

How Useful are Implied Distributions?

Evidence from Stock-Index Options

by

Gordon Gemmill and Apostolos Saflekos*

Abstract

Option prices can be used to construct implied (risk-neutral) distributions, but it remains to be proven whether these are useful either in relation to forecasting subsequent market movements or in revealing investor sentiment. We estimate the implied distribution as a mixture of two lognormals and then test its one-day-ahead forecasting performance, using 1987-97 data on LIFFE's FTSE-100 index options. We find that the two-lognormal method is much better than the one-lognormal (Black/Scholes) approach at fitting observed option prices, but it is only marginally better at predicting out-of-sample prices. A closer analysis of four "crash" periods confirms that the shape of the implied distribution does not anticipate such events but merely reflects their passing. Similarly, during three British elections the implied distributions take on interesting shapes but these are not closely related to prior information about the likely outcomes. In short, while we cannot reject the hypothesis that implied distributions reflect market sentiment, we find that sentiment (thus measured) has little forecasting ability.

Keywords: option pricing, implied distribution, volatility smile, market sentiment, crashes, elections.

*The authors are grateful for comments from Robert Bliss and Paul Dawson.

Introduction

Investors, risk-managers and policy-makers all need to forecast the probability distribution of prices if they are to take rational decisions. Conventionally, an estimate of the variance is obtained from recent data on returns. A month of data may give a reasonable estimate of the variance, but observations over several months are required if the skewness and kurtosis are to be measured accurately. Another approach is to use options data to construct implied distributions. These are the so-called risk-neutral distributions (RNDs) which traders are using when they set the prices of the options and which relate to the period until an option expires. One day's options can reveal not only the forecast variance, but also the whole shape of the risk-neutral distribution (including skewness and kurtosis). Using options is therefore an efficient way in which to forecast the whole distribution.¹ Several alternative methods have been suggested for extracting the risk-neutral distribution from option prices, the main difference between them being the extent to which they constrain its shape. At one extreme, Longstaff (1995) imposes no constraints but the result can be a rather "badly-behaved" or spiky distribution. At the other extreme, Rubinstein (1994) and Jackwerth and Rubinstein (1996) constrain their distributions to be those with the smallest possible deviations from the lognormal. Somewhere in between these two extremes is the assumption of this paper, which is based on the work of Ritchey (1990), Melick and Thomas (1997) and Bahra (1997). We assume that the distribution can take any shape which may be approximated by a mixture of two lognormals.

The first objective of this study is to examine whether an option pricing model, based upon two lognormal distributions, performs well for equity-index options (having previously been applied only to oil futures and interest-rates). The performance of the method is determined not only by measuring the (ex-post) fit of the implied distribution, but also ex-ante by testing how well it forecasts option prices out-of-sample. We find for LIFFE's FTSE-100 index options over the 1987-97 period that although the model fits the data significantly better than the Black/Scholes model, the out-of-sample performance is only marginally better. This is consistent with work on US index options by Dumas, Fleming and Whaley (1998), who found that taking account of volatility smiles did not help in forecasting one-day-ahead option prices.²

If implied distributions are of rather limited use in normal periods, it might still be possible that they help to forecast market movements during exceptionally turbulent periods. Our second objective is therefore to examine the performance of the method around crashes (of October 1987, October 1989 and October 1997), British general elections (of May 1987, April 1992 and May 1997), and

¹ Of course, it is important to remember that these implied distributions are different from directly sampled distributions, as they reflect risk-neutral processes (see Harrison and Kreps, 1979). In the absence of transactions costs the shapes should be the same, but the location of the implied distribution reflects only a risk-free rate of return.

² There is a one-to-one relationship between the volatility smile and the implied distribution, as demonstrated explicitly by Shimko (1994), so forecasting with the volatility smile is equivalent to forecasting with the implied distribution.

extraordinary events (the sterling currency crisis of September 1992). Our chosen periods are more general than those examined (in a related way) by other researchers. For example, Bates (1991) and Gemmill (1996) have examined volatility smiles for the US and UK around the 1987 crash and Malz (1997), Campa and Chang (1995) have examined options on sterling in the ERM crisis of 1992. Coutant, Jondeau and Rockinger (1998) have examined implied distributions for interest rates at the time of the snap general election in France in 1997. We find that the method does not help to reveal the probability of crashes, because increased left-skewness **follows** rather than precedes these events. Nevertheless, the method can help to reveal the divergent expectations which arise immediately after crashes and during election campaigns. In other words, the method helps to reveal “market sentiment”, which could be useful for the policy-stance of a central bank (e.g. Federal Reserve, Bank of England) and for investors who may wish to take positions based upon the difference between their forecast of the distribution and the consensus of the market.

The paper is written as follows. The next section provides a theoretical presentation of the main techniques used to derive the implied distribution of future asset prices from option prices. The two-lognormal mixture distribution method is described, as well as the practical issues of the implementation (namely the data used and the selection of the studied periods). Section 3 presents the empirical results from applying the method to FTSE-100 index options, at the same time assessing its usefulness. Conclusions and suggestions for further research are given in the fourth and final section of the paper.

2. Theoretical framework

2.1 Review of literature and development of the model

Option prices reflect forward-looking distributions of asset prices. In the absence of market frictions, it is possible to take a set of option prices, for a single maturity and at various exercise prices, and imply the underlying risk-neutral distribution (RND). Breeden and Litzenberger (1978) first showed how the second partial derivative of the call-pricing function with respect to the exercise price is directly proportional to the RND function. However, since observed option prices are only available at discretely spaced intervals rather than being continuous, some approximation for the second derivative is necessary and more than one implied distribution can be implied depending on the approximation chosen. As Jackwerth and Rubinstein (1996) observe, selecting among the competing distributions then amounts to a choice of how to interpolate or extrapolate option prices across exercise prices.

The most direct way of estimating the implied distribution is by simple application of the Breeden and Litzenberger result to a function relating the call price to exercise prices. This has been done by Longstaff (1995) and Ait-Sahalia and Lo (1995). The former implements a procedure that attributes a

probability to an option mid-way between two adjacent exercise prices, then uses this to solve for the next probability, and so on. Ait-Sahalia and Lo first smooth the pricing function with a set of polynomials and then proceed in a similar way.

Shimko (1993, 1994) proposes an alternative approach by interpolating in the implied-volatility domain instead of the call-price domain. He begins by fitting a quadratic relationship between implied volatility and exercise price. The Black/Scholes formula is then used to invert the smoothed volatilities into option prices. At this point he has a continuous spectrum of call prices as a function of the exercise prices and the application of the Breeden and Litzenberger result is straightforward, generating the implied probability distribution.

The main limitation of the above techniques is the need for a relatively wide range of exercise prices. This can be overcome by imposing some form of prior structure on the problem. One such prior (used by Bates (1991, 1996) and Malz (1997)) is to assume a particular stochastic process for the price dynamics of the underlying asset. In their papers the asset price is assumed to follow a jump-diffusion process. In other words, the basic probability distribution is lognormal, but it can jump up or down³.

