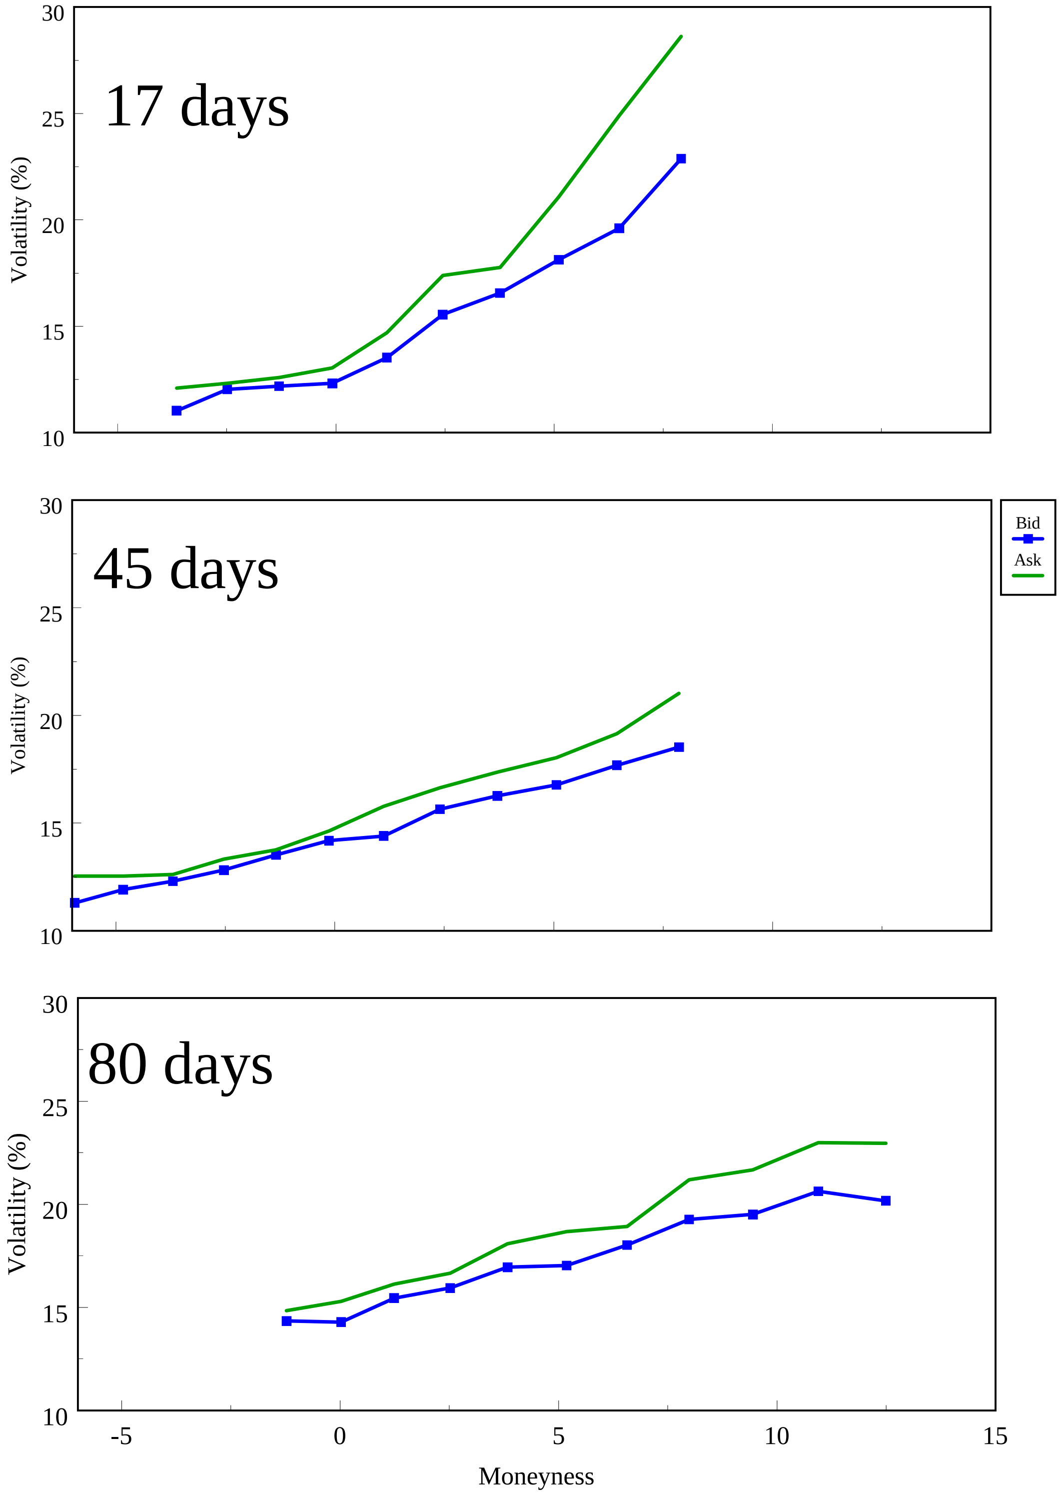

Figure 2: Cross-section of S&P 500 call and put option prices on April 1, 1992. The dashed line corresponds to the actual bid/ask midpoints of the S&P 500 options, while the solid line corresponds to their fitted values using the deterministic volatility option valuation approach. The notation C1, C2, and C3 (P1, P2, and P3) indicates call (put) options with the shortest, second shortest, and third shortest times to expiration, respectively. For each time to expiration category, calls are arranged from deep in-the-money to deep out-of-the-money, and puts are arranged in the reverse.

