MAXIMUM ENTROPY IN OPTION PRICING: A CONVEX-SPLINE SMOOTHING METHOD

WEIYU GUO

Applying the principle of maximum entropy (PME) to infer an implied probability density from option prices is appealing from a theoretical standpoint because the resulting density will be the least prejudiced estimate, as "it will be maximally noncommittal with respect to missing or unknown information."[1] Buchen and Kelly (1996) showed that, with a set of well-spread simulated exact-option prices, the maximum-entropy distribution (MED) approximates a risk-neutral distribution to a high degree of accuracy. However, when random noise is added to the simulated option prices, the MED poorly fits the exact distribution. Motivated by the characteristic that a call price is a convex function of the option's strike price, this study suggests a simple convex-spline procedure to reduce the impact of noise on observed option prices before inferring the MED. Numerical examples show that the convex-spline smoothing method yields satisfactory empirical results that are consistent with prior studies. © 2001 John Wiley & Sons, Inc. Jrl Fut Mark 21:819–832, 2001

[1]Buchen & Kelly (1996), p. 143.

The author would like to express her deepest appreciation to Charles Corrado, Stephen Ferris, Les Gulko, Vivek Mande, Thomas Miller, and Mark Wohar for advice and constructive comments. She would also like to acknowledge the valuable comments and suggestions from the reviewers and thank Jackie Lynch for her assistance in formatting the final version of this paper.

For correspondence, Weiyu Guo, Assistant Professor of Finance, Department of Finance, Banking, and Law, College of Business Administration, University of Nebraska at Omaha, Omaha, Nebraska 68182-0048.

Received September 1999; Accepted January 2001

■ *Weiyu Guo is an Assistant Professor of Finance in the Department of Finance, Banking, and Law at the University of Nebraska at Omaha, in Omaha, Nebraska.*

The Journal of Futures Markets, Vol. 21, No. 9, 819–832 (2001)

INTRODUCTION

The empirical deficiencies of the Black–Scholes option-pricing model are well known and documented extensively by Black (1975), MacBeth and Merville (1979), Emanuel and MacBeth (1982), and Rubinstein (1985, 1994). In particular, Rubinstein (1994) documented two interesting and related phenomena. First, before the October 1987 market crash, the Black–Scholes model appeared to be an effective tool for pricing options on the S&P 500 index. Second, since the October 1987 crash, Rubinstein found that the Black–Scholes option-pricing formula systematically misprices deep in-the-money and out-of-the-money options. To practitioners, these mispricing patterns are known as "volatility smiles." Volatility smiles generally are believed to result from violations of one or more assumptions underlying the Black–Scholes option-pricing model. In particular, the Black–Scholes model assumes that the natural logarithm of the underlying stock price follows a normal distribution, with a variance that increases exactly linearly with time. In contrast to these assumptions, Engle (1982) and Bollerslev (1986) found that stock-return variances are volatile. In fact, stock-return variances are stochastic and generally correlated negatively with stock price levels. Consequently, Rubinstein (1994) demonstrated that implied probability distributions obtained from stock-index options are highly negatively skewed, kurtotic, and quite unlike a normal or lognormal distribution. Hull and White (1987) and Heston (1993) showed that the departure of stock-return distributions from a normal specification resulted in systematic pricing errors from the Black–Scholes option-pricing formula.

Breeden and Litzenberger (1978) showed that the risk-neutral density of an asset was proportional to the second derivative of a related call-option formula with respect to its strike price. This is the theoretical basis for the "inverse" process: starting from a series of market option prices, one works backward to identify a risk-neutral distribution for the underlying asset. A good solution to the inverse problem produces a distribution that accurately prices options at strikes and maturities not included in the data set, that is, it is a good implicit interpolator of option prices as a function of moneyness and maturity. Furthermore, the solution should help enlighten investigators about the relationship between the markets for the option and its underlying asset, for example, the difference between the actual and risk-neutral volatilities for the underlying assets. Empirically, however, the procedure is difficult, as strike prices are discrete and limited to a range near the current asset price. Breeden and Litzenberger relied on interpolation to recover a whole distribution. As might be expected, the results were sensitive to the interpolation scheme.