Alternatively, the imposed structure may apply to the distribution of the future asset price itself, instead of the asset-price dynamics. This approach proves to be more general than making assumptions about the stochastic process of the underlying asset price, because any given RND function is consistent with many different stochastic processes, whereas a given stochastic price process implies a unique RND function (Melick and Thomas, 1997). The approach of Rubinstein (1994) and Jackwerth and Rubinstein (1996) falls into this category. They employ an optimisation algorithm to find that RND function which is closest to lognormal, taking account of bid/ask bounds on the observed option prices.⁴

The framework used in the current paper follows Ritchey (1990), who notes that a wide variety of shapes may be approximated with a mixture of lognormal distributions. He assumes that the implied density function, $f(S_T)$, of the underlying asset terminal price, S_T , comprises a weighted sum of k individual lognormal density functions:

$$f(S_T) = \sum_{i=1}^k [\theta_i L(a_i, b_i, S_T)] \quad \text{Equation 1}$$

where $L(\alpha_i, b_i, S_T)$ is the i^{th} lognormal density function with parameters α_i, b_i :

$$L(a_i, b_i, S_T) = \frac{1}{S_T b_i \sqrt{2\pi}} e^{\{- (\ln S_T - a_i)^2\} / 2 b_i^2} \quad \text{Equation 2}$$

³ Malz assumes that there is either no jump or just one jump over the life of the option, in which case the terminal RND function is a mixture of two lognormal distributions (Bahra, 1997)

⁴ The distance criteria used in the two papers are, respectively, a quadratic difference and a smoothness function.

$$a_i = \ln S + (\mu_i - \sigma_i^2 / 2) \tau \text{ and } b_i = \sigma_i \sqrt{\tau} \quad \text{Equation 3}$$

In the above equations, S is the spot price of the underlying asset, $\tau (=T-t)$ is the time remaining to maturity and μ and σ are the parameters of the normal RND function of the underlying returns. The weights θ_i are positive and sum to unity.

Melick and Thomas (1997) apply this framework to options on crude oil futures, using a mixture of three lognormal distributions. Bahra (1997), Butler and Davies (1998) and Soderlund and Svensson (1997) use a mixture of two lognormals on interest-rate futures. Since our data on FTSE-100 options cover a limited range of exercise prices for each maturity, it seems more appropriate to use two lognormals, which require only five parameters: the mean of each lognormal, α_1, α_2 , the standard deviation of each lognormal, b_1, b_2 and the weighting coefficient, θ .

2.2 The two-lognormal mixture distribution method for equity index options

Let the terminal pay-off on a European call maturing at time T be $\max(S_T - X, 0)$, given terminal asset price S_T and exercise price X . Assuming that the risk-free interest rate r is constant, the life of the option is τ and the asset price is S , then the price of the call is the discounted expected payoff (conditional upon finishing in the money) times the probability of finishing in the money:

$$c(X, \tau) = e^{-r\tau} \int_X^{\infty} f(S_T)(S_T - X) dS_T \quad \text{Equation 4}$$

where $f(S_T)$ is the risk-neutral probability density function of the terminal asset price at time T . Similarly, the terminal payoff on a European put is $\max(X - S_T, 0)$ and its current price is:

$$p(X, \tau) = e^{-r\tau} \int_0^X f(S_T)(X - S_T) dS_T \quad \text{Equation 5}$$

Under the assumption that the probability density function is a mixture of two lognormals (with weights θ and $(1-\theta)$), the above equations for call and put prices can be rewritten as:

$$c(X, \tau) = e^{-r\tau} \int_X^{\infty} [\theta L(a_1, b_1, S_T) + (1-\theta)L(a_2, b_2, S_T)](S_T - X) dS_T \quad \text{Equation 6}$$

$$p(X, \tau) = e^{-r\tau} \int_0^X [\theta L(a_1, b_1, S_T) + (1-\theta)L(a_2, b_2, S_T)](X - S_T) dS_T \quad \text{Equation 7}$$

These equations can be used iteratively to minimise the deviation of estimated prices from observed prices, a search being made over the five parameters. We use both puts and calls across five exercise prices and minimise the total sum of squared errors for the ten options:

$$\sum_{i=1}^n [c(X_i, \tau) - \hat{c}_i]^2 + \sum_{i=1}^n [p(X_i, \tau) - \hat{p}_i]^2 \quad \text{Equation 8}$$

where $b_1, b_2 > 0, 0 \leq \theta \leq 1$, subscript i denotes an observation and $\hat{}$ denotes an estimate.

Bahra (1997) shows that equations 6 and 7 have the following closed-form solutions:

$$c(X, \tau) = e^{-r\tau} \left\{ \theta \left[e^{a_1 + b_1^2/2} N(d_1) - X N(d_2) \right] + \right. \\ \left. (1 - \theta) \left[e^{a_2 + b_2^2/2} N(d_3) - X N(d_4) \right] \right\} \quad \text{Equation 9}$$

and

$$p(X, \tau) = e^{-r\tau} \left\{ \theta \left[-e^{a_1 + b_1^2/2} N(-d_1) - X N(-d_2) \right] + \right. \\ \left. (1 - \theta) \left[-e^{a_2 + b_2^2/2} N(-d_3) - X N(-d_4) \right] \right\} \quad \text{Equation 10}$$

where

$$d_1 = \frac{-\ln X + a_1 + b_1^2}{b_1}, \quad d_2 = d_1 - b_1 \\ d_3 = \frac{-\ln X + a_2 + b_2^2}{b_2}, \quad d_4 = d_3 - b_2$$

Equations 9 and 10 have a very simple interpretation: the model prices are just weighted sums of two Black/Scholes solutions, each having its own mean and variance.

2.3 Data sources

The empirical research in this study is based on the daily settlement prices of FTSE-100 calls and puts covering up to five exercise prices and four maturities for each day from January 1st 1987 to December 31st 1997⁵. Our sample contains American-style options for the period to March 1992 and European-style options thereafter, the switch being made because the latter were only thinly traded

⁵ We gratefully acknowledge LIFFE for financial assistance in collecting some of the data and for providing the other data.

before then.⁶ Exercise prices have been chosen such that one is at-the-money, two are in-the-money and two are out-of-the-money. The interest rates used are UK Sterling 3-month interbank deposit rates, retrieved from Datastream.

2.4 Hypotheses to be tested

Hypothesis 1: *The two-lognormal model performs better than a one-lognormal (Black/Scholes) model*

This paper is primarily a critical examination of the two-lognormal model. On one day per month for the period 1/87 to 11/97 the distribution is implied and then used to price options on the next day. This allows us to test the method's forecasting performance relative to the Black/Scholes model over quite a long period.⁷

Hypothesis 2: *The option market anticipates crashes*

One particularly interesting question is whether option markets have any usefulness in predicting extreme events, such as stock market crashes. The crash of October 1987, the mini crash of October 1989 and the market turmoil of October 1997 were chosen as examples of such occasional events. Although other studies have looked at some of these periods (e.g. Bates, 1991, for the US and Gemmill, 1996, for the UK), the use of the two-lognormal mixture is original. We also examine the European monetary crisis of September 1992 (when sterling left the Exchange Rate Mechanism (ERM) and was devalued by more than 10%) as another period of great uncertainty for the British stockmarket.

Hypothesis 3: *A bimodal distribution is appropriate during elections*

The two-lognormal mixture may prove particularly useful in periods when a market jump is expected but the direction of the jump is unknown. Such is the case during political elections. Assuming that there are only two possible outcomes (for example, Labour victory or Conservative victory) and that investors prefer one to the other, a stock-index option which matures after an election should reflect a bimodal underlying distribution. The mixture of two separate lognormal distributions should therefore fit the observed option prices particularly well at such times. We have included the British elections of 1987, 1992 and 1997 in our analysis in order to test this hypothesis.