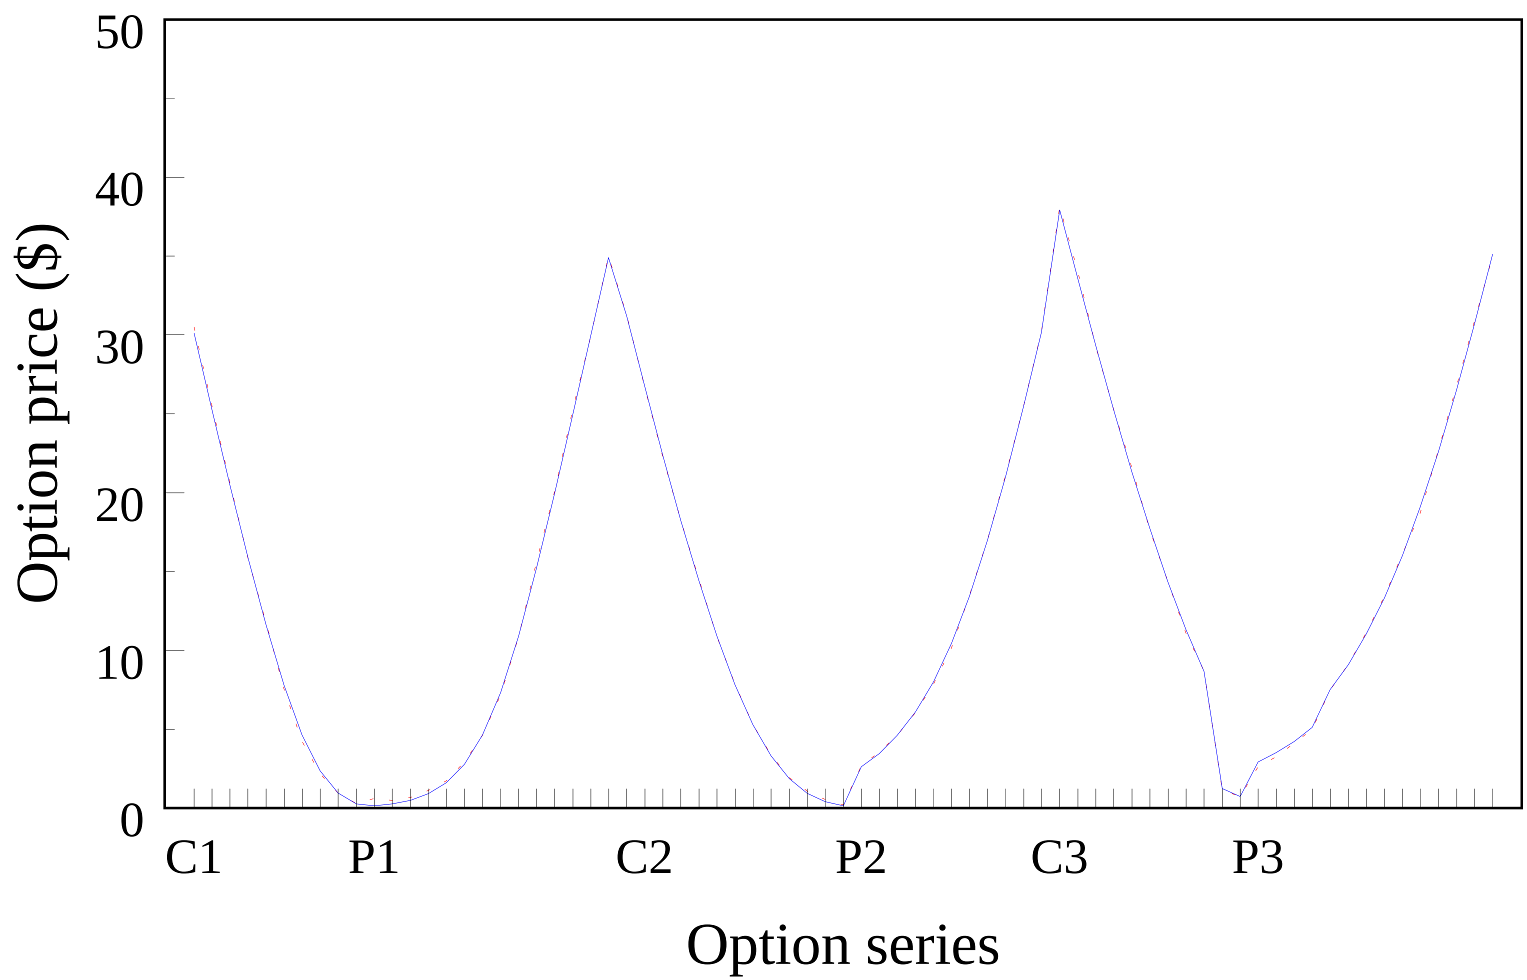

Figure 3: Cross-section of valuation errors for S&P 500 call and put option prices on April 1, 1992. The line corresponds to the difference between actual bid/ask midpoints of the S&P 500 options and their fitted values using the deterministic volatility option valuation approach. The notation C1, C2, and C3 (P1, P2, and P3) indicates call (put) options with the shortest, second shortest, and third shortest