On the usefulness of implied risk-neutral distributions – evidence from the Korean KOSPI 200 Index options market

In Joon Kim
Graduate School of Management, Korea Advanced Institute of Science and Technology, 207-43 Chongyang-ni, Dongdaemoon-gu, Seoul, Korea

Sol Kim
Graduate School of Management, Korea Advanced Institute of Science and Technology, 207-43 Chongyang-ni, Dongdaemoon-gu, Seoul, Korea

This study focuses on the usefulness of implied risk-neutral distributions. We compare the empirical performance of the Black and Scholes model, which assumes single lognormal distribution, with that of the option pricing model, which assumes a mixture of two lognormal distributions using three metrics: (1) in-sample performance, (2) out-of-sample performance, and (3) hedging performance.

We find that the option pricing formula using the two-lognormal mixture distributions model shows the best in-sample and out-of-sample pricing performance for short-term and long-term forecasting periods. For hedging, the differences between each model are not so large, but the Black and Scholes model is better than the two-lognormal mixture model, especially in the long term.

1 Introduction

Cross-sections of option prices have generally been used to estimate the implied risk-neutral distribution (henceforth RND). Because this RND represents the forward-looking view of the distribution of prices of the underlying asset, traders and policy-makers have been using the RND to assess market beliefs on future movements of the underlying asset.

Several research studies have proposed alternative methods to extract the RND from option prices. Estimation methods are divided into two categories. First, parametric methods assume specific distributional forms. In addition to the

We thank Tong Suk Kim, Suk Joon Byun, Chang Hyun Yoon, Geun Hyuk Chang and the anonymous referee for helpful comments and suggestions. We gratefully thank the Editor-in-Chief, Philippe Jorion, for detailed comments that have substantially improved the paper. Any remaining errors are our responsibility.

[1] Some functional forms that have been applied include the Burr III polynomial, a Hermite polynomial expansion around the lognormal density, a generalized beta, a normal inverse Gaussian and a Bayesian maximum entropy estimate using the lognormal density as the prior distribution.

Black–Scholes lognormal density, the most frequently used functional form for the RND has been the mixture of lognormal densities.[1] Second, non-parametric methods, attributed to Shimko (1991, 1993), are based on the equation of Breenden and Litzenberger (1978), implying that the risk-neutral density function is equal to the second derivative of the option price relative to the strike price. Putting this technique into operation has been far from straightforward, however, as the range of available strikes generally falls short of the continuity needed to produce reasonably smooth RND.[2]

Many papers focus on whether the RND provides an unbiased forecast of realized probability density functions, with mixed results.[3] The problem is that risk premiums may cause the RND to differ from the actual, physical distribution. Another line of research is to test whether such a RND is useful for forecasting option prices and for hedging options. This is the approach taken here.

Gemmill and Saflekos (2000) estimated a mixture of two lognormals from UK index options. They found that this method was much better than the single lognormal approach at fitting observed option prices, predicting, and hedging out-of-sample prices. However, an ad hoc procedure that merely smoothes Black–Scholes (1973) implied volatilities across exercise prices and times-to-expiration showed better performance than a mixture of two lognormals. In a study of US index options, Dumas, Fleming, and Whaley (1998) assessed the out-of-sample pricing and hedging validity of assuming that volatility is a deterministic function to deduce the shape of the RND. They found that an ad hoc Black and Scholes procedure (henceforth AHBS) performed better than a more complex model where volatility is modeled as a deterministic function for out-of-sample pricing performance. For hedging, the Black and Scholes model (henceforth BS) performed best of all the deterministic volatility specifications. Their contribution was to test whether the RND has forecasting and hedging potential in each market. In spite of using different methods to extract the RND, they showed similar results that AHBS showed relatively better performance. Our study verifies the results for the US and the UK, focusing on the long-term performance and delta-hedging strategies.

The purpose of this paper is to examine the usefulness of the RND estimated from the KOSPI 200 Index option prices. Introduced on July 7, 1997, the KOSPI 200 Index options market has become one of the biggest option markets in the world, despite its short history. During the three-year period from 1999 to 2001 the KOSPI 200 options market ranked first in the world in terms of trading volume. In 2002 its trading volume reached 1.9 billion contracts. This is the first article to address the usefulness of RND for this market.

---

[2] Examples of non-parametric estimation include Shimko (1993), Jackwerth and Rubinstein (1996), Malz (1997), Aït-Sahalia and Lo (1998), Aparicio and Hodges (1998), and Campa, Chang and Refalo (1999).

[3] See, for instance, Melick and Thomas (1997) for crude oil prices, Weinberg (2001) for S&P500 and currency prices, Anagnou, Bedendo, Hodges and Thompkins (2002) for S&P500 and the sterling/US dollar rate, Shiratsuka (2001) for Japanese stock index prices, Bliss and Panigirtzoglou (2002) for the FTSE 100 and S&P500.

Focusing on this market is also useful because there is an excellent liquidity in the near contract and because options are European-style, which facilitates pricing. This makes the KOSPI 200 Index options market an excellent market to investigate the potential mispricing of short-term options. As Bakshi, Cao and Chen (1997) mention, “The volatility smiles are the strongest for short-term options (both calls and puts), indicating that short-term options are the most severely mispriced by the BS [Black–Scholes model] and present perhaps the greatest challenge to any alternative option pricing model.”

We expand on the study of Gemmill and Saflekos (2000), but with several differences. We estimate the two-lognormal mixture (TLM) model using different parameters without the need to use parameters that are adjusted for the horizon. For reference, we also compare results given by the adjusted TLM method (ATLM) used by Gemmill and Saflekos (2000). We also use the traditional Black–Scholes method, as well as the ad hoc Black–Scholes (AHBS) method with an empirical fit to the volatility smile, as in Dumas, Fleming and Whaley (1998). Second, we consider one-week out-of-sample pricing and hedging performance to gauge parameter stability over long time periods and control for possible overfitting to the data. Third, Gemmill and Saflekos (2000) defined hedging errors as the difference between the change in the theoretical option price and the change in the market option price. In practice, option traders usually focus on the risk due to the underlying asset price variation alone and carry out a delta-neutral hedge employing only the underlying asset as a hedging instrument. We perform our hedging exercise consistently with this implementation.

The outline of the present work is as follows. The two-lognormal mixture model is reviewed in Section 2. Section 3 describes estimation methods. The data used for analysis are described in Section 4. Section 5 describes the pricing and hedging performance of each model. Finally, Section 6 concludes our study by summarizing the results.

2 Model

2.1 Two-lognormal mixture model

Since KOSPI 200 Index options have a limited range of strike prices on each day and lack the continuity to produce the smooth RND using non-parametric methods, parametric methods are more suitable and stable. Among them, many studies indicate the mixture of lognormal distributions as a good candidate to represent the RND, given its flexible specification that allows the approximation of quite a wide range of shapes. Bahra (1997), Melick and Thomas (2000) and Anagnou *et al.* (2002) used this distribution to extract the RND.