⁶ Strictly speaking the method is only applicable to European options, because we are attributing all of an option's value to the terminal distribution and not to early exercise. However, the value of early exercise on these index options is likely to be small: see Dawson (1994) for an analysis on FTSE-100 options.

⁷ Dumas, Fleming and Whaley (1998) use a similar approach on S&P 500 options, but based upon the volatility smile rather than the implied distribution.

3. Results

3.1 Forecasting Performance over the Whole Period, January 1987-December 1997

The options used in this part of the analysis are chosen on one day per month (from the middle of the week) such that they have approximately 45 days to maturity. Table 1 gives a summary of the conventional dispersion and shape statistics of the implied distribution over the whole period. The results leave no doubt about two features. First, the implied distributions have fatter tails than those of lognormal distributions, with kurtosis positive in each subperiod and averaging 1.54. This result is expected, as it is the corollary of the well-known volatility smile which is found for many different options (e.g. Bahra, 1997, on interest rates, Melick and Thomas, 1997, on crude oil futures, Malz, 1996, on foreign exchange).

Second, and more importantly, the generated distributions exhibit consistently negative skewness (that is, they have a more pronounced tail to the left) averaging -0.26 over the whole period and becoming more pronounced over time. Figure 1 plots the monthly results, which indicate that after March 1991 there is no month in which skewness is positive, although variation is quite large. This result differentiates equity-index options from options on other underlying assets and is also well documented from volatility smiles (see e.g. Gemmill, 1996).

The performance of the two-lognormal model can be compared with Black/Scholes in two ways; first, in how well (ex-post) the two models fit observed option prices within sample, and second, in forecasting (ex-ante) the price of an option on day $t+1$. To do the latter we obtain the model parameters (a_1 , b_1 , a_2 , b_2 and θ) that best fit the option prices observed on day t . Then for forecasting we update the means of the distributions to take account of changes in the stock price.⁸

$$a_i^{t+1} = a_i^t + \ln \frac{S^{t+1}}{S^t} \quad \text{Equation 11}$$

Similarly, we update the variances to take account of one day's less time to maturity:

$$b_i^{t+1} = b_i^t \sqrt{\frac{\tau^{t+1}}{\tau^t}} \quad \text{Equation 12}$$

Table 2 shows the errors obtained by the two-lognormal and Black/Scholes methods, both within sample (ex-post) and out-of-sample (ex-ante). Results are given separately for 1/87 to 2/92, for which

⁸ Strictly speaking, the adjustment should take account of the change in the forward price, for which the change in spot price is a good approximation unless an ex-dividend date is straddled (which does not occur in our sample).

American options were used, and for 3/92 to 11/97, for which European options were used. The two-lognormal method has an in-sample performance which is considerably better than that of the Black/Scholes model, being 29% better in terms of root-mean-squared error for the American options in the earlier period and 89% better for the European options in the more recent period. This is to be expected since it uses five parameters ($a_1, b_1, a_2, b_2, \theta$) as compared with the two parameters (a, b) of Black/Scholes.⁹ The out-of-sample (forecasting) test shows a root-mean-squared-error improvement of only 11% for the American options (54 out of 61 observations show an improvement), but a more impressive gain of 43% for the European options (all 68 observations show an improvement). However, these *relatively* large RMSE improvements for European options translate into *absolute* gains of about 1-2 index points per option, which are small when compared with a bid/ask spread of at least 2 points. Hence the method gives a consistent but small improvement in forecasting performance on average across the eleven year period.

3.2 Market crashes

The crash of 1987

Even if the method gives only small benefits in most periods, it may fit the data and forecast better than Black/Scholes in periods when there are significant events. On Monday, October 19th 1987 the FTSE-100 fell by 10.9% to 2052. The following day, after the news from Wall Street convinced everybody that this was a global crash, there was a further drop of 12.2%. The decline continued and the index dropped to 1684 on October 26th, 1608 on November 4th and to the 1987 low of 1565 on November 9th. This represented a fall of 32% in three weeks. The London stock market did not recover these losses for more than 18 months.

We have divided our analysis of implied distributions around this time into two distinct periods: the first is the two trading weeks immediately before the crash and the second is the month immediately after the crash. Results on shape and goodness of fit are summarised in Table 3 (first segment) and representative implied distributions are plotted in Figure 2.¹⁰ It should be noted that in this and subsequent figures, the first distribution is plotted as estimated but the other distributions have been adjusted to give the same time to maturity as the first. Without such an adjustment there would be a narrowing of the distributions as maturity approached.¹¹ Prior to the crash (13th October in Figure 2)

⁹ In fact, Black/Scholes normally has only one unknown parameter, the volatility of the distribution. However, since we do not use the forward price of the underlying asset as a parameter in the minimisation procedure but imply it as the mean of the distribution, the number of B/S parameters is two and the number of additional parameters of the two-lognormal model is therefore three.

¹⁰ Because averaging across distributions tends to remove deviations from lognormality, we have plotted the distribution on that particular day which has skewness nearest to the period average.

¹¹ The adjustment to the mean takes account of the change in spot price and the size of the contango (forward price less spot price), treating each distribution separately. We have $\alpha_i^* = \ln S + (\alpha_i - \ln S) (t_1 / t_2)$, where α_i^* is the adjusted value for the i th mean, t_1 is the maturity observed and t_2 is the maturity required for purposes of comparison. Each variance is adjusted as in Equation 12.

the distribution is unimodal, with fat tails and slightly positive skewness. The means of the two component distributions are close together ($a_1 \approx a_2$), but there is a large difference in their standard deviations (b_1, b_2) which generates the fat tails. Option prices on the three trading days from October 19th to October 21st have been excluded, because both methods lead to huge errors when fitted.¹² The two lognormal distributions thereafter move apart and, on some days, give a bimodal composite, as illustrated for November 6th in the middle segment of Figure 2. The representative distribution for the whole month after the crash, as shown by November 26th in Figure 2, is not bimodal but does show quite widely separate means for the two component distributions ($a_1 \neq a_2$). In this period the mode of the composite function exhibits great instability, jumping regularly between 1300 and 1600 which reflects the difficulty which investors had in reaching a new consensus.

From before to after the crash, average volatility jumps from 20.8% to 50.7%¹³ and skewness falls from positive (0.36) to negative (-0.26). However, kurtosis falls compared to the pre-crash period (from 1.74 to 0.02). Finally, it is interesting to note that the mean of the implied distribution is below the spot price for much of this period, consistent with the observed backwardation in the futures market.

The mini crash of 1989

On October 16th 1989, two years after the 1987 crash, the FTSE-100 dropped by 74 points (3.3%) to 2163. This was the most dramatic plunge of the index for more than a year, but this time the market reacted in a more muted way than in 1987, avoiding panic and quickly recovering its prior level. Our analysis is divided into a before-crash period of September 11th to October 15th and an after-crash period of October 16th to November 10th.¹⁴ There are only slight changes in the implied distributions from the first to the second period (see Table 3, second segment, and Figure 3), which is in stark contrast to what happened in 1987. Volatility increases from 19.3% to 27.8%, but that is a small change relative to events in October 1987. Left-skewness increases from -0.26 to -0.35, but there is a reduction in kurtosis just as in 1987 (from 1.74 to 0.02). In both the pre-crash and post-crash periods the two component distributions move closely together, preserving the unimodal nature of the composite RND function.