Of the existing solutions that have been explored, maximum-entropy methods extract an asset's probability distribution by maximizing all unknown information, subject to the constraint of being consistent with all "known" information. This is appealing from a theoretical standpoint because the resulting density provides the least prejudiced estimate. Stutzer (1996) applied the maximum-entropy concept in canonical valuation to estimate risk-neutral probabilities for underlying assets based on historical returns. Jackwerth and Rubinstein (1996) examined the "maximum entropy function" as one of the objective functions in an optimization method to recover the implied probability distribution from observed S&P 500 index options.

With a set of well-spread simulated exact-option prices, Buchen and Kelly (1996) demonstrated that the maximum entropy distribution (MED) approximated a risk-neutral distribution to a high degree of accuracy. However, when random noise was added to the simulated option prices, the MED fit poorly the exact distribution. Unfortunately, under real market conditions, observed option prices often are noisy. Motivated by the theoretical non-arbitrage restriction that a call price is a convex function of the option's strike price, this study proposes a simple smoothing algorithm to reduce the adverse impact of noise to accommodate the proper application of PME. The proposed procedure fits a smoothed cubic-spline to the data, subject to the constraint that the fitted line is convex. It then infers the MED from the smoothed prices. The proposed procedure has two objectives. First, and most important, is the removal of non-convexities from the data. Second is the smoothing of the data. Numerical examples show that the smoothing method yields satisfactory empirical results.

This article contributes to the current literature by extending the maximum-entropy proposal of Buchen and Kelly (1996) to cope better with possible measurement errors in the chosen subset of market-observed option prices through smoothing. Smoothing was used by Jackwerth and Rubinstein (1996); however, the specific method this article proposes has not been suggested previously. Jackwerth and Rubinstein (1996) integrated smoothing into their estimation procedures to generate a smoothed distribution through a complicated procedure that included carefully constructing an implied volatility smile, executing the optimization, and monitoring the convergence. The proposed procedure in this study seeks to reduce noise before recovering the MED. Furthermore, varying the smoothness parameter in the algorithm, which involves a subjective decision, can alter the degree of smoothness in Jackwerth and Rubinstein (1996). The smoothing procedure examined in this study is not subjective, but based on a criterion specified in the original IMSL algorithm, which includes the constraint of removing non-convexities from the data. Although the IMSL algorithm could be modified to allow subjective smoothing, the major purpose of this study is to suggest a smoothing procedure that does not require subjective judgment.

THE MAXIMUM ENTROPY VALUATION

Option strike prices are discrete and limited to a range near the current asset price; consequently, limited observed-option prices provide only partial information with which to recover implied probability densities representing the whole price domain. In a situation like this, the data are considered "ill-posed" for the problem and the application of the principle of maximum entropy is a good solution.

"Entropy" is a classical thermodynamic concept first utilized by Shannon (1948) as a measure of uncertainty. Jaynes (1957) applied the concept in estimating the probability density for the occurrence of a future event. According to the PME, in estimating the unknown probability distribution based on partial information, one chooses the distribution for which the data is sufficient to determine the probability so that the entropy (the unknown) is maximized. Consequently, the distribution recovered from this principle maximizes all unknown information while being consistent with all "known" information. A distribution recovered following this principle is called a MED.