We assume that the density function $f(S_T)$ is given by a mixture of two lognormal density functions:

$$f(S_T) = \sum_{i=1}^{2} \left[\theta_i L(\alpha_i, \beta_i; S_T)\right] \tag{1}$$

where $L(\alpha_i, \beta_i; S_T)$ is the ith lognormal density with parameters α_i and β_i:

$$\alpha_i = \ln(S_0)\left(\mu_i - \frac{1}{2}\sigma_i^2\right)\tau \quad \text{and} \quad \beta_i = \sigma_i\sqrt{\tau} \quad \text{for} \quad i = 1, 2 \tag{2}$$

Given the assumption made on $f(S_T)$, the prices of European call and put options at time t can be expressed as follows:

$$C(K,\tau) = \mathrm{e}^{-r\tau}\int_K^{\infty}\left[\theta L(\alpha_1, \beta_1; S_T) + (1-\theta)L(\alpha_2, \beta_2; S_T)\right](S_T - K)\,\mathrm{d}S_T \tag{3}$$

$$P(K,\tau) = \mathrm{e}^{-r\tau}\int_0^{K}\left[\theta L(\alpha_1, \beta_1; S_T) + (1-\theta)L(\alpha_2, \beta_2; S_T)\right](S_T - K)\,\mathrm{d}S_T \tag{4}$$

Bahra (1997) derived closed-form solutions to Equations (3) and (4) as follows:

$$C(K,\tau) = \mathrm{e}^{-r\tau}\left\{\begin{array}{l} \theta\left[\exp\left(\alpha_1 + \frac{1}{2}\beta_1^2\right)N(d_1) - KN(d_2)\right] + \\ (1-\theta)\left[\exp\left(\alpha_2 + \frac{1}{2}\beta_2^2\right)N(d_3) - KN(d_4)\right] \end{array}\right\} \tag{5}$$

$$P(K,\tau) = \mathrm{e}^{-r\tau}\left\{\begin{array}{l} \theta\left[-\exp\left(\alpha_1 + \frac{1}{2}\beta_1^2\right)N(-d_1) + KN(-d_2)\right] + \\ (1-\theta)\left[-\exp\left(\alpha_2 + \frac{1}{2}\beta_2^2\right)N(-d_3) + KN(-d_4)\right] \end{array}\right\} \tag{6}$$

where

$$d_1 = \frac{-\ln K + \alpha_1 + \beta_1^2}{\beta_1}, \quad d_2 = d_1 - \beta_1, \quad d_3 = \frac{-\ln K + \alpha_2 + \beta_2^2}{\beta_2}, \quad d_4 = d_3 - \beta_2$$

In the absence of arbitrage opportunities, the mean of the RND function should equal the forward price of the underlying asset. In this sense, Bahra (1997) and Anagnou *et al.* (2002) treat the underlying asset as a zero-strike option and use the incremental information it provides by including its forward price as an additional observation in the minimization procedure. On the other hand, Gemmill and Saflekos (2000) and Bliss and Panigirtzoglou (2000) did not impose the constraint. The reason is that not imposing the constraint allowed them to see how closely the estimated RND conformed to the theoretical restriction on the mean – in effect, how well the underlying conditions for no arbitrage hold. Moreover this constraint will usually be binding and will degrade the goodness-of-fit. We also do not impose constraints.

The closed-form option pricing formula makes it possible to derive comparative statics and hedge ratios analytically. In the present context, we consider only a source of stochastic variations over time, the price risk S_t. The delta of call options is

$$\Delta_S = \theta e^{(\mu_1 - r)\tau} N(d_1) + (1-\theta) e^{(\mu_2 - r)\tau} N(d_3) \tag{7}$$

This two-lognormal mixture (TLM) model is simply the weighted sum of two Black–Scholes solutions, where θ is the weighting parameter, and α_1, β_1, α_2 and β_2 are the parameters of each of the component lognormal RND functions. Instead we estimate μ_1, σ_1, μ_2 and σ_2, which are parameters of the normal RND functions of the underlying asset return. Using parameters of the underlying asset return makes the out-of-sample pricing and hedging strategy easy. This is what we refer to as the TLM method.

In their out-of-sample pricing, Gemmill and Saflekos (2000) update the means and variances of the distributions to take account of changes in the stock price and decreasing time-to-maturity. We call an out-of-sample pricing model reflecting these adjustments the adjusted two-lognormal mixture model (ATLM).

$$\alpha_{i,t} = \alpha_{i,t-1} + \ln \frac{S_t}{S_{t-1}}, \quad \beta_{i,t} = \beta_{i,t-1} \sqrt{\frac{\tau_t}{\tau_{t-1}}} \tag{8}$$

The above adjustments are not perfect, however, because α_i does not consider decreasing time-to-maturity. Using parameters of the underlying asset return, as in Equation (2), makes the above adjustments unnecessary because α_i and β_i then automatically reflect changes in the stock price and decreasing time-to-maturity. This is why we believe the TLM method should be superior.

2.2 Ad hoc Black–Scholes procedure

Since the TLM method has more parameters than the Black–Scholes, it may have an unfair advantage over the latter for in-sample fitting. Therefore, we follow Dumas, Fleming and Whaley (1998) and construct the ad hoc Black–Scholes (AHBS) procedure, in which each option has its own implied volatility depending on the strike price.

Specifically, we adopt the following specification for the Black–Scholes implied volatilities:

$$\sigma_n = \beta_1 + \beta_2 \cdot \left(S/K_n\right) + \beta_3 \cdot \left(S/K_n\right)^2 \tag{9}$$

where σ_n is the implied volatility for the nth option with strike price K_n and spot price S.

We follow a four-step procedure. First, we abstract the Black–Scholes implied volatility from each option. Second, we estimate the β_i, $i = 1, 2, 3$, by ordinary least squares. Third, using estimated parameters from the second step, we plug in each option's moneyness into the equation and obtain the model-implied volatility for each option. Finally, we use volatility estimates computed in the third step to price options with the Black–Scholes formula. The AHBS, though theoretically inconsistent, can be a more challenging benchmark than the simple Black–Scholes for any competing option valuation model.

3 Estimation procedure

As is the standard practice, we estimate the parameters of each model every sample day. Since closed-form solutions are available for the option price, a natural candidate for the estimation of parameters which enter the pricing formula is a non-linear least-squares procedure involving minimization of the sum of squared errors between the model and market prices.