The ERM crisis of 1992

The period of our study is mid-August to mid-October 1992. During the second and third quarter of 1992 a number of European currencies – including sterling – were subject to strong pressures which

An alternative to such an approach would have been to do the analysis with options of two maturities and then synthesise an implied distribution for a constant maturity. This is discussed by Butler and Davies (1998).

¹² Both methods give large fitting errors from October 22nd to October 28th (daily MSE in the range 3 to 8), but these results are included as the period is particularly important.

¹³ Volatility is measured empirically by integrating across the composite distribution.

¹⁴ November index options were used for the analysis.

eventually pushed them to the ERM floor. The British government tried to hold sterling's value against the D-Mark through a series of interest rates increases. As rates rose, so the stockmarket fell by about 15% from early May to late August. The last effort to defend sterling was an unprecedented 5 percentage point rise in interest rates, announced in two steps, on September 16th 1992. Before the end of that day sterling had left the ERM and the second of the two interest rate rises did not come into effect. The devaluation of sterling and lower expected interest rates then pushed the FTSE-100 up by 7.7% over two days.

Despite the rise in share prices, the shape of the implied distribution is only mildly changed from before to after the devaluation, becoming almost bimodal (see Table 3, third segment, and Figure 4). Average volatility and kurtosis both increase slightly (volatility from 22.9% to 23.5%, kurtosis from 1.61 to 2.50). Skewness, which is extremely negative in both cases, moves from -0.60 to -0.74. It seems that exit from the ERM, an *upward* "shock" for the stockmarket, had a very muted impact on the implied distribution, which contrasts with the impact of the *downward* shocks in 1987 and 1989 which increased left-skewness and volatility. This result is consistent with time-series models of volatility (such as EGARCH) which find significant asymmetry: volatility rises by much more when the market falls than it does when the market rises (see, for example, Crouhy and Rockinger, 1993).

The Asian crash of 1997

Our analysis covers 59 trading days, from 9/9/97 to 28/11/97, and is based on options maturing in December. In the five weeks before the 20th October volatility is relatively high (19.6% – see Table 3, fourth segment). The implied distribution is highly skewed to the left (-0.81) but its kurtosis is of lognormal size (0.03). After October 20th, volatility almost doubles (34.9%) and kurtosis increases significantly (1.05). The implied distribution is still skewed to the left, but on average less so than before (-0.56). The representative plots in Figure 5 show how the distribution changes from being almost bimodal on 16th September, to being stretched and clearly bimodal on 29th October. This suggests that investors hold widely difference views about the potential level of the index at this time. Our analysis of crashes can be summarised as follows. The largest effect is simply an asymmetric impact on volatility, which responds more to a market fall than to a market rise. Left-skewness rises hugely after the 1987 crash, but does not change much when smaller shocks occur. There is no consistent pattern in kurtosis during these events, but there is a tendency for a bimodal distribution to appear after the event. Finally, there is no pattern in the results to suggest that any of these four events was anticipated by participants in the options market.

3.3 British Elections

The 1987 election

The 1987 election was called on May 11th and held one month later on June 11th. During the whole campaign the Conservatives maintained a clear lead in the opinion polls over Labour and the stockmarket rose by 5.5%. Investors were said to be awaiting a “Japanese wall of money” which would arrive after a Conservative victory (Financial Times, May 30th 1987).

Our study covers 30 trading days from 12/5/87 to 23/6/87 and is based on the prices of the July (American-style) options. During the campaign the distribution is volatile, averaging 26.5% as given in Table 4, and quite left-skewed, averaging -0.62. The representative distribution plotted for 20th May in Figure 6 indicates a nearly bimodal distribution. Taking account of the information available at the time (see Gemmill, 1992) it is reasonable to assert that the mode at 2010 represents a Labour win while the mode at 2300 represents a Conservative win. Relative to the average futures price for this period, these represent an anticipated 3.4% rise if the Conservatives win and a 9.6% fall if Labour wins. After the election the distribution resumes its familiar shape, becoming unimodal and with a more modest dispersion (volatility=21.9%) and skewness (-0.09).

The two-lognormal model proves to be useful during the election campaign in revealing the market’s sentiment, confirming that two distinct outcomes are perceived to be possible. However, that perception is itself difficult to explain, given the extremely high probability of a Conservative win which the opinion polls forecast throughout this period (see Gemmill, 1992) and the absence of any upward movement in the market after the Conservative win. The method appears to have captured market sentiment in advance of a known event, but that sentiment also appears to have been a rather misleading forecast of the election outcome.

The 1992 election

The 1992 general election was called on March 11th and held four weeks later, on April 9th. Unlike the previous election, when the Conservatives had been the strong favourites, the outcome of this election was very uncertain. The opinion polls gave Labour a narrow lead over the Conservatives, which, if confirmed on election day, might have led to a hung parliament. Therefore, the possible scenarios were three: a Labour government, a Labour-Liberal Democrat coalition and a Conservative government.

Just as in the previous election, the market’s disposition in favour of Conservatives was often heard and the prospect of a change in government caused the stock market to slide. The index eventually fell to its 1992 low a few days before the election. On the day of the election, the index gained considerable ground, since a published opinion poll created “last-minute hopes” for a Conservative

victory. The following day the Conservatives returned to power once again and the stock market gained 136 points (5.4%) in one session.

Our analysis covers the period from March 12th to April 30th 1992 and is based on the May (American-style) options. During the campaign volatility is 23%, skewness -0.20 and kurtosis 1.96 (see Table 4). This time we do not observe a bimodal pattern (see Figure 7), even though the polls indicate that the final winner is less clear than in 1987. After the election volatility falls (to 17%), skewness becomes less negative (-0.12) and kurtosis rises to 5.53. In sum, the Conservative win brings forth a smile rather than a sneer, but the options do not forecast that the market will rise on a Conservative victory.

The 1997 election

The election was announced on March 17th and held six weeks later, on May 1st. During the unusually long campaign, the Labour party maintained an estimated lead over the Conservatives which ranged from 28% at the start to 5% a week before voting. Unlike 1992, this time the actual outcome confirmed the predictions of the opinion polls and Tony Blair became the first Labour prime minister in 18 years.

Our analysis covers 50 trading days, from 18/3/97 to 30/5/97, and is based on the June (European-style) options. During the campaign the implied distribution is skewed to the left (-0.53 , see Table 4) and kurtotic ($+0.53$). After the election, volatility decreases (from 15.4% to 13.0%) and both kurtosis and skewness decline (skewness= -0.45 , kurtosis= 0.29). The representative distributions for 24th April and 15th May in Figure 8 are very similar in shape. In sum, the 1997 election is almost a “non-event” for the index-options market.

What has been learnt from the study of election periods? In principle, they provide an ideal test of the informational content of implied distributions, since a known event is certain to occur on a specific date but its impact has to be forecast. The two-lognormal method gives much smaller root-mean-squared errors (in sample) than does the Black/Scholes model, particularly in 1997. It reveals a rather bimodal distribution in 1987, but not in 1992 or 1997 when such a distribution would have seemed more plausible given the more balanced contests. In sum, the analysis helps to “tell a story” about investors’ expectations, but it is not a story which is supported by subsequent outcomes: if investors’ expectations are revealed by the implied distributions during election campaigns then those expectations do not seem to have much forecasting power.