Consistent with the notation of Buchen and Kelly (1996), the application of maximum entropy for option pricing can be presented as follows. Let X be a continuous random variable that represents the price of an asset at some fixed time, T, in the future. To estimate a risk-neutral probability density $p(x)$ of X on $0 < x < +\infty$, based only on price information from options written on X, which expire at time T, the entropy of the distribution $p(x)$ is defined as:

$$S(p) = -\int_0^{+\infty} p(x) \log p(x)\, dx$$

To maximize all missing information, the objective function is:

$$\max_p S(p) = -\int_0^{+\infty} p(x) \log p(x)\, dx$$

The objective function is subject to the normalization–additivity constraint:

$$\int_0^{+\infty} p(x)\, dx = 1$$

More than one frequency function may satisfy the above objective function. To choose the most likely distribution function, the one generated in the greatest number of ways consistent with the data constraints is preferred. In a risk-neutral world, observed option prices are modeled as a simple discounted pay-off function at expiration in the form of:

$$c_i(X) = D(T) \int_0^{+\infty} \max\{x - k_i, 0\} p(x)\, dx$$

where $D(T)$ is the discount factor corresponding to maturity T; x is underlying asset price; k_i is strike price and $\max\{x - k_i, 0\}$ is the ith option payoff function at expiration.

Set up the Lagrangian function,

$$L = -\int_0^{+\infty} p(x) \log p(x)\, dx + \lambda_0 \left(1 - \int_0^{+\infty} p(x)\, dx\right) + \sum_{i=1}^{N} \lambda_i D(T) \int_0^{+\infty} \max\{x - k_i, 0\} p(x)\, dx$$

where $\lambda_0, \lambda_1, \ldots \lambda_n$ are Lagrangian parameters. The first order condition of the Lagrangian function leads to the solution of $p(x)$ as follows:

$$p(x) = \frac{\exp\left(\sum_{i=1}^{N} \lambda_i \max\{x - k_i, 0\}\right)}{\int_0^{+\infty} \exp\left(\sum_{i=1}^{N} \lambda_i \max\{x - k_i, 0\}\right) dx}$$

is always positive due to the exponential form. Knowing the probability density function of X, an entropy call price $c_{ent-i}(X)$ is determined as:

$$c_{ent-i}(X) = D(T) \int_0^{+\infty} \max\{x - k_i, 0\} \frac{\exp\left(\sum_{i=1}^{N} \lambda_i \max\{x - k_i, 0\}\right)}{\int_0^{+\infty} \exp\left(\sum_{i=1}^{N} \lambda_i \max\{x - k_i, 0\}\right) dx} dx$$

where N is the number of strikes (which includes a notional strike price of zero).

It is important that the current asset price, which corresponds to a strike price of zero, should be included in the calculation. If the current asset price is not included in the data constraints, the MED is uniform below the lowest strike because, on a finite interval without prior information, all future outcomes are equally likely. Current asset price information serves to tie down the distribution in this range.

SMOOTHING MODIFICATION

When there is no noise in the simulated option prices, Buchen and Kelly (1996) demonstrated that, with a set of well-spread strikes over the full asset price domain, the MED is able to fit accurately a known density. However, when random noise was added to the simulated option prices, the maximum entropy distribution fit poorly the known distribution. This was because when there was noise in the data, data constraints became inconsistent. The maximum entropy approach may run into numerical difficulties, such as a solution becoming trapped in a local maximum, which can be far from a global maximum.

Unfortunately, in practice, noise is inevitable. Noise may come from (1) data-processing errors; (2) price discreteness; (3) simple supply and demand for particular contracts causing prices to depart from their "true" values; and (4) nonsynchronous trading, situations where option prices across strikes may not be quoted simultaneously, or option prices quoted may not correspond exactly to the underlying index.

To reduce the impact of noise, natural approaches include fitting prices constrained to lie within the bid-ask range, as done in Rubinstein (1994). However, for the Buchen and Kelly procedure described above, observed option prices were used as data constraints, which could result in the inability to set up the Lagrange function, thereby making the model intractable. Consequently, this study proposes smoothing the data before applying the entropy procedure. Motivated by the theoretical non-arbitrage restriction that a call price is a convex function of the option's strike price, the proposed procedure fits a smoothed cubic-spline to the data, subject to the constraint that the fitted line is convex. It then infers the maximum entropy distribution from the smoothed prices. The smoothing algorithm is explained below in more detail.