Let $O_i^*(t, \tau; K)$ denote the model price of option i on day t, and let $O_i(t, \tau; K)$ denote the market price of the option i on day t. We minimize the sum of squared errors between the model and market prices:

$$\min_{\phi_t} \sum_{i=1}^{N} \left[O_i^*(t, \tau; K) - O_i(t, \tau; K) \right]^2, \qquad t = 1, \ldots, T \tag{10}$$

where N denotes the number of options on day t, T denotes the number of days in the sample.

For the Black–Scholes model, the volatility parameter, σ, is estimated. For the TLM, we estimate the structural parameters $\{\mu_1, \sigma_1, \mu_2, \sigma_2, \theta\}$ of the normal RND functions of the underlying asset return. For the AHBS, coefficients are estimated via ordinary least squares, minimizing the sum of squared errors between the Black–Scholes implied volatilities across different strikes and the model's functional form of the implied volatility.

4 Data

The KOSPI 200 Index options market has three consecutive near-term delivery months and one additional month from the quarterly cycle (March, June, September and December). The expiration day is the second Thursday of each contract month. Each options contract month has at least five strike prices. The number of strike prices may, however, increase according to the price movement. Trading in KOSPI 200 options is fully automated. The exercise style of the options is European and thus contracts can be exercised only on the expiration dates. Hence, our test results are not affected by the complication that arises due to the early exercise feature of American options. Moreover, it is important to note that liquidity is concentrated in the nearest expiration contract.

The sample period extends from January 3, 1999, through December 26, 2000. Minute-by-minute transaction prices for the KOSPI 200 options are obtained from the Korea Stock Exchange. The three-month treasury yields were used as risk-free interest rates.[4] Because KOSPI 200 contracts are European-style, index levels were adjusted for dividend payments before each option's expiration date.[5] The KOSPI 200 Index pays dividends only at the end of March, June, September and December, which are used for adjustment dates.

The following rules are applied to filter data needed for the empirical test. For each day in the sample only the last reported transaction price, which has to occur between 2.30 and 2.50pm[6] of each option contract, is employed in the empirical test. The tight time window is chosen to minimize problems stemming from intra-day variation in volatility.[7]

An option of a particular moneyness and maturity is represented only once in the sample. In other words, although the same option may be quoted again during the time window, only the last record of that option is included in our sample.

As options with less than six days to expiration may induce biases due to their low prices and bid–ask spreads, they are excluded from the sample. The maturity of options is identical each day because the estimation of the RND must be applied to a specific maturity. Because the liquidity of KOSPI 200 option contracts is concentrated in the nearest expiration contract, the maturity of options is not more than 37 days.

To mitigate the impact of price discreteness on option valuation, prices lower than 0.5 are not included.

Prices not satisfying the arbitrage restriction are excluded:[8]

$$C_{t,\tau} \geq S_t - \sum_{s=1}^{\tau} \mathrm{e}^{-R_{t,s}s} D_{t+S} - KB_{t,\tau} \tag{11}$$

$$P_{t,\tau} \geq KB_{t,\tau} - S_t + \sum_{s=1} \mathrm{e}^{-R_{t,s}s} D_{t+S} \tag{12}$$

where $B_{t,\tau}$ is a zero-coupon bond that pays 1 in τ periods from time t and D_t is daily dividend at time t.

4 Korea does not have a liquid Treasury bill market, so the three-month Treasury yield is used in spite of the mismatch of maturity of options and interest rates.

5 We assume that traders have perfect knowledge about future dividend payments because options in this study have short time-to-maturities.

6 In the Korea stock market there are simultaneous bids and offers from 2.50pm.

7 Because the recorded KOPSI 200 Index values are not the daily closing index levels, there is no non-synchronous price issue here, except that the KOSPI 200 Index level itself may contain stale component stock prices at each point in time.

8 Based on this criterion, 77 observations (approximately 1.1% of the original sample) are eliminated.

TABLE 1 KOSPI 200 Index options data.

	1999		2000		All	
Moneyness (S/K)	**Call**	**Put**	**Call**	**Put**	**Call**	**Put**
<0.94	1.62	11.48	1.20	12.69	1.38	12.30
	(447)	(264)	(611)	(556)	(1058)	(820)
0.94–0.97	2.84	6.49	2.19	6.28	2.52	6.38
	(271)	(244)	(273)	(265)	(544)	(509)
0.97–1.00	4.04	4.73	3.24	4.37	3.64	4.56
	(283)	(279)	(282)	(273)	(565)	(552)
1.00–1.03	2.39	3.29	4.62	3.05	5.02	3.17
	(255)	(261)	(235)	(245)	(490)	(506)
1.03–1.06	7.06	2.18	6.11	2.04	6.59	2.11
	(197)	(243)	(193)	(232)	(390)	(475)
≥1.06	11.95	1.30	9.28	1.25	10.88	1.28
	(371)	(643)	(250)	(447)	(621)	(1081)
Subtotal	5.39	4.23	3.70	5.79	4.54	5.03
	(1824)	(1925)	(1844)	(2018)	(3668)	(3943)

This table reports average option prices, and the number of options (shown in parentheses), for each moneyness and type (call or put) category. The sample period is January 3, 1999, to December 26, 2000. Daily information from the last transaction prices (prior to 2.50pm) of each option contract is used to obtain the summary statistics. The moneyness of an option is defined as S/K, where S denotes the spot price and K denotes the strike price.

We divide the option data into several categories according to the moneyness, S/K. Table 1 describes certain sample properties of the KOSPI 200 option prices used in the study. Summary statistics are reported for the option price and the total number of observations, according to each moneyness/option-type category. Note that there are 3,668 call- and 3,943 put-option observations, with deep out-of-the-money options, respectively, taking up 29% for call and 27% for put.

5 Empirical findings

Figure 1 depicts the representative probability density functions obtained from the Black–Scholes and TLM methods; the top panel uses estimates from the January contract on December 20, and the bottom panel uses estimates from the July contract on June 26. Non-normal skewness and kurtosis are reflected in the shapes of the RND from the TLM.

We compare the empirical performance of alternative models using three metrics: in-sample performance, out-of-sample performance, and hedging performance.

The analysis is based on four measures: mean absolute errors (MAE), mean percentage errors (MPE), mean absolute percentage errors (MAPE), and mean squared errors (MSE). The MAE and MAPE measure the magnitude of the pricing error, while the MPE indicates the direction of the pricing error. The MSE measures the volatility of errors. In the remaining sections, we mainly deal with

FIGURE 1 Implied risk-neutral distributions.