4. Conclusions

In this study we have examined whether implied distributions are informative with respect to subsequent stockmarket moves and to what extent they may be used to reveal investor sentiment. To do this we have applied the mixture-of-two-lognormals technique to London’s FTSE-100 index

options and critically examined the in-sample and out-of-sample performance of this model in a variety of periods.

The analysis was intended to test three general hypotheses: 1) the two-lognormal model performs better than Black/Scholes; 2) implied distributions indicate that the option market anticipates crashes; and 3) the method is particularly useful in periods when a bimodal distribution is to be expected.

We accept the first hypothesis (better than Black/Scholes), but with reservations. The method gives a better in-sample fit to observed option prices and its forecasting performance out-of-sample over 1987 to 1997 is also better but not by enough to be economically useful.

We reject hypothesis 2 (that the options market anticipates crashes). Neither before the large crash of 1987 nor before the much smaller crashes of 1989 and 1997 did the options market become more left-skewed. The upward adjustment of the stockmarket after sterling left the ERM in September 1992 was also not anticipated. Generally, we can say that the index-option market *reacts* to crucial events such as stock market crashes, it does not *predict* them.

We weakly accept hypothesis 3: the method does help to reveal market sentiment during elections. In particular, during the 1987 election the method allows us to reveal the development of a bimodal distribution, reflecting widely different potential outcomes. Nevertheless, while this may help in telling a “market story”, it is not one which is consistent with rational expectations: in 1992 and 1997 the election outcome was much more uncertain than in 1987, but a bimodal shape failed to appear. In particular, the market rose on the unexpected Conservative win in 1992, but the options had not shown that a jump was at all likely.

In sum, implied distributions (recovered by using the two-lognormal mixture technique) provide some potential insight into stockmarket sentiment, but their forecasting performance is not markedly better than that of Black/Scholes. Similar conclusions were reached for the US market (using different methods) by Dumas, Fleming and Whaley (1998). These empirical results cast doubt on the view that the shape of the implied distribution is a rational expectation. Fundamentally, what has to be explained is why the implied distribution is so left-skewed and why its shape changes so frequently? The most plausible explanation is portfolio-insuring behaviour (see Grossman and Zhou, 1996) and that does not require implied distributions to be good forecasts: they just need to reflect recent moves in the stockmarket and particular investor preferences.

Table 1: Dispersion and shape statistics of implied distributions for the period 1987-1997

Period	Volatility (%)	skewness	kurtosis
1987-89	22.6	-0.075	1.970
1990-91	21.9	-0.064	2.710
1992-93	17.0	-0.252	1.261
1994-95	15.6	-0.353	0.417
1996-97	14.5	-0.652	1.134
All years	18.7	-0.257	1.544

Notes:

The data are for one day per month, averaged over the periods shown. Data up to (and including) 1993 are for American options, thereafter for European options. The skewness and kurtosis have been measured by deducting the appropriate values for a lognormal distribution, hence the null hypothesis for each is zero. Data run until November of 1997 only.

Table 2: Root-mean-squared errors of the two-lognormal and the Black/Scholes models

<i>Period</i>	<i>Fitting (ex-post)</i>		<i>Forecasting (ex-ante)</i>	
	<i>Two-lognormal</i>	<i>Black/Scholes</i>	<i>Two-lognormal</i>	<i>Black/Scholes</i>
<i>1/87-2/92 (American-style options)</i>	2.15	3.02	4.39	4.94
<i>3/92-11/97 (European-style options)</i>	0.28	2.46	1.83	3.22

Notes:

The root-mean-squared errors are measured in index points for the option prices. The data are for one day per month, averaged over the periods shown.

Table 3: Period average statistics on implied distributions for four crash periods

Event	Trading Dates	No. of trading days	change in spot	volatility	RMSE of B/S	RMSE of 2-lognormal	skewness relative to lognormal	kurtosis relative to lognormal
crash of 1987	1/10/87-16/10/87	12	-	0.208	2.20	2.00	0.363	1.736
	22/10/87-10/11/87	28	-28.2%	0.507	3.72	3.14	-0.263	0.020
mini-crash of 1989	11/9/89-13/10/89	25	-	0.193	3.65	2.58	-0.264	3.695
	16/10/89-10/11/89	20	-7.0%	0.278	4.33	2.19	-0.350	1.762
ERM crisis of 1992	24/8/92-15/9/92	16	-	0.229	2.91	0.70	-0.601	1.615
	16/9/92-15/10/92	22	+9.0%	0.234	2.81	0.75	-0.735	2.504
Asian crash of 1997	9/9/97-17/10/97	29	-	0.196	4.13	0.26	-0.814	0.025
	20/10/97-28/11/97	30	-5.9%	0.349	5.55	0.86	-0.556	1.050

Notes:

Data up to (and including) 2/92 are for American options, thereafter for European options. The skewness and kurtosis have been measured by deducting the appropriate values for a lognormal distribution, hence the null hypothesis for each is zero.

Table 4: Period average statistics on implied distributions for three election periods

Election	Trading Dates	No of trading days	change in spot	volatility	RMSE of B/S	RMSE of 2-lognormal	skewness relative to lognormal	kurtosis relative to lognormal
1987	12/5/87-11/6/87	22	+4.4%	0.265	3.63	2.14	-0.606	0.620
	12/6/87-23/6/87	8	+4.1%	0.219	2.83	1.38	-0.090	0.893
1992	12/3/92-9/4/92	21	-4.1%	0.234	2.80	1.77	-0.200	1.960
	10/4/92-30/4/92	13	+7.6%	0.168	3.84	3.27	-0.118	5.527
1997	18/3/97-1/5/97	31	-1.9%	0.154	4.59	0.16	-0.532	0.527
	2/5/97-30/5/97	19	+7.5%	0.130	2.86	0.29	-0.445	0.295

Notes:

Data up to (and including) 2/92 are for American options, thereafter for European options. The skewness and kurtosis have been measured by deducting the appropriate values for a lognormal distribution, hence the null hypothesis for each is zero.