A cubic-spline is a commonly used smoothing method in mathematics. This is an interpolation method that produces a third-order polynomial between each pair of data points. If every two consecutive data points set an interval, the first derivative of the third-order polynomial may not be continuous across the boundary between two intervals. The idea of a cubic-spline is to require this continuity. If a connecting point between two third-order polynomials is called a "knot," then typically the number of "knots" in a cubic-spline coincide with the number of data points.

However, a simple cubic-spline only makes the connections at the "knot" smooth. To smooth the entire spline, a smoothed cubic-spline appears to be a better idea. A smoothed cubic-spline applies a smoothness criterion to the interpolating function and smoothes out the spline function. In this case, an algorithm generates new "knots," and a smoothed spline typically has fewer "knots" than a simple cubic-spline. In order to apply properly the smoothed cubic-spline method, finding the right smoothness criterion is important. A convexity constraint makes economic sense as a non-arbitrage restriction to reduce noise.

When noise is embedded in option prices, the convexity constraint is violated, and the violation can be detected by the cost of forming a long butterfly spread strategy, a strategy of buying two mid-strike calls and selling one of each high- and low-strike calls. Since the payoff to this strategy always is non-negative, the cost of establishing the strategy always should be positive as well. Therefore, a negative long butterfly cost indicates the existence of noise in the observed option prices. It also is worth pointing out that when no negative butterfly cost is identified, it does not necessarily suggest that there is no noise in the data. It may suggest that the degree of noise is not as significant. In this case, convex-spline smoothing may still help to reduce the impact of noise in the data.

The convex-spline algorithm is easy to implement. A routine for cubic-spline smoothing called "CONF" is available from the IMSL Fortran Library. The routine CONF produces a constrained, weighted least-squares fit to data from a spline subspace. The specific constraint required here is that the fitted spline be convex.

To obtain the probability density, $p(x)$, values for the Lagrangian multipliers, λ_i are needed. These Lagrangian multipliers are estimated by numerical procedures. Buchen and Kelly use a variant of a multidimensional Newton–Raphson procedure suggested by Agmon, Alhassid, and Levine (1981).[2] Here, a similar least-squares approach is adopted. Specifically, when a smoothing method is involved, λ_i is estimated by minimizing the following sum of squares with respect to λ_i:

$$\min_{\lambda_i} \sum_{i=1}^{N} [c_{fitted,\, i} - c_{entropy,\, i}(\lambda_i)]^2$$

[2]Buchen & Kelly (1996), p. 150.

where N is number of price constraints across available strikes; $c_{fitted,\, i}$ is a fitted price generated by observed market prices; and $c_{entropy,\, i}$ is a theoretic entropy call price. The values of λ_i obtained from the above procedure minimize the sum of squared errors between the entropy call prices and the fitted prices. If no smoothing method is used, λ_i is estimated by $\min_{\lambda_i} \Sigma_{i=1}^{N} [c_{observed,\, i} - c_{entropy,\, i}(\lambda_i)]^2$, where $c_{fitted,\, i}$ is replaced by an observed call option price $c_{observed,\, i}$.

NUMERICAL EXAMPLES

Initial Test with Rubinstein's S&P 500 Index Option Prices

The S&P 500 index data in Rubinstein (1994) are used for an initial example. The data represent S&P 500 index options maturing in June 1990, with prices quoted at 11:00 am on January 2, 1990. The SPX options have 164 days remaining to maturity.[3] Figure 1 shows a jagged implied-probability density resulting from the raw Buchen and Kelly procedure. Figure 2 shows that, after convex-spline smoothing, the probability density appears much smoother. The density is highly skewed, kurtotic, and slightly bimodal, which is consistent with the results reported in Rubinstein (1994), yet still shows slight improvement with less bimodality at the lower tail.

It might be interesting to compare results with those from the maximum-entropy criterion in Jackwerth and Rubinstein (1996). However, Jackwerth and Rubinstein (1996) did not reveal the exact degree of smoothing applied. Since the degree of smoothing used is subjective across samples, it could vary between too little or too much, and therefore introduce "noise" in the estimation procedure. Consequently, it is impossible to replicate their results. A rough comparison seems to reveal that the probability density in Jackwerth and Rubinstein (1996) is smoother. A possible explanation for this is that the smoothing method this article suggests is less forceful.