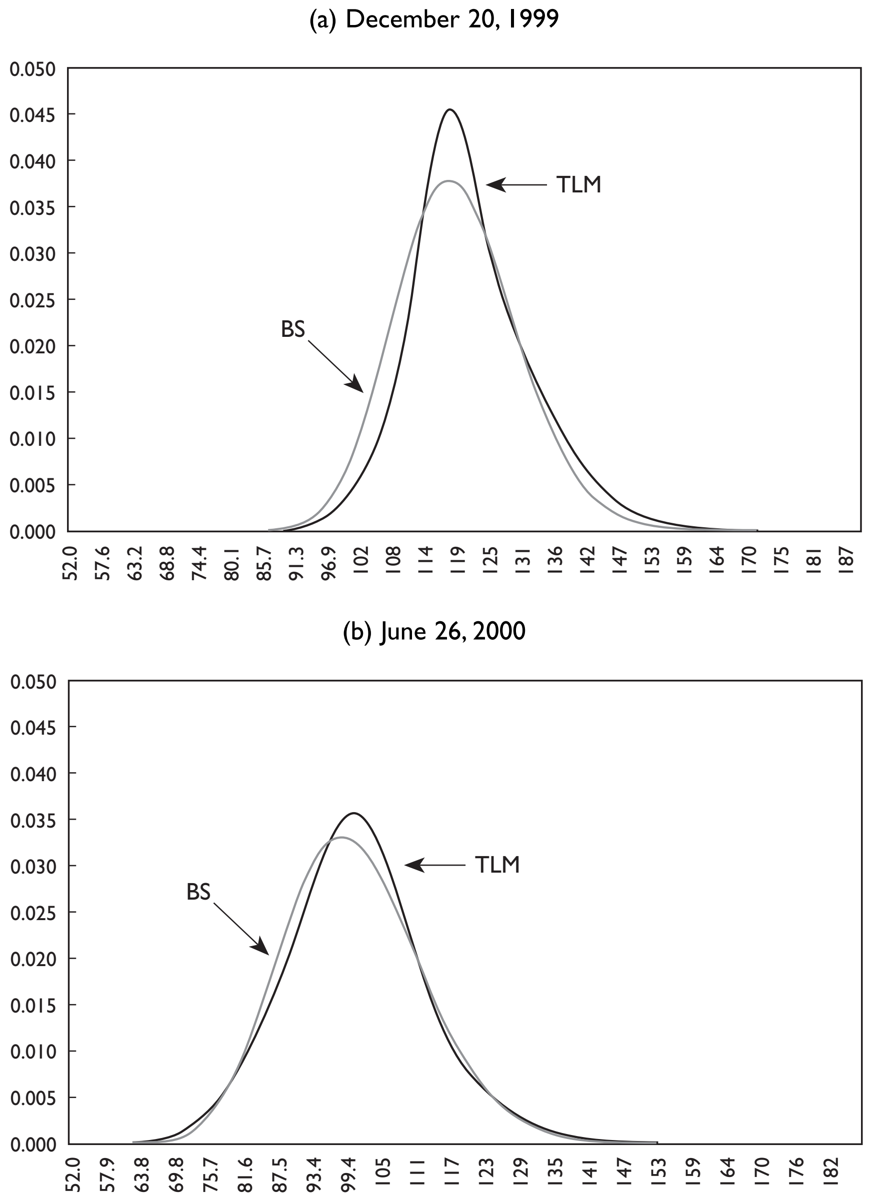

Figure 1 depicts the representative probability density functions obtained from the Black–Scholes (BS) and the two-lognormal mixture (TLM) models; the top panel uses estimates from the January contract on December 20, 1999, and the bottom panel uses estimates from the July contract on June 26, 2000.

the MAPE because the relative comparison considering each option price is important above all else.

5.1 In-sample performance

We evaluate the in-sample performance of each model by comparing the market prices with the model prices computed by using the parameter estimates from the current day. Table 2 reports the in-sample valuation errors for alternative models computed over the whole sample of options as well as across six moneyness and two option-type categories. Results from the analysis are as follows.

First, with respect to all measures, the TLM shows the best performance, followed by the AHBS. The risk-neutral distribution of options is explained by a mixture of two lognormal distributions better than by a single lognormal distribution. Next, the AHBS is not much better than the Black–Scholes method even though it has more parameters. This result can be explained by the lower R^2 compared to other option markets. In our study, the R^2 of the AHBS is 31% on average, which is quite low. In the study by Kirgiz (2001) using S&P500 data, the R^2 was 93%. Because of the low R^2, the AHBS seems to lead to relatively large in-sample errors.

Second, all models show moneyness-based valuation errors. The fit of the models, as measured by the MAPE, is worst for out-of-the-money options and steadily improves as we move from out-of-the-money to in-the-money options. The worse fit for out-of-the-money options is partly the result of the objective function used to estimate structural parameters of models. This function gives more weight to the relatively expensive in-the-money options.[9]

To sum up, the TLM shows the best in-sample performance. But the AHBS used by practitioners does not show much better performance than Black–Scholes in spite of the large number of parameters.

5.2 Out-of-sample performance

In-sample performance can be biased due to the potential problem of overfitting to the data. A good in-sample fit might be a consequence of having an increasingly large number of structural parameters. This is why we turn to examining the model's out-of-sample pricing performance. In out-of-sample pricing, the presence of more parameters may actually cause overfitting. We can also check each model's parameter stability over time using the current day's estimated structural parameters to price options on the next day and on the next week. For out-of-sample pricing, we use the actual asset price, interest rate and time-to-maturity on the target date.

Tables 3 and 4 report one-day and one-week ahead out-of-sample valuation errors for alternative models computed. Consider one-day ahead out-of-sample pricing errors. The pure TLM shows the best performance, closely followed by the ATLM. Using adjusted structural parameters does not show better performance than using untouched parameters in this study. In the moneyness-based error,

[9] The re-estimation that minimizes the sum of percentage-squared errors does not settle the highest pricing errors of out-of-the-money options.

the TLM is also the best for all moneyness. The ATLM performs better than other models except the deep in-the-money for calls and puts. In one-week ahead out-of-sample pricing, the order of the models is changed. The TLM shows the best performance as usual. But the Black–Scholes model is the second and the ATLM

TABLE 2 In-sample pricing errors.