Figure 1 Skewness of the Implied Distributions




Figure 2 Implied Distributions Around the Crash of 1987






Figure 3 Implied Distributions Around the Crash of 1989





Figure 4 Implied Distributions Around the Sterling Crisis of 1992





Figure 5 Implied Distributions Around the Asian Crash of 1997





Figure 6 Implied Distributions Around the Election of 1987





Figure 7 Implied Distributions Around the Election of 1992







Figure 8 Implied Distributions Around the Election of 1997



References

- Ait-Sahalia, Y and A Lo (1995): "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices". NBER Working Paper No. 5351.
- Bahra, B (1997): "Implied Risk-Neutral Probability Density Functions From Option Prices: Theory and Application". Bank of England Working Paper No. 66, July.
- Bates, D (1991): "The Crash of '87: Was it expected? The evidence from Options Markets". *Journal of Finance*, 46, July, pp 1009–44.
- Bates, D (1996): "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in PHLX Deutschemark Options", *Review of Financial Studies*, 9, pp 69–107.
- Black, F and M Scholes (1973): "The Pricing of Options and Corporate Liabilities". *Journal of Political Economy*, 81, pp 637–59.
- Breeden, D and R H Litzenberger (1978): "Prices of State-Contingent Claims Implicit in Option Prices". *Journal of Business*, 9 (4), pp 621–51.
- Butler, C and H Davies (1998): "Assessing Market Views on Monetary Policy: the Use of Implied Risk Neutral Probability Distributions". Working Paper, Monetary Instruments and Markets Division, Bank of England.
- Campa, J and P Chang (1996): "Arbitrage-Based Tests Of Target-Zone Credibility: Evidence From ERM Cross-Rate Options". *American Economic Review*, 86 (4), pp 726–40.
- Coutant, S, E Jondeau and M Rockinger (1998): "Reading Interest Rate and Bond Futures Options' Smiles: How PIBOR and Notional Operators Appreciated the 1997 French Snap Election". Working Paper, HEC School of Management, Jouy-en-Josas, France.
- Crouhy, M and G Rockinger. (1993): "Asymmetric Conditional Heteroskedastic Processes with Hysteresis Effects". Working Paper, September, HEC, Paris.
- Dawson, P (1994): "Comparative Pricing of European and American Index Options: an Empirical Analysis". *Journal of Futures Markets*, 14 (3), pp 363–78.
- Dumas, B, J Fleming. and R Whaley (1998): "Implied Volatility Functions: Empirical Tests". *Journal of Finance*, 53, (6), pp 2059–2106.
- Gemmill, G (1992): "Political risk and market efficiency: Tests based in British stock and options markets in the 1987 election". *Journal of Banking and Finance*, 16, pp 211–31.
- Gemmill, G (1995): "Stockmarket Behaviour and Information in British Elections", Working Paper, City University Business School.
- Gemmill G (1996): "Did Option Traders Anticipate the Crash? Evidence from Volatility Smiles in the UK with US comparisons". *Journal of Futures Markets*, 16, December, pp 881–97.
- Grossman, S and Z Zhou (1996): "Equilibrium Analysis of Portfolio Insurance". *Journal of Finance*, 51 (4), pp 1379–1403.
- Harrison, J M and D M Kreps (1979): "Martingales and Arbitrage in Multi-Period Securities Markets". *Journal of Economic Theory*, 20, pp 381–408.
- Jackwerth, J (1997): "Implied Risk Aversion". Working Paper, London Business School.
- Jackwerth, J and M Rubinstein (1996): "Recovering Probability Distributions from Option Prices". *Journal of Finance*, 51, December, pp 1611–31.
- Longstaff, F (1995): "Option pricing and the Martingale Restriction". *Review of Financial Studies*, 8 (4), pp 1091–1124.
- Malz, A (1997): "Estimating the Probability Distribution of the Future Exchange Rate from Option Prices". *Journal of Derivatives*, 5 (2), pp 18-36.

- Malz, A (1996): “Option-Based Estimates of the Probability Distribution of Exchange Rates and Currency Excess Returns”, mimeo, Federal Reserve Bank of New York.
- Melick, W R and C P Thomas (1997): “Recovering an Asset’s Implied PDF from Option Prices: An Application to Crude Oil during the Gulf Crisis”. *Journal of Financial and Quantitative Analysis*, March, pp 91–115.
- Ritchey, R J (1990): “Call Option Valuation for Discrete Normal Mixtures”. *Journal of Financial Research*, 13 (4), 285–96.
- Rubinstein, M (1994): “Implied Binomial Trees”. *Journal of Finance*, 49 (3), July, pp 771–818.
- Shimko, D (1993): “Bounds of Probability”. *Risk*, 6 (4), April, 33–7.
- Shimko, D (1994): “A Tail of Two Distributions”, *Risk*, 7 (9), September, pp 39–45.
- Söderlind, P and L Svensson (1997): “New Techniques to Extract Market Expectations from Financial Instruments”. Working Paper No. 5877, NBER.

Discussion of Gemmill and Saflekos paper “How Useful are Implied Distributions? Evidence from Stock-Index Options”

Discussant: Shigenori Shiratsuka

The purpose of this paper is to estimate implied probability distributions (IPDs) as a mixture of two lognormal distributions by using British stock price index option data. It also empirically tests (1) the performance of pricing formula, (2) the forecastability of crashes, and (3) the revelation of market sentiments. The conclusions are as follows: (1) the two-lognormal method is better than the Black-Scholes model in fitting observed option prices (in-sample estimation), but there is no significant difference in predicting out-of-sample prices; (2) the shape of the IPDs does not indicate the forecastability of market crashes; and (3) the IPDs help to reveal changes in market sentiment.

Since I think that these conclusions are reasonable, I will make comments with a view to improving their robustness. My comments are divided into two parts: one is concerned with the technical issue of empirical procedures and the other is concerned with the more general issue of the case study methodology used to infer market sentiment from the IPDs.

Empirical Procedures

Regarding empirical procedures, I would like to point out two problems in this paper. First, I recommend that the authors use data of whole trading days in the forecasting exercise to compare the performance of different pricing formulas. The current estimations are done with data of one day in each month, probably because of the possible problem of autocorrelation caused by overlapping sample periods. However, considering that a standard method of adjusting for the effects of autocorrelation is available, I think it would be better to test forecasting power by using a larger sample of whole trading days.¹

Second, I am wondering if the methodology used for estimating the IPDs to examine the forecastability of crashes is appropriate. According to this paper, we can see large estimation errors in market turbulence, suggesting that the approximation of the IPDs by a mixture of two lognormal distributions might be too rigid a restriction for the period of market stress, which is thought to be an important time for case studies. Therefore, I am not sure that the estimation methodology of the IPDs

¹ See Newey and West (1987) for details of how to estimate autocorrelation-heteroskedasticity robust standard errors.

is appropriately selected for examining whether the IPDs might be useful for anticipating crashes and gauging market sentiment? In this case, it might be that a simpler but less restricted methodology would be better for examining market sentiment during the stress period.

Case Study Method

Let me turn next to the method of case studies used in examining market sentiment. Here, I would like to propose another way of deriving market sentiment or expectations from the estimated IPDs, based on my own recent research (Nakamura and Shiratsuka, 1999).

It is important to recognize that there is a risk of misreading the estimation errors as indicators of changes in market expectations, if we look at the shape of the IPD on a specific date. Alternatively, we can avoid this risk by observing the general trend of changes in each summary statistic that is shown by the distribution shape, and focus on examining the relationship between asset price movements and changes in distribution shape.² In this case, I would like to emphasize the effectiveness of a time-series plot of summary statistics for the estimated IPDs in the case studies for examining the changes in market sentiment over time.

We estimated the IPDs for the Nikkei 225 Stock Price Index Option from mid-1989 to 1996 on a daily basis (see Figure 1). By observing the time-series movements of underlying asset prices and summary statistics of IPDs, we found the typical patterns in linking large fluctuations of asset prices and response in the shape of IPDs as follows (see also Figure 2):

- (1) The standard deviation rises during a period of sharp decline in stock prices (high positive correlation with lagged absolute changes in stock prices), and a shift in level occurred at the end of 1989.
- (2) The skewness moves in the opposite direction to changes in market level, reflecting lags in the adjustment of market participants' confidence to the market levels (high negative correlation with simultaneous and lagged changes in stock prices).
- (3) The excess kurtosis jumps in the case of extreme price changes (positive correlation with simultaneous absolute changes in stock prices).