Table I presents the bid–ask averages of the Rubinstein data, the costs to establish a long butterfly spread strategy, fitted option prices after convex-spline smoothing, and the corresponding entropy prices

[3]According to Rubinstein (1994, p. 784), the 13 puts and 15 calls used to construct these bid and ask prices did not trade simultaneously at 11:00 am. To surmount this problem, elaborate procedures were followed to adjust nonsimultaneously observed option prices to their probable 11:00 am levels, had they actually been quoted at that time.

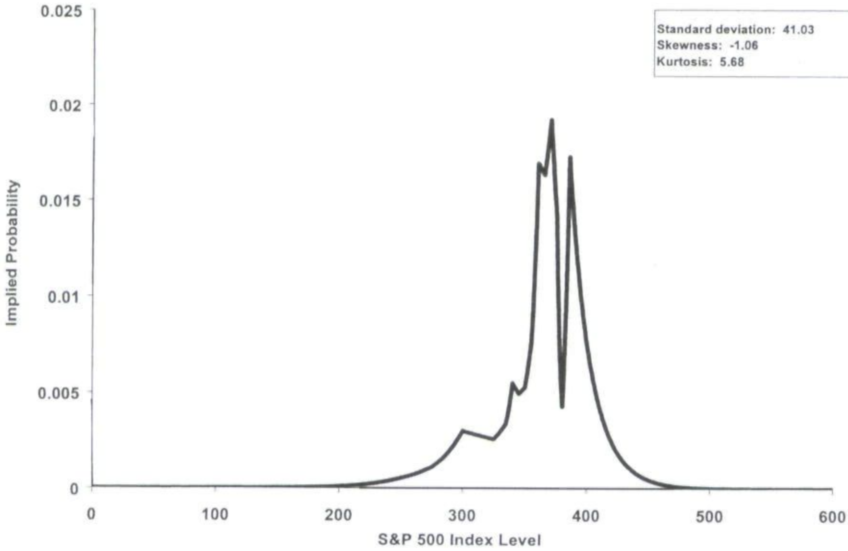

FIGURE 1
Risk-neutral 164-day probabilities of S&P 500 index options, no smoothing, for January 2, 1990, 11 am.

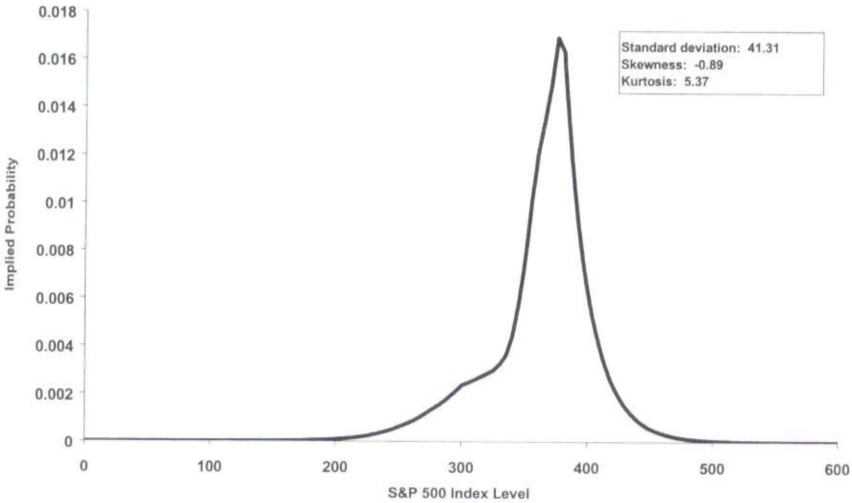

FIGURE 2
Risk-neutral 164-day probabilities of S&P 500 index options, convex-spline smoothing, for January 2, 1990, 11 am.