		<0.94	0.94–0.97	0.97–1.00	1.00–1.03	1.03–1.06	≥1.06	All
				Panel A: Calls				
	BS	0.0183	0.0087	–0.0076	–0.0076	0.0037	0.0203	0.0082
MPE	AHBS	–0.0560	0.0363	0.0248	0.0091	0.0038	0.0000	–0.0053
	TLM	–0.0102	0.0060	–0.0003	–0.0083	–0.0085	–0.0016	–0.0044
	BS	0.1802	0.1431	0.1126	0.0935	0.0769	0.0585	0.1211
MAPE	AHBS	0.2257	0.1707	0.1200	0.0862	0.0644	0.0372	0.1336
	TLM	0.0764	0.0425	0.0315	0.0309	0.0280	0.0262	0.0447
	BS	0.2368	0.3442	0.3896	0.4620	0.5065	0.6585	0.4064
MAE	AHBS	0.2668	0.3787	0.4119	0.4386	0.4312	0.3719	0.3640
	TLM	0.0844	0.0925	0.1116	0.1548	0.1926	0.3181	0.1503
	BS	0.1119	0.2242	0.2887	0.3910	0.5026	1.4244	0.4568
MSE	AHBS	0.1431	0.2412	0.3020	0.3642	0.4009	0.3070	0.2668
	TLM	0.0162	0.0200	0.0320	0.0646	0.1219	0.5744	0.1314
				Panel B: Puts				
	BS	0.0101	–0.0389	–0.0615	–0.0746	–0.0825	0.0482	–0.0179
MPE	AHBS	0.0007	–0.0280	–0.0371	–0.0488	–0.0983	–0.0425	–0.0384
	TLM	–0.0021	–0.0030	–0.0037	0.0070	0.0047	0.0101	0.0029
	BS	0.0465	0.0822	0.1067	0.1359	0.1565	0.1791	0.1206
MAPE	AHBS	0.0343	0.0723	0.0993	0.1436	0.1995	0.1673	0.1187
	TLM	0.0197	0.0270	0.0311	0.0335	0.0533	0.0902	0.0474
	BS	0.5523	0.4981	0.4397	0.3730	0.2734	0.2106	0.3793
MAE	AHBS	0.3809	0.4487	0.4270	0.4002	0.3264	0.1929	0.3405
	TLM	0.2327	0.1708	0.1383	0.1004	0.0913	0.0898	0.1383
	BS	0.5847	0.4543	0.3378	0.2457	0.1369	0.0828	0.2983
MSE	AHBS	0.3135	0.4034	0.3164	0.2644	0.1884	0.0824	0.2408
	TLM	0.1417	0.0714	0.0437	0.0257	0.0195	0.0201	0.0560

This table reports in-sample pricing errors for KOSPI 200 Index options with respect to moneyness, S/K, where S is the asset price and K is the strike price. Each model is estimated every day during the sample period and in-sample pricing errors are computed using estimated parameters from the current day. Denoting $\varepsilon_n = O_n^* - O_n$, where O_n^* is the model price and O_n is the market price, pricing performance is evaluated by: (1) mean percentage error (MPE), $\left(\sum_{n=1}^{N}\varepsilon_n/O_n\right)/N$; (2) mean absolute percentage error (MAPE), $\left(\sum_{n=1}^{N}|\varepsilon_n|/O_n\right)/N$; (3) mean absolute error (MAE), $\left(\sum_{n=1}^{N}|\varepsilon_n|\right)/N$; and (4) mean squared error (MSE), $\left(\sum_{n=1}^{N}(\varepsilon_n)^2\right)/N$; where N is the total number of options in a particular moneyness category. BS denotes the Black and Scholes model, AHBS denotes the ad hoc Black and Scholes procedure that fits the implied volatility surface and TLM denotes the two-lognormal mixture model.

TABLE 3 One-day ahead out-of-sample pricing errors.

		<0.94	0.94–0.97	0.97–1.00	1.00–1.03	1.03–1.06	≥1.06	All
					Panel A: Calls			
MPE	BS	0.0384	0.0123	–0.0066	–0.0086	0.0050	0.0190	0.0144
	AHBS	–0.0239	0.0392	0.0231	0.0075	0.0061	0.0026	0.0047
	TLM	0.0160	0.0138	–0.0001	–0.0105	–0.0096	0.0048	0.0049
	ATLM	0.0611	0.0118	–0.0026	–0.0118	–0.0089	0.0102	0.0179
MAPE	BS	0.2169	0.1575	0.1250	0.0983	0.0810	0.0602	0.1363
	AHBS	0.2588	0.1692	0.1269	0.0966	0.0772	0.0554	0.1485
	TLM	0.1868	0.1274	0.1013	0.0816	0.0681	0.0507	0.1143
	ATLM	0.1955	0.1458	0.1187	0.0932	0.0789	0.0575	0.1261
MAE	BS	0.2725	0.3782	0.4284	0.4873	0.5366	0.6761	0.4395
	AHBS	0.3044	0.3992	0.4432	0.4870	0.5152	0.6223	0.4423
	TLM	0.2245	0.2870	0.3359	0.3893	0.4405	0.5623	0.3551
	ATLM	0.2393	0.3283	0.3918	0.4421	0.5061	0.6304	0.3998
MSE	BS	0.1327	0.2544	0.3140	0.4190	0.5368	1.4657	0.4923
	AHBS	0.1653	0.2894	0.3455	0.4263	0.5160	1.2857	0.4793
	TLM	0.1022	0.1508	0.2043	0.2662	0.3924	1.1460	0.3602
	ATLM	0.1138	0.2001	0.2937	0.3461	0.4845	1.2613	0.4248
					Panel B: Puts			
MPE	BS	0.0122	–0.0385	–0.0608	–0.0782	–0.0878	0.0477	–0.0190
	AHBS	0.0053	–0.0288	–0.0381	–0.0542	–0.0983	–0.0311	–0.0351
	TLM	0.0065	–0.0037	–0.0079	–0.0012	–0.0121	0.0063	–0.0002
	ATLM	0.0066	–0.0077	–0.0192	–0.0188	–0.0287	–0.0032	–0.0090
MAPE	BS	0.0505	0.0893	0.1133	0.1522	0.1826	0.1968	0.1323
	AHBS	0.0488	0.0850	0.1069	0.1518	0.2136	0.2080	0.1372
	TLM	0.0475	0.0669	0.0838	0.1087	0.1399	0.1770	0.1086
	ATLM	0.0535	0.0782	0.0939	0.1166	0.1469	0.1804	0.1156
MAE	BS	0.5979	0.5439	0.4722	0.4168	0.3188	0.2300	0.4183
	AHBS	0.5747	0.5265	0.4607	0.4262	0.3532	0.2339	0.4160
	TLM	0.5597	0.4205	0.3703	0.3157	0.2554	0.1963	0.3502
	ATLM	0.6254	0.4862	0.4113	0.3412	0.2756	0.2028	0.3859
MSE	BS	0.7232	0.5541	0.4077	0.3121	0.1868	0.1066	0.3751
	AHBS	0.6752	0.5498	0.3952	0.3295	0.2335	0.1150	0.3728
	TLM	0.5828	0.3319	0.2460	0.1858	0.1213	0.0746	0.2608
	ATLM	0.6875	0.3941	0.3010	0.2133	0.1402	0.0774	0.3054

This table reports one-day ahead out-of-sample pricing errors for KOSPI 200 Index options with respect to moneyness, S/K. Each model is estimated every day during the sample period and one-day ahead out-of-sample pricing errors are computed using estimated parameters from the previous trading day. Pricing performance evaluated as in Table 2. BS, Black–Scholes model; AHBS, ad hoc Black–Scholes procedure that fits the implied volatility surface; TLM, two-lognormal mixture model; ATLM, adjusted TLM model.