By comparing the actual movement of summary statistics with the aforementioned their typical patterns in response to the market fluctuations, the estimated IPDs provide us with a lot of information on the market sentiments and expectations. For example, the magnitude of changes in summary

² This paper computed summary statistics on original strike prices, and standardized summary statistics of IPDs for comparison. However, in our paper (Nakamura and Shiratsuka, 1999), we computed summary statistics on log-transformed strike prices, and did not need to standardize computed summary statistics for evaluating the divergence from normal position.

statistics tells us of the impact of external shocks. The length of adjustment periods provides us with information on the smoothness of adjustment of market expectations. Through the case studies of various episodes in Japanese financial markets during the period from 1989 to 1996, we have shown such usefulness of the IPDs as an information variable for monetary policy.

Conclusions

In summary, market participants' expectations are too diverse and informative to be captured merely by using a single summary statistic, i.e., the mean, because the same mean value implies different market expectations and policy implications, depending on the shape of the probability distribution of the expected outcome. In particular, since market participants' confidence in stock prices differs substantially depending on timing, we can expect to capture more market information, both in quality and in quantity, by carefully examining the changes in market participants' expectations that lie behind stock price fluctuations.

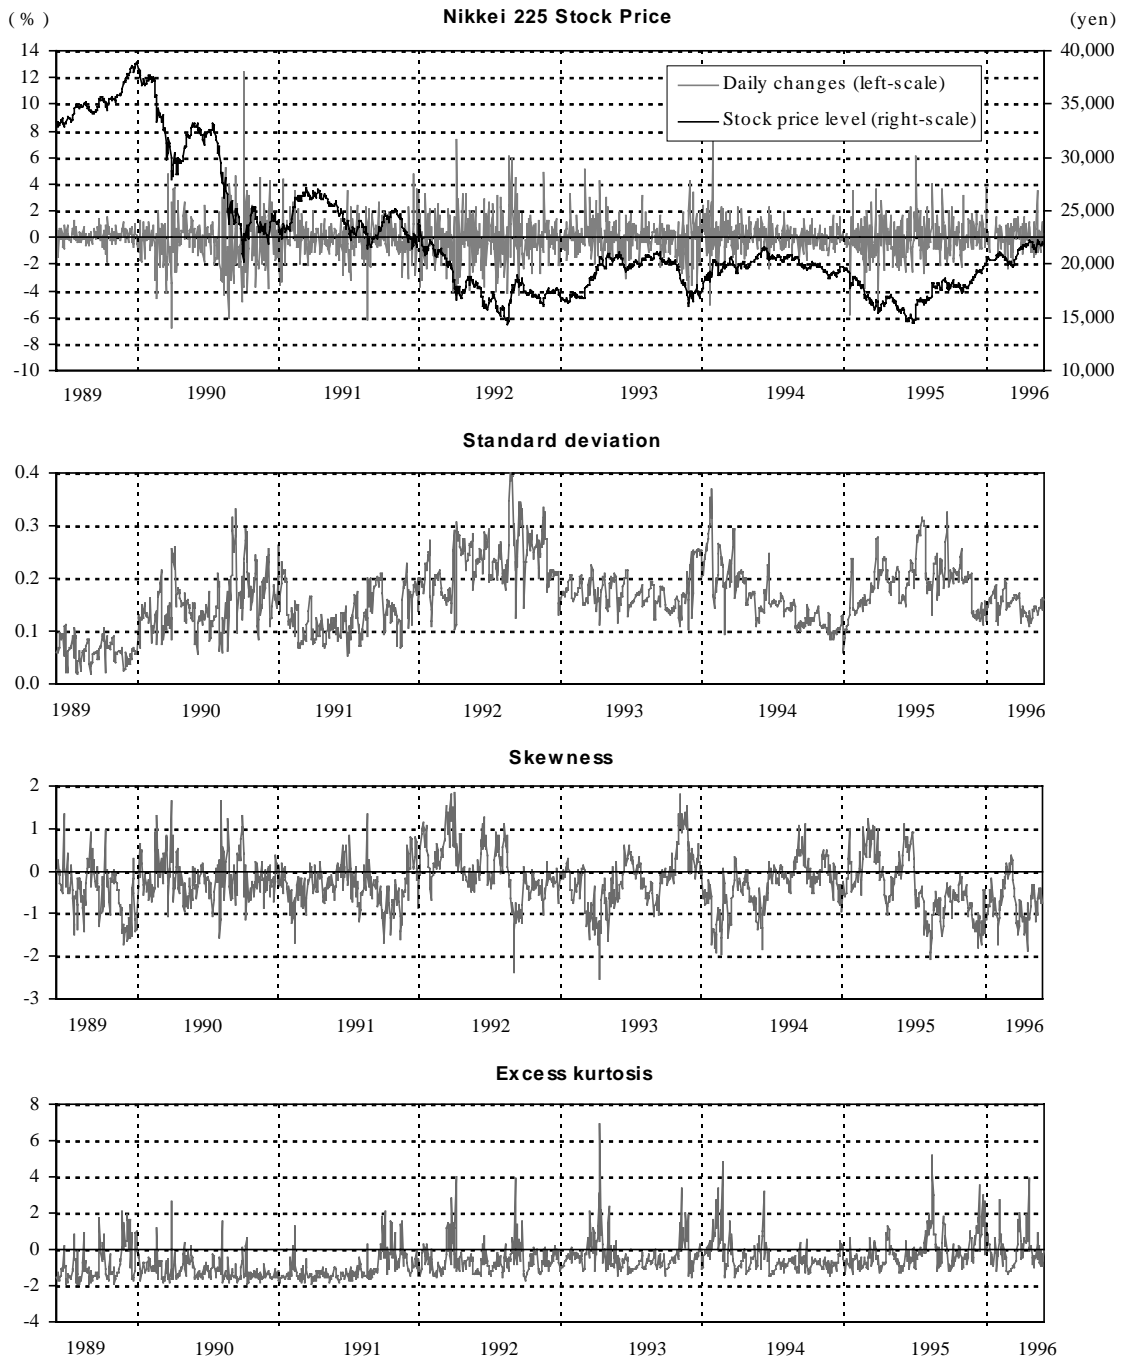
In this sense, as the papers and discussions contributed to this workshop suggest, the implied probability distributions extracted from option prices will provide useful information for the conduct of monetary policy. However, studies on how to make use of the information extracted from option prices in policy judgments have only just begun, and further research is necessary in this area.

References

Nakamura, Hisashi, and Shigenori Shiratsuka: "Extracting Market Expectations from Option Prices: Case Studies in Japanese Option Markets." *BOJ Monetary and Economic Studies*, 17 (1), 1999, pp. 1-43.

Newey, Whitney K and Kenneth D. West: "A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55 (3), 1987, pp. 703-08.