generated by the fitted prices. Entropy prices generated from the raw Buchen and Kelly procedure and Rubinstein's optimization values also are reported for comparison. Table I shows two negative long butterfly costs at the 345 and 380 strike levels. These negative costs indicate that noise is imbedded in the data as the convexity of option price property is

TABLE I
Rubinstein's S&P 500 Index Options*

Strike	*Bid–Ask Average*	*Butterfly Cost*	*Fitted Price*	*Entropy w/ Smoothing*	*Entropy w/o Smoothing*	*Optimization Value*
0	349.21	–	349.22	349.22	349.21	349.16
250	109.59	–	109.52	109.54	109.60	109.71
275	86.19	–	86.19	86.16	86.18	86.70
300	63.52	–	63.59	63.58	63.51	63.90
325	42.38	–	42.35	42.35	42.36	42.09
330	38.29	0.01	38.31	38.31	38.33	37.99
335	34.33	0.04	34.34	34.34	34.37	34.01
340	30.42	0.29	30.45	30.47	30.49	30.16
345	26.79	−0.03	26.67	26.70	26.74	26.45
350	23.14	0.12	23.04	23.07	23.11	22.90
355	19.60	0.20	19.61	19.62	19.61	19.55
360	16.26	0.52	16.41	16.41	16.31	16.45
365	13.44	0.22	13.49	13.50	13.38	13.65
370	10.84	0.51	10.90	10.90	10.85	11.15
375	8.75	0.38	8.67	8.66	8.75	8.96
380	7.03	−0.06	6.84	6.82	6.96	7.07
385	5.26	–	5.39	5.37	5.33	5.45

Description: Bid–ask averages and optimization values are from Rubinstein (1994).

*Data are from S&P 500 index options maturing in June 1990 across all strikes at 11:00 am on January 2, 1990. Options have 164 days to maturity. Butterfly cost is the cost to establish a long butterfly option strategy (buy two mid-strike calls and sell one of each high- and low-strike calls). Fitted prices are generated from observed option prices by convex-spline smoothing. Entropy w/ smoothing refers to entropy call option prices generated when convex-spline smoothing is used. Entropy w/o smoothing refers to entropy call option prices generated when no smoothing method is involved. Optimization values are the theoretical option prices generated by Rubinstein using his optimization method.

violated. This justifies the need for data smoothing to improve the performance of entropy valuation.

Dow Jones Industrial Average Options

The second example is for options on the Dow Jones Industrial Average (DJIA). Similar to options on the S&P 500 index, options on the DJIA are European style, cash-settled, do not have a "wildcard" feature, and have futures and futures options traded on the index. Unlike options on the S&P 500 index, where the underlying index is value-weighted, the underlying index for options on the DJIA is price-weighted. While the current literature has documented the shape of implied probability densities from S&P 500 index options, no prior study has looked into the implied density that might be obtained from options on the DJIA.

DJIA option price quotes and concurrent Dow Jones index levels were downloaded from the Chicago Board Option Exchange web site