TABLE 4 One-week ahead out-of-sample pricing errors.

		<0.94	0.94–0.97	0.97–1.00	1.00–1.03	1.03–1.06	≥1.06	All
					Panel A: Calls			
MPE	BS	0.1118	0.0319	0.0123	0.0064	0.0167	0.0250	0.0431
	AHBS	0.0782	0.0549	0.0342	0.0187	0.0191	0.0208	0.0423
	TLM	0.1159	0.0523	0.0358	0.0170	0.0194	0.0232	0.0522
	ATLM	0.1162	–0.0056	–0.0159	–0.0168	0.0002	0.0222	0.0288
MAPE	BS	0.3049	0.1936	0.1533	0.1089	0.0860	0.0581	0.1668
	AHBS	0.3457	0.2031	0.1557	0.1090	0.0856	0.0579	0.1791
	TLM	0.2997	0.1894	0.1487	0.1001	0.0800	0.0534	0.1613
	ATLM	0.3139	0.2260	0.1789	0.1280	0.1014	0.0699	0.1844
MAE	BS	0.3464	0.4120	0.4831	0.5141	0.5561	0.6660	0.4844
	AHBS	0.3941	0.4369	0.4967	0.5187	0.5558	0.6629	0.5025
	TLM	0.3400	0.3967	0.4628	0.4578	0.5060	0.6073	0.4528
	ATLM	0.3636	0.4543	0.5350	0.5656	0.6289	0.7597	0.5362
MSE	BS	0.2085	0.3132	0.3842	0.4385	0.4986	1.5537	0.5713
	AHBS	0.2856	0.3624	0.4139	0.4608	0.5197	1.5432	0.6061
	TLM	0.2049	0.3000	0.3520	0.3433	0.4208	1.3272	0.4985
	ATLM	0.2325	0.3709	0.4851	0.5396	0.6478	1.6719	0.6546
					Panel B: Puts			
MPE	BS	0.0210	–0.0248	–0.0383	–0.0665	–0.0711	0.1063	0.0027
	AHBS	0.0175	–0.0175	–0.0234	–0.0480	–0.0633	0.0429	–0.0060
	TLM	0.0160	0.0031	0.0096	–0.0003	0.0008	0.0731	0.0223
	ATLM	0.0044	–0.0197	–0.0318	–0.0736	–0.0958	–0.0146	–0.0311
MAPE	BS	0.0539	0.1003	0.1296	0.1962	0.2433	0.3089	0.1717
	AHBS	0.0547	0.0999	0.1306	0.1982	0.2542	0.3235	0.1770
	TLM	0.0564	0.0942	0.1160	0.1783	0.2199	0.2997	0.1621
	ATLM	0.0708	0.1189	0.1480	0.2110	0.2571	0.3302	0.1894
MAE	BS	0.6231	0.6187	0.5362	0.5332	0.4269	0.3531	0.5121
	AHBS	0.6226	0.6230	0.5484	0.5432	0.4426	0.3615	0.5195
	TLM	0.6504	0.6032	0.5125	0.5124	0.4054	0.3422	0.5050
	ATLM	0.7904	0.7348	0.6190	0.5675	0.4531	0.3623	0.5890
MSE	BS	0.7284	0.6824	0.4837	0.4779	0.3081	0.2120	0.4836
	AHBS	0.7331	0.7182	0.5039	0.5013	0.3477	0.3054	0.5219
	TLM	0.7460	0.6233	0.4574	0.4671	0.3239	0.2213	0.4783
	ATLM	1.0132	0.8727	0.6559	0.5682	0.3723	0.2452	0.6286

This table reports one-week ahead out-of-sample pricing errors for KOSPI 200 Index options with respect to moneyness, S/K. Each model is estimated every day during the sample period and one-week ahead out-of-sample pricing errors are computed using estimated parameters from one week ago. Pricing performance evaluated as in Table 2. BS, Black–Scholes model; AHBS, ad hoc Black–Scholes procedure that fits the implied volatility surface; TLM, two-lognormal mixture model; ATLM, adjusted TLM model.

is the worst. The performance of the ATLM deteriorates because the adjusted parameters of the ATLM do not reflect the decreasing time to maturity. The Black–Scholes model exhibits a good fit for the one-week out-of-sample pricing, which demonstrates the robustness of simplicity.

The one-week forecasts, however, are overlapping. There are about 500 working days in the sample, which should give about 499 independent one-day ahead forecasts. However, we have only 25 option contracts with different maturities from January 1999 to January 2001.[10] The fact that one model forecasts badly for the February options today may not be independent of its bad forecasts for the February options yesterday. We check whether forecasting results are consistent by contracts using the MAPE metric. For one-day out-of-sample pricing errors, the TLM is better than the Black–Scholes except for three contracts (July 1999, February 2000 and November 2000). For one-week pricing errors, the TLM is better than Black–Scholes except for four contracts (August 1999, May 2000, August 2000 and October 2000). The results, therefore, do not seem to be affected by the overlap.

We also find that the difference between the Black–Scholes and the TLM methods becomes smaller for out-of-sample pricing. The ratio of MAPE from the Black–Scholes to the TLM is 2.709 (2.544) for in-sample errors of call (put) options. This ratio changes to 1.192 and 1.218 (1.034 and 1.059) for one-day (one-week) ahead out-of-sample errors. As the term of the out-of-sample pricing gets longer, the difference between the two models decreases. The strong pricing performance of the TLM does not keep up as the term of out-of-sample pricing gets longer, which implies that the market consensus of the RND is very volatile and structural parameters must be changed frequently.

5.3 Hedging performance

Hedging performance is particularly important to market practitioners as it is often used as a tool of risk management to cover the positions in the underlying asset and option markets. Gemmill and Saflekos (2000) defined hedging errors as the difference between the change in model option price and the change in market option price. We examine hedges in which a single instrument (ie, the underlying asset) can be employed. In practice, option traders usually focus on the risk due to the underlying asset price volatility alone and carry out a delta-neutral hedge employing only the underlying asset as the hedging instrument.

We implement hedging with the following method. Consider hedging a short position in an option, $O(t, \tau; K)$, with τ periods to maturity and strike price K. Let $\Delta_S = \partial O(t, \tau; K)/\partial S_t$ be the number of shares of the underlying asset to be purchased and $\Delta_0 = O(t, \tau; K) - \Delta_S(t) S_t$ be the residual cash positions.