Figure 1. Stock Prices and Market expectations (Overview)



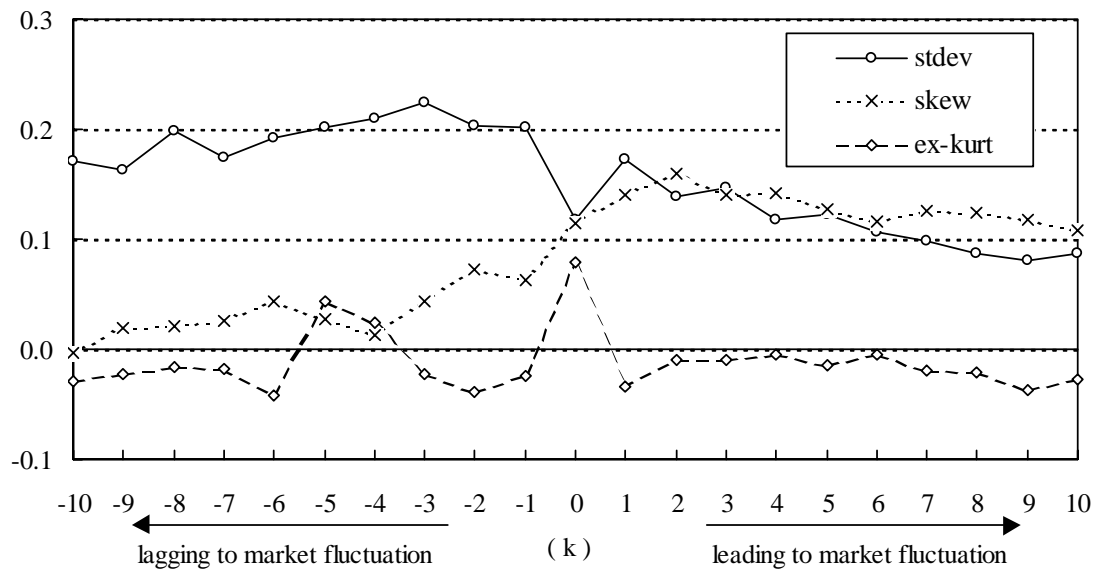
Source: Nakamura and Shiratsuka (1999)

Figure 2. Dynamic Cross-Correlation

(1) Market Changes (t) vs. Summary Statistics (t-k)



(2) Absolute Market Changes (t) vs Summary Statistics (t-k)



Discussion of
“How useful are implied distributions ? Evidence from stock-index options”
by G. Gemmill and A. Saflekos
Discussant: Raf Wouters

It is a pleasure for me to discuss this paper. The paper covers several different topics :

- the estimation method of the implied distribution ;
- the test of the information content of the implied distribution ;
- the effect of specific events on the implied distribution.

This diversity of topics makes it somewhat easier to discuss the paper as I can choose the topics that look most interesting to me. Since the estimation technique was discussed during the morning session, I will concentrate my remarks on the test of the information content of the implied distribution (i.d.).

The authors use an out-of-sample pricing prediction to test the information quality of the estimated i.d. This is a strong test for the i.d. as it depends on the whole structure of the distribution.

Traditional tests of the information content of option prices have concentrated on the test of the volatility measure only: they analysed whether the implied volatility was an unbiased predictor of the ex-post realisation of the average volatility over the option maturity. The general conclusion was that the implied volatility is a biased estimator because systematically overpredicting the realised volatility. On the other hand, the relation between implied volatility and the realised volatility was proved to be very significant and the implied volatility outperformed the forecasts of time-series based volatility measures such as Garch estimation techniques.

The out of sample pricing test used in this paper is a more general test : by pricing the whole range of options, the error will depend on both the volatility forecast but also on the higher moments describing the whole shape of the i.d.

The authors summarise these results by the RMSE of the forecast and compare this figure with the in sample errors. The mixture of two-lognormals delivers a strong improvement both in sample (but without correction for the number of parameters in the model) and out of sample compared to the B.S. model. These statistical gains are impressive at least for the period where European options are used (indicating perhaps that the method as applied here is not well suited for American style options).

This is a strong result and confirms similar conclusions in the literature (e.g. Dumas a.o. JF 1998). However, the authors minimise this finding: the improvement in value terms is not economic meaningful as it falls within the typical bid/ask spread in the market. They conclude from this that the forecasting performance is not markedly better than B.S. and that these results cast doubt on the hypothesis that the shape of the i.d. is a rational expectation.

In my opinion this conclusion is drawn somewhat too quickly : I would like to see more detailed results and add more tests:

- the average RMSE is too general: I would suggest for instance some statistic on the number of cases and the size of improvements in the fit in excess of the bid/ask spread (the outcome of some investment strategy would even be more conclusive). I would like to add here that by minimising the sample size to 10 observations, one probably makes it more difficult to find important improvements : so why not take the whole set of available prices into account?

- I wonder whether the errors are systematically related to strike prices and whether such systematic errors are higher / lower compared to the B.S. model

- I would also prefer to see predictions over a longer horizon: this can increase the possibility to find economic significant improvements;

- I would also compare the pricing error to a third model, for instance a simple approach based on a martingale assumption in which the pricing in the forecast is based on the observed volatility smile of the previous day.

Whatever the answers to these remarks, the important question remains as to why the prediction error is so much bigger out-of-sample as compared to the in-sample estimation error:

- a first explanation is the misspecification of the model based on the mixture of two-lognormals. However the small in-sample errors, typical for methods based on option prices of one maturity and at one point in time, provide not much hope that further improvements can be achieved within this approach;

- a second answer is the arrival of new information that changes not only the price of the underlying asset but also the volatility and the whole shape of the distribution. Under this hypothesis one can still test whether the forecast is unbiased and rational, but it will be statistically more difficult to find the answer. But more important, under this hypothesis the implied volatility and the i.d. can no longer be considered as constant and one should go in the direction of modelling the behaviour of these variables over time (Examples like Ait-Sahalia JF 1998 give promising results in this respect). But if one accepts that volatility is non-constant over time and that this variability is very important, one arrives quickly at the limit of methods based on a specification of the terminal distribution or other “non-structural” methods to estimate the option pricing function. These methods can never explain the economic logic behind the changes in the volatility or in the higher moments of the distribution, neither can they explain, in economic terms, the observed negative skewness and high kurtosis. Therefore one should move to structural approaches that model explicitly the dynamic process of the asset price and the volatility process. Within such a framework one can look for alternative explanations for these observations in the direction of :

- the role of the leverage effect;
- the role of non-normally distributed shocks;
- the variable risk-aversion, etc;

The test of such a model typically covers a broad range of option prices both cross-section and over time. The existence of higher in-sample errors in such application leaves more room to find economic significant differences in and out of sample;

- a third alternative explanation for the high prediction errors can be found in specific factors or characteristics of the option market. For instance through the existence of :

- bid-ask spreads that generate pricing errors;
- important liquidity problems due to the lower number of transactions in far out or in the money options;
- specific exposure or insurance arguments in the option market that distort the option prices.

This last type of explanations for the pricing errors make the use of the i.d. to investigate their information content problematic. Such distortions make the i.d. less useful or at least more difficult to interpret. Only option investors would still be interested to study these effects.

Now perhaps it is too strong to separate these different explanations underlying high prediction errors : if option market behaviour has a feedback to the underlying asset market, as one should expect in a general equilibrium framework, the different explanations can no longer be separated, and the interpretation becomes difficult in any case.

As long as our knowledge on the mechanisms that drive the movements in the distribution remains limited, we should be careful in using this information, and indeed consider it, as the authors suggest, only as some indication of market sentiment:

- simple volatility measures can be used to give an indication of the uncertainty of investors outlook. The use of implied volatility for analysing credibility of monetary policy typically falls in this category ;
- but the interpretation of the higher moments should be made very cautiously.

This conclusion also follows from the review in the paper of the i.d. behaviour around specific events. The i.d. does not show any systematic behaviour during these periods. But the exercises, and in particular the in-sample statistics, illustrate that the mixture of two lognormals does a reasonable job in describing the information of option prices during periods of high negative skewness, high kurtosis or bimodal distribution.