times to expiration, respectively. For each time to expiration category, calls are arranged from deep in-the-money to deep out-of-the-money, and puts are arranged in the reverse.

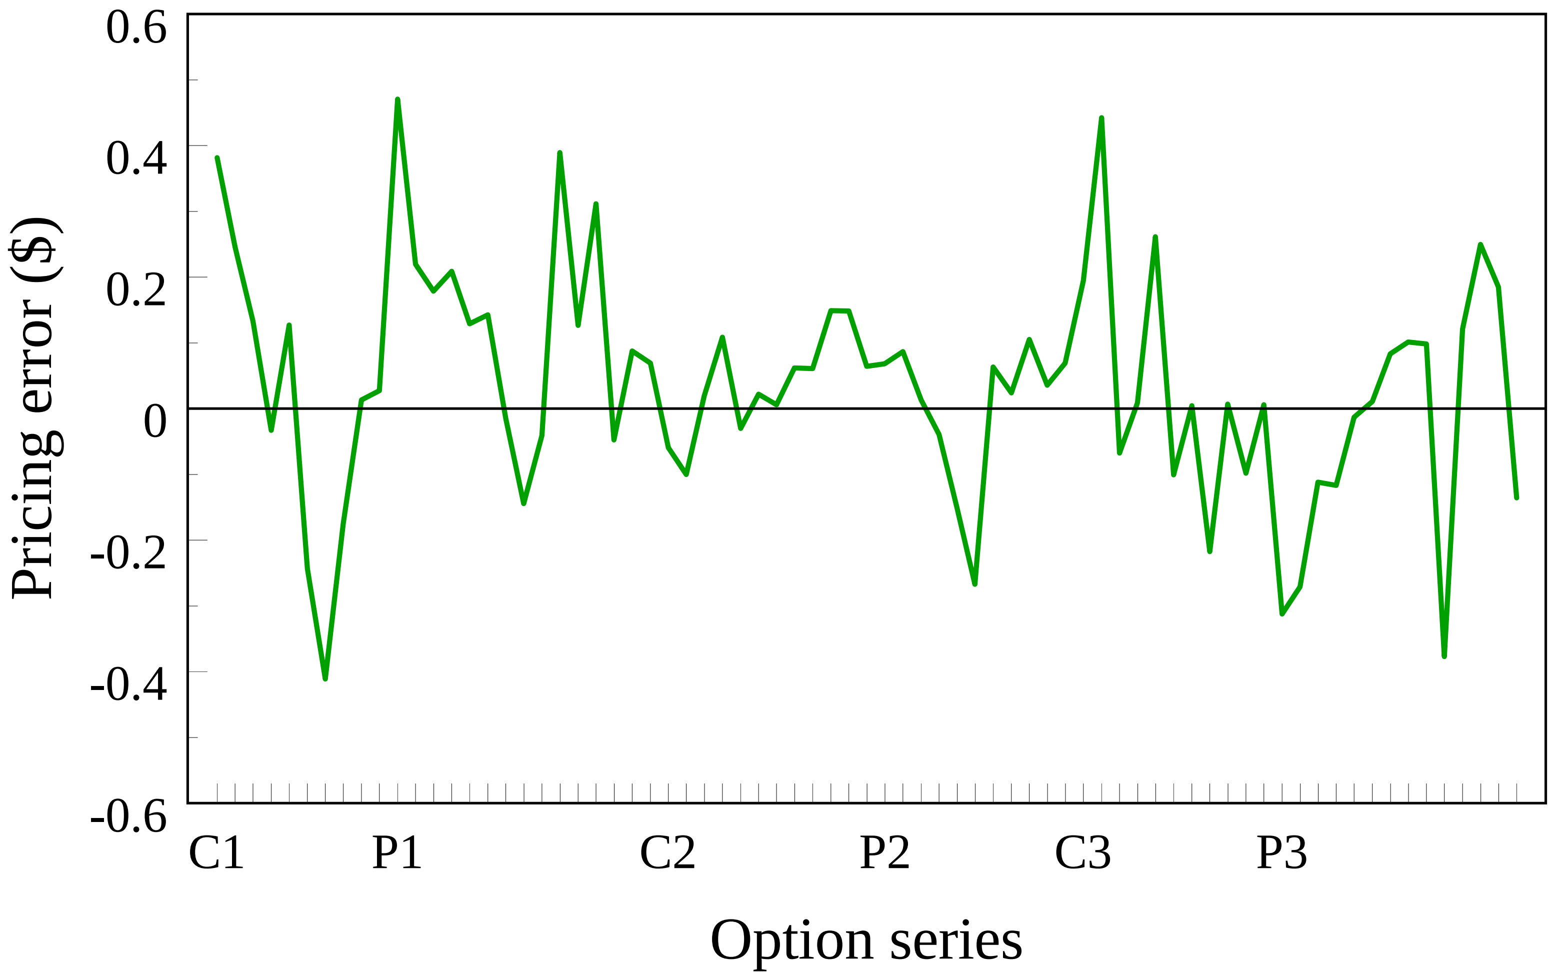

Figure 4: Estimated volatility function on April 1, 1992. The surface displays the annualized percent volatility rate $\sigma(S,t)$ for different index levels and different times to expiration, as implied by the deterministic volatility option valuation approach and S&P 500 option prices.