To examine the hedging performance, we use the following steps. First, on day t we short an option and construct a hedging portfolio by buying $\Delta_S(t)$ shares of

[10] In the case of January 1999 and January 2001 contracts, part of the data is included in our sample because the sample period extends from January 3, 1999, through December 26, 2000.

the underlying asset[11] and investing $\Delta_0(t)$ in a risk-free bond. To compute $\Delta_S(t)$, we use estimated structural parameters from the previous trading day and the current day's asset price. Second, we liquidate the position after the next trading day or the next week. Then we compute the hedging error as the difference between the value of the replicating portfolio and the option price at the time of liquidation:

$$\varepsilon_t = \Delta_S \cdot S_{t+\Delta t} + \Delta_0 e^{r\Delta t} - O(t + \Delta t, \tau - \Delta t; K) \tag{13}$$

Tables 5 and 6 present one-day and one-week ahead hedging errors over alternative moneyness categories, respectively. For one-day ahead errors, the Black–Scholes and the AHBS have better hedging performance for call options. For put options, the TLM is a little better than the other models. As a rule, differences among the models are not so large for the short term. For one-week ahead errors, however, the Black–Scholes and the AHBS are better than the other models. This finding seems somewhat surprising, especially given the TLM's better out-of-sample pricing performance. As discussed by Dumas, Fleming and Whaley (1998) in a different context, a possible explanation is as follows. Although the Black–Scholes option values are systematically incorrect, its errors are stable (or, at least, strongly serially dependent as suits a specification error), unlike the less parsimonious model, the TLM.

6 Conclusion

We have studied the usefulness of implied risk-neutral distributions. To extract the risk-neutral distributions from option prices, we compared the performance of the Black–Scholes model, which assumes a single lognormal distribution, with that of a mixture of two lognormal distributions. We have also reported results for the ad hoc Black–Scholes model with an empirical fit of the volatility smile. Comparisons are based on three metrics: in-sample pricing performance, out-of-sample pricing performance, and hedging performance.

We find that the option pricing formula using the two-lognormal mixture distributions shows the best pricing performance for the short- and long-term forecasting periods. The results are reversed for hedging. Although differences across models are small, we find the Black–Scholes and AHBS models to be better than other models, especially for long-term horizons.

These results point to several conclusions. First, we obtain opposite results for pricing and hedging considerations. Models that perform well for pricing, based on the risk-neutral distributions, may not do so well for hedging, which involves actual distributions. Second, the differences across the models decrease with the horizon, which suggests instability in the distribution parameters. Finally, for hedging purposes, simple models such as the Black–Scholes are superior to more complicated models that may overfit the data.

---

[11] In the case of put options, some shares of the underlying asset are shorted because $\Delta_S(t)$ is negative.

TABLE 5 One-day ahead hedging errors.

		<0.94	0.94–0.97	0.97–1.00	1.00–1.03	1.03–1.06	≥1.06	All
					Panel A: Calls			
MPE	BS	–0.0355	0.0236	0.0265	0.0216	0.0147	–0.0028	0.0019
	AHBS	–0.0293	0.0256	0.0265	0.0217	0.0154	–0.0028	0.0042
	TLM	–0.0345	0.0162	0.0166	0.0179	0.0157	–0.0006	–0.0008
	ATLM	–0.0375	0.0210	0.0216	0.0201	0.0166	–0.0008	0.0003
MAPE	BS	0.2393	0.1455	0.1161	0.0868	0.0777	0.0567	0.1420
	AHBS	0.2427	0.1469	0.1162	0.0870	0.0783	0.0564	0.1433
	TLM	0.2439	0.1497	0.1196	0.0899	0.0776	0.0562	0.1450
	ATLM	0.2481	0.1531	0.1213	0.0890	0.0772	0.0565	0.1469
MAE	BS	0.2880	0.3145	0.3673	0.3939	0.4836	0.5499	0.3739
	AHBS	0.2901	0.3163	0.3677	0.3948	0.4874	0.5478	0.3751
	TLM	0.2890	0.3259	0.3832	0.4084	0.4830	0.5420	0.3797
	ATLM	0.2945	0.3327	0.3871	0.4033	0.4800	0.5439	0.3823
MSE	BS	0.2105	0.1916	0.2571	0.2681	0.4361	0.7050	0.3089
	AHBS	0.2102	0.1928	0.2580	0.2679	0.4484	0.7025	0.3100
	TLM	0.2204	0.2102	0.2879	0.2828	0.4212	0.6864	0.3183
	ATLM	0.2333	0.2209	0.2858	0.2793	0.4186	0.6738	0.3211
					Panel B: Puts			
MPE	BS	–0.0059	0.0050	0.0205	0.0229	0.0280	0.0223	0.0160
	AHBS	–0.0050	0.0056	0.0200	0.0235	0.0308	0.0290	0.0184
	TLM	–0.0061	0.0023	0.0126	0.0128	0.0330	0.0413	0.0186
	ATLM	–0.0069	0.0038	0.0154	0.0179	0.0390	0.0416	0.0208
MAPE	BS	0.0592	0.0802	0.1014	0.1328	0.1827	0.1916	0.1312
	AHBS	0.0594	0.0799	0.1015	0.1340	0.1833	0.1916	0.1315
	TLM	0.0593	0.0806	0.1029	0.1375	0.1795	0.1802	0.1287
	ATLM	0.0599	0.0816	0.1024	0.1358	0.1806	0.1852	0.1301
MAE	BS	0.6517	0.4947	0.4337	0.3664	0.3087	0.2223	0.3944
	AHBS	0.6535	0.4931	0.4342	0.3683	0.3089	0.2208	0.3945
	TLM	0.6509	0.4981	0.4375	0.3756	0.3079	0.2135	0.3942
	ATLM	0.6627	0.5044	0.4358	0.3744	0.3102	0.2182	0.3981
MSE	BS	0.8192	0.4265	0.3246	0.2514	0.1857	0.1147	0.3336
	AHBS	0.8349	0.4274	0.3252	0.2522	0.1850	0.1124	0.3358
	TLM	0.8304	0.4309	0.3274	0.2658	0.1873	0.1085	0.3371
	ATLM	0.8830	0.4487	0.3247	0.2671	0.1865	0.1120	0.3486

This table reports one-day ahead hedging errors for KOSPI 200 Index options with respect to moneyness, S/K. All hedges use only the underlying asset as the hedging instrument. Parameters and spot volatility implied by all options of the previous day are used to establish the current day's hedge portfolio, which is then liquidated on the following day. For each option, its hedging error is the difference between the replicating portfolio value and its market price. Denoting hedging error by ε_n, hedging performance is evaluated using the same error expressions as in Table 2. BS, Black–Scholes model; AHBS, ad hoc Black–Scholes procedure that fits the implied volatility surface; TLM, two-lognormal mixture model; ATLM, adjusted TLM model.

TABLE 6 One-week ahead hedging errors.