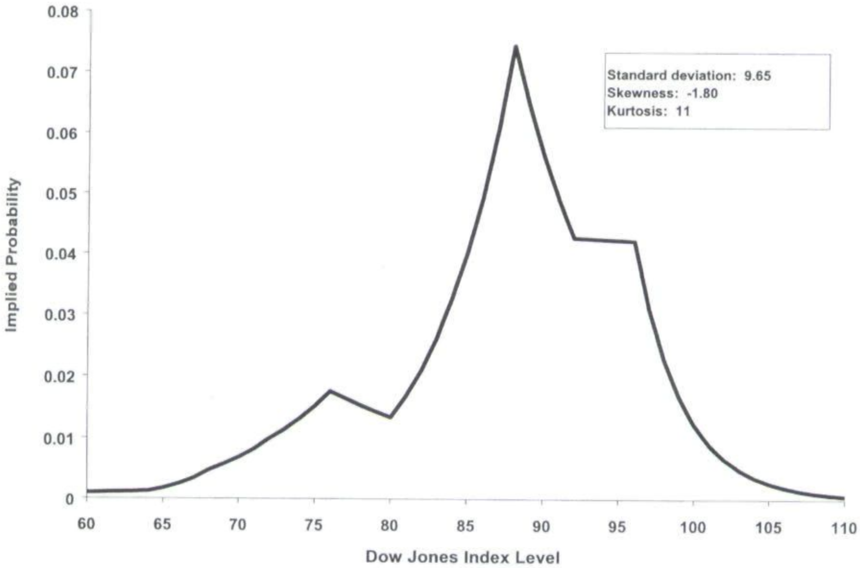

FIGURE 3
Risk-neutral 99-day probabilities of Dow Jones Industrial Average options, no smoothing, for January 12, 1998, 11:02 am.

at approximately 11:22 am on March 12, 1998.[4] The options mature in June 1999 and have 99 days remaining to maturity. The dividend adjusted index level is calculated by $S - (\Sigma_{i=1}^{j} D_i e^{-rT_i}/divisor)$, where S is a current index level; D_i is a dividend paid by a Dow component between March 12, 1998 and the option expiration date, June 19, 1998; T_i is the time to the ex-dividend date; r is a Treasury-bill rate with a matching maturity; and *divisor* is the concurrent divisor for the DJIA.

Figure 3 shows that, without smoothing, the implied probability density from the DJIA option prices is jagged. However, it is smoother than that obtained from the S&P 500 index options. There are two possible explanations for this: there is less noise in the Dow Jones option data, as no negative butterfly cost is reported in Table II. Also, there are fewer strikes in the Dow sample, so the recovered probability density does not have to go through as many inconsistent data points. Figure 4 shows that the recovered density is improved after the smoothing. It is negatively skewed and kurtotic, which is consistent with what is generally expected.

[4]Since there is a 20-minute delay for quotes posted on the web site, the data corresponds to real-time quotes at approximately 11:02 am of the day.

TABLE II
Dow Jones Industrial Average Options*

Strike	*Bid–Ask Average*	*Butterfly Cost*	*Fitted Price*	*Entropy w/ Smoothing*	*Entropy w/o Smoothing*
0	86.09	–	86.09	86.10	86.09
64	23.25	–	23.23	23.24	23.25
68	19.38	0.13	19.40	19.40	19.40
72	15.63	0.13	15.63	15.63	15.62
76	12.00	0.25	12.00	12.00	12.00
80	8.63	0.25	8.59	8.59	8.62
84	5.50	0.56	5.52	5.52	5.51
88	2.94	0.94	2.98	2.98	2.94
92	1.31	0.75	1.27	1.26	1.31
96	0.44	–	0.45	0.45	0.44

*Option prices are from options on the Dow Jones Industrial Average (DJX) at 11:02 am on March 12, 1998. Options have 99 days remaining to maturity. The Treasury-bill rate with a matching maturity is 5.07%. The corresponding dividend adjusted index level is 86.09. Butterfly cost is the cost to establish a long butterfly spread option strategy (buy two mid-strike calls and sell one of each high- and low-strike calls). Fitted prices are generated from observed option prices by convex-spline smoothing. Entropy w/ smoothing refers to entropy call option prices generated when convex-spline smoothing is used. Entropy w/o smoothing refers to entropy call option prices generated when no smoothing method is involved.