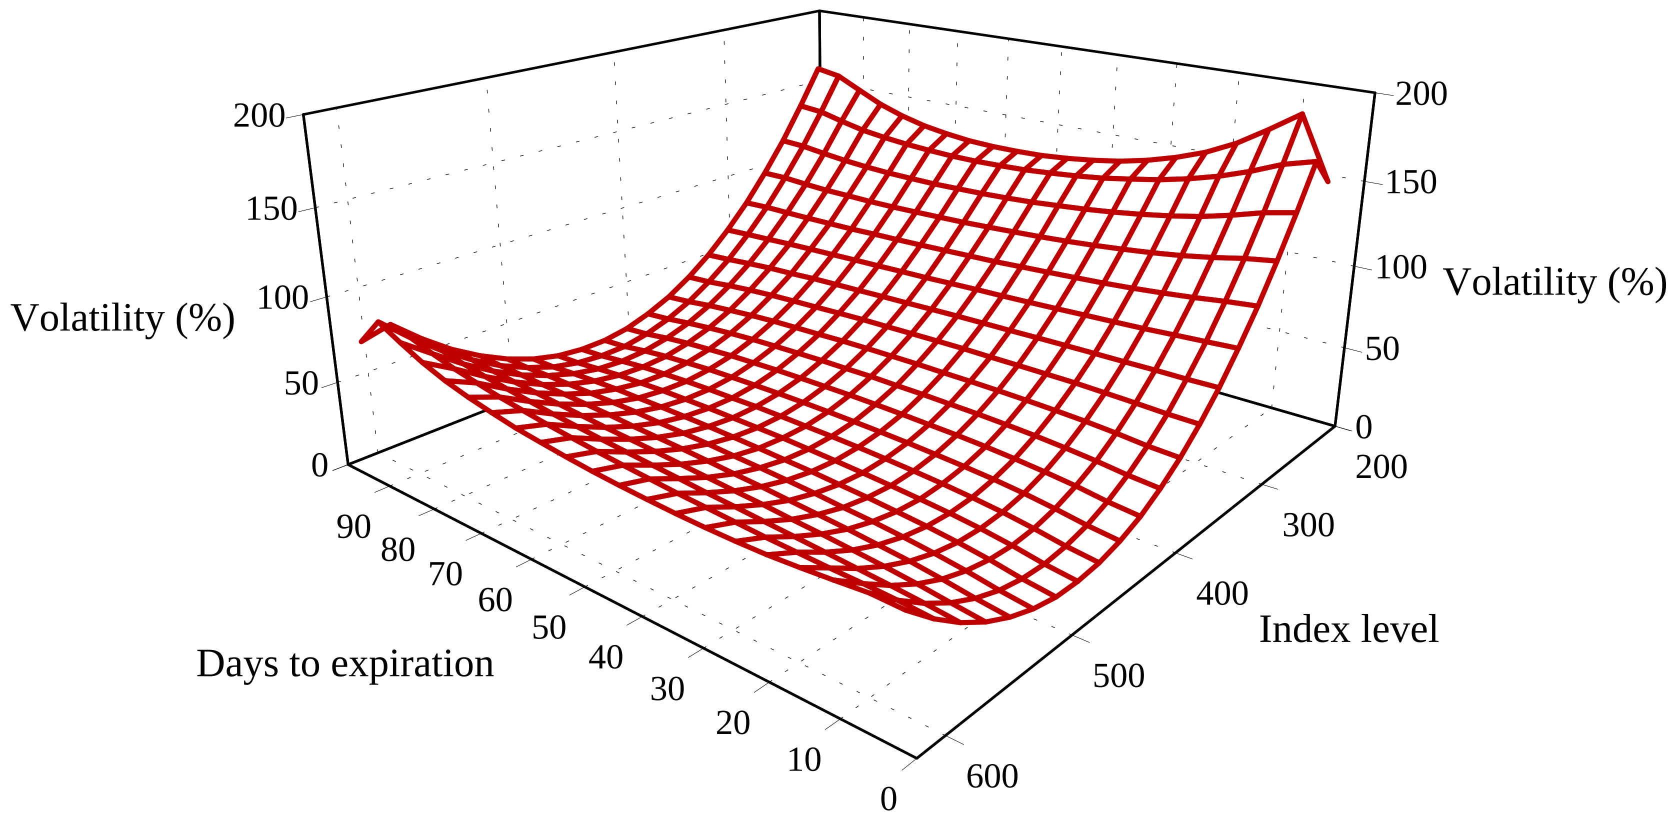

Figure 5: Risk-neutral probability density functions for April, May, and June S&P 500 option expirations. The distributions are based on a deterministic volatility function estimated on April 1, 1992.

Figure 6: Cross-section of predicted S&P 500 call and put option prices on April 1, 1992. The dashed lines correspond to the actual bid and ask prices of the S&P 500 options, while the solid line corresponds to their predicted values using the deterministic volatility option valuation approach. The notation C1, C2, and C3 (P1, P2, and P3) indicates call (put) options with the shortest, second shortest, and third shortest times to expiration, respectively. For each time to expiration category, calls are arranged from deep in-the-money to deep out-of-the-money, and puts are arranged in the reverse.