		<0.94	0.94–0.97	0.97–1.00	1.00–1.03	1.03–1.06	≥1.06	All
					Panel A: Calls			
MPE	BS	–0.2951	0.1872	0.1934	0.1316	0.0624	–0.0084	–0.0006
	AHBS	–0.2799	0.1952	0.1966	0.1339	0.0631	–0.0084	0.0061
	TLM	–0.3397	0.1482	0.1684	0.1194	0.0567	–0.0063	–0.0263
	ATLM	–0.3461	0.1558	0.1837	0.1217	0.0616	–0.0043	–0.0230
MAPE	BS	0.9081	0.4471	0.2906	0.1890	0.1287	0.0844	0.4364
	AHBS	0.9126	0.4530	0.2907	0.1899	0.1284	0.0850	0.4388
	TLM	0.9623	0.4650	0.2901	0.1834	0.1319	0.0881	0.4558
	ATLM	0.9523	0.4519	0.2904	0.1859	0.1311	0.0872	0.4502
MAE	BS	0.8930	0.7715	0.7663	0.7683	0.7654	0.8242	0.8127
	AHBS	0.8990	0.7791	0.7649	0.7726	0.7642	0.8287	0.8166
	TLM	0.9493	0.8020	0.7638	0.7441	0.7874	0.8594	0.8378
	ATLM	0.9397	0.7876	0.7567	0.7543	0.7830	0.8440	0.8300
MSE	BS	1.8093	0.9674	0.9328	0.9087	0.9968	1.1761	1.2323
	AHBS	1.7970	0.9663	0.9300	0.9322	1.0022	1.1916	1.2341
	TLM	2.2102	1.1262	0.9518	0.8548	1.0159	1.2448	1.3836
	ATLM	2.1727	1.1032	0.9263	0.8605	1.0156	1.1885	1.3563
					Panel B: Puts			
MPE	BS	–0.0313	0.0327	0.1061	0.1780	0.2494	0.2354	0.1260
	AHBS	–0.0292	0.0347	0.1099	0.1826	0.2550	0.2439	0.1306
	TLM	–0.0352	0.0157	0.0846	0.1592	0.2285	0.2478	0.1169
	ATLM	–0.0341	0.0178	0.0929	0.1599	0.2399	0.2586	0.1225
MAPE	BS	0.1158	0.1520	0.2081	0.3114	0.5048	0.5623	0.3141
	AHBS	0.1166	0.1527	0.2063	0.3070	0.5017	0.5693	0.3146
	TLM	0.1219	0.1638	0.2183	0.3124	0.4909	0.5636	0.3176
	ATLM	0.1204	0.1611	0.2171	0.3127	0.5092	0.5658	0.3193
MAE	BS	1.2736	0.9153	0.8330	0.7536	0.7251	0.5600	0.8470
	AHBS	1.2805	0.9191	0.8240	0.7425	0.7188	0.5644	0.8460
	TLM	1.3370	0.9910	0.8693	0.7570	0.7125	0.5580	0.8741
	ATLM	1.3167	0.9745	0.8612	0.7565	0.7364	0.5638	0.8713
MSE	BS	3.2940	1.4483	1.1820	0.8881	0.8076	0.5072	1.4027
	AHBS	3.2854	1.4419	1.1610	0.8650	0.7905	0.5154	1.3929
	TLM	3.8056	1.7192	1.2535	0.9071	0.8045	0.5154	1.5575
	ATLM	3.6695	1.6578	1.2331	0.9037	0.8467	0.5333	1.5292

This table reports one-week ahead hedging errors for KOSPI 200 Index options with respect to moneyness, S/K. All hedges use only the underlying asset as the hedging instrument. Parameters and spot volatility implied by all options of the previous day are used to establish the current day's hedge portfolio, which is then liquidated on the following week. For each option, its hedging error is the difference between the replicating portfolio value and its market price. Denoting hedging error by ε_n, hedging performance is evaluated using the same error expressions as in Table 2. BS, Black–Scholes model; AHBS, ad hoc Black–Scholes procedure that fits the implied volatility surface; TLM, two-lognormal mixture model; ATLM, adjusted TLM model.

REFERENCES

Aït-Sahalia, Y., and Lo, A. (1998). Nonparametric estimation of state–price densities implicit in financial asset prices. *Journal of Finance* **53**, 499–547.

Anagnou, I., Bedendo, M., Hodges, S., and Tompkins, R. (2002). The relation between implied and realized probability density functions. Working paper, University of Warwick.

Aparicio, S. D., and Hodges, S. (1998). Implied risk-neutral distribution: a comparison of estimation methods. Working paper, University of Warwick.

Bahra, B. (1997). Implied risk-neutral probability density functions from option prices: theory and application. Working paper, Bank of England.

Bakshi, G. S., Cao, C., and Chen, Z. W. (1997). Empirical performance of alternative option pricing models. *Journal of Finance* **52**, 2003–49.

Black, F., and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy* **81**, 634–59.

Bliss, R. R., and Panigirtzoglou, N. (2000). Testing the stability of implied probability density functions. Working paper, Bank of England.

Bliss, R. R., and Panigirtzoglou, N. (2002). Option-implied risk aversion estimates: robustness and patterns. Working paper, Federal Reserve Bank of Chicago.

Breeden, D. T., and Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business* **51**, 621–51.

Campa, J. M., Chang, P. H. K., and Refalo, J. F. (1999). An option-based analysis of emerging market exchange rate expectations: Brazil's real plan, 1994–1997. Working paper, University of Southern California.

Dumas, B., Fleming, J., and Whaley, R. (1998). Implied volatility functions: empirical tests. *Journal of Finance* **53**, 2059–106.

Gemmill, G., and Saflekos, A. (2000). How useful are implied distributions? Evidence from stock index options. *Journal of Derivatives* **22**, 83–98.

Jackwerth, J. C., and Rubinstein, M. (1996). Recovering probability distributions from option prices. *Journal of Finance* **51**, 1611–31.

Kirgiz, Ì. (2001). An empirical comparison of alternative stochastic volatility option pricing models. Working paper, University of Maryland.

Malz, A. M. (1997). Estimating the probability distribution of future exchange rates from option prices. *Journal of Derivatives* **5**, 20–36.

Melick, W. R., and Thomas, C. P. (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* **32**, 91–115.

Shimko, D. (1991). Beyond implied volatility: probability distribution and hedge ratios implied by option prices. Working paper, University of Southern California.

Shimko, D. (1993). Bounds of probability. *Risk* **6**, 34–7.

Shiratsuka, S. (2001). Information content of implied probability distributions: empirical studies on Japanese stock price index options. Working paper, Bank of Japan.

Weinberg, S. A. (2001). Interpreting the volatility smile: an examination of the information content of option prices. Working paper, Board of Governors of the Federal Reserve System.