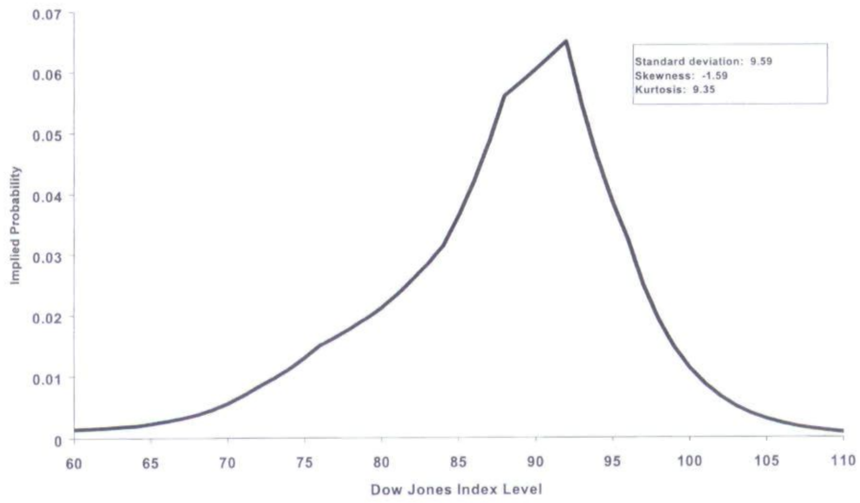

FIGURE 4
Risk-neutral 99-day probabilities of Dow Jones Industrial Average options, convex-spline smoothing, for January 12, 1998, 11:02 am.

SUMMARY

Maximum entropy methods extract an asset's probability distribution by maximizing all unknown information, subject to the constraint of being consistent with all "known" information. This is appealing from a theoretical standpoint because the resulting density provides the least prejudiced estimate. Buchen and Kelly (1996) showed that, with a set of well-spread simulated exact option prices, the MED approximated a risk-neutral distribution to a high degree of accuracy. However, when random noise was added to the simulated option prices, the MED fit poorly the exact distribution. Motivated by the theoretical constraint that a call price is a convex function of its strike price, this study has suggested a simple convex-spline smoothing algorithm to reduce the impact of noise before inferring the MED. Applying the smoothing procedure to Rubinstein's adjusted SPX option prices demonstrates that the implied probability density is much improved over that generated from the raw Buchen and Kelly procedure and is consistent with that in Rubinstein (1994). Similarly, with convex-spline smoothing, the implied density recovered from options on the DJIA also is smoother. In general, the resulting density is negatively skewed, kurtotic, and consistent with what is generally expected from index options.

BIBLIOGRAPHY

Agmon, N., Alhassid, Y., & Levine, R. D. (1981). An algorithm for determining the Lagrange parameters in the maximal entropy foundation. In R. D. Levine & M. Tribus (Eds.), The maximum entropy formalism (pp. 207–209), Cambridge, MA: MIT Press.

Black, F. (1975). Fact and fantasy in the use of options. Financial Analysts Journal, 31, 36–72.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31, 307–327.

Breeden, D., & Litzenberger, R. (1978). Prices of state-contingent claims implicit in options prices. Journal of Business, 51, 621–651.

Buchen, P., & Kelly, M. (1996). The maximum entropy distribution of an asset inferred from option prices. Journal of Financial and Quantitative Analysis, 31, 143–159.

Emanuel, D., & MacBeth, J. (1982). Further results on the constant elasticity of variance call option pricing model. Journal of Financial and Quantitative Analysis, 18, 533–554.

Engle, R. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. Econometrica, 50, 987–1008.

Heston, S. (1993). A closed form solution for options with stochastic volatility with applications to bond and currency options. Review of Financial Studies, 6, 327–344.

Hull, J., & White, A. 1987. The pricing of options on assets with stochastic volatilities. Journal of Finance, 42, 281–300.

Jackwerth, J., & Rubinstein, M. (1996). Recovering probability distributions from option prices. Journal of Finance, 51, 1611–1631.

Jaynes, E. T. (1957). Information theory and statistical mechanics. Physics Review, 106, 620–630.

MacBeth, J., & Merville, L. (1979). An empirical examination of the Black–Scholes call option pricing model. Journal of Finance, 34, 1173–1186.

Rubinstein, M. (1985). Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active CBOE option classes from August 23, 1976 through August 31, 1978. Journal of Finance, 40, 455–480.

Rubinstein, M. (1994). Implied binomial trees. Journal of Finance, 69, 771–818.

Shannon, C. E. (1948). The mathematical theory of communication. Bell Systems Technical Journal, 27, 379–423.

Stutzer, M. (1996). A simple nonparametric approach to derivative security valuation. Journal of Finance, 51, 1633–1651.