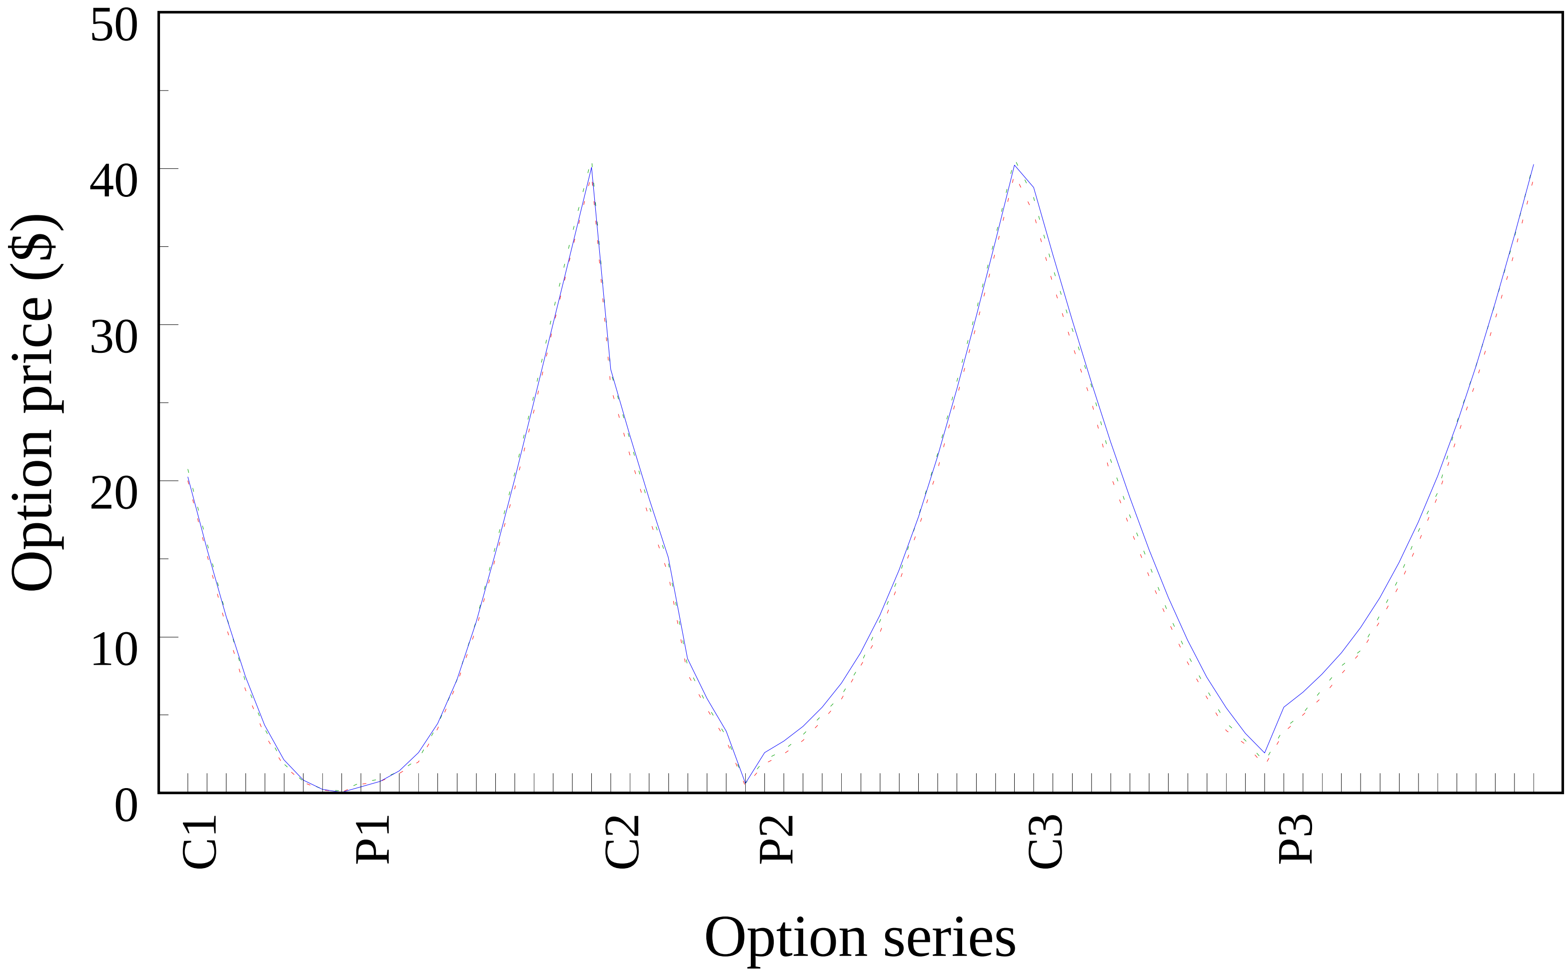

Figure 7: Cross-section of errors for predicted S&P 500 call and put option prices on April 1, 1992. The dashed lines correspond to the actual bid and ask prices of the S&P 500 options normalized by the bid/ask price midpoint. The solid line corresponds to the difference between the fitted option value using the deterministic volatility option valuation approach and based on the volatility function estimated the previous week and the bid/ask price. The notation C1, C2, and C3 (P1, P2, and P3) indicates call (put) options with the shortest, second shortest, and third shortest times to expiration, respectively. For each time to expiration category, calls are arranged from deep in-the-money to deep out-of-the-money, and puts are arranged in the reverse.