ESTIMATING IMPLIED PDFs FROM AMERICAN OPTIONS ON FUTURES: A NEW SEMIPARAMETRIC APPROACH

DIMITRIS FLAMOURIS
DANIEL GIAMOURIDIS*

This article develops a new methodology for estimating implied probability density functions for futures prices from American options. The restricting Black–Scholes assumption of a lognormal distribution for the underlying asset is relaxed with the use of the more flexible distributional form of an Edgeworth series expansion around a lognormal distribution. The model is applied to the crude oil market. The results provide strong evidence that the market consensus can be accurately reflected in the risk-neutral densities recovered from observed option prices. The recovered distributions are tested and found to differ significantly from a single lognormal distribution. In addition, the recovered distributions are more robust than those

The authors thank Gordon Gemmill (City University Business School), Roberto Violi (Bank of Italy), two anonymous referees, and participants at the Chicago Board of Trade 13th Annual European Futures Symposium in Glasgow for their helpful discussions. D. Giamouridis acknowledges the support of the "Alexander S. Onassis" Foundation and the Astron Maritime Company SA.
*Correspondence author, Department of Shipping, Trade and Finance, City University Business School, Frobisher Crescent, Barbican Centre, London EC2Y 8HB, U.K.; e-mail: d.giamouridis@city. ac.uk

Received October 2000; Accepted June 2001

- *Dimitris Flamouris is a Risk Analyst at Bear Stearns.*
- *Daniel Giamouridis is a Research Student in the Department of Shipping, Trade and Finance at City University Business School in London, United Kingdom.*

The Journal of Futures Markets, Vol. 22, No. 1, 1–30 (2002)
© 2002 by John Wiley & Sons, Inc.

recovered with a model, which assumes a mixture of two lognormal distributions. © 2002 John Wiley & Sons, Inc. Jrl Fut Mark 22:1–30, 2002

INTRODUCTION

The estimation of the risk-neutral probability distribution implied in option prices has been the objective of numerous research papers in academic literature. Implied distributions are, however, being used extensively by traders and policy makers to assess market beliefs on the future movements of the underlying asset. Having a valid tool for the extraction of the information implied in option prices is, therefore, an issue of vital importance.

This article introduces a new methodology for recovering implied distributions from American option prices. The restricting assumption of normally distributed returns is relaxed, and a more realistic form for the distribution of the underlying asset is proposed. This is the first time that the Edgeworth series expansion (ESE) method has been used in an American option framework.

The method can generate distributions exhibiting bimodal patterns, a typical characteristic of distributions belonging to the family of two-component mixtures of lognormal distributions. The proposed model, however, is more parsimonious, requiring the estimation of four parameters only, the respective number being five for a mixture of two lognormals model. This advantage can prove significant in markets where the number of available actively traded strike prices is relatively small. Hence, possible situations of overparameterization of the data are avoided.

The information conveyed by the implied probability distributions, recovered by the proposed method, proved to be consistent with the market commentary at the time. The implied distributions were shown to be able to capture the general market sentiment and also incorporate isolated events causing a significant impact on the market.

In addition, the model's explanatory ability was tested against that of the corresponding nested single lognormal. It offered a statistically significant better fit to observed option prices and can, therefore, be considered a superior means of extracting information implied in option prices.

Finally, the robustness of the ESE model was investigated, and the model was found to be capable of estimating on average probability density functions (PDFs) whose summary statistics converge to the original solution; it is also capable of estimating statistics with a relatively low dispersion with respect to an alternative model that assumes a mixture of two lognormals distribution.

The work of Melick and Thomas (1997) is used as a starting point. The use of a mixture of three lognormal distributions involves the estimation of a large number of parameters (eight total) and so is difficult to implement in situations where the number of available traded strike prices is not large. Furthermore, the estimation of a large number of parameters is associated with the risk of overparameterization and overfitting of the data. The ESE of a known distribution can be used alternatively to represent the functional form of the true risk-neutral distribution of asset returns.

The ESE was first introduced in the literature of option valuation by Jarrow and Rudd (1982). Jarrow and Rudd (1983) performed empirical tests to examine the validity of the methodology when used to price individual stock options. Corrado and Su (1996, 1997) also tested the model, using options data on the Standard & Poor's 500 Index. They showed that a model of this type outperforms the simple Black–Scholes pricing model by providing a significantly better fit to observed option prices and by performing better when used to predict option prices one period ahead. Backus, Foresi, Li, and Wu (1997) used a similar expansion scheme (the Gram–Charlier series expansion) to deal with the excess kurtosis present in the Foreign Exchange (FX) market. They also identified the potential of such models to deal with nonnormally distributed security returns.

Jarrow and Rudd (1983), Corrado and Su (1997), and Backus et al. (1997) used expansion schemes in the context of option pricing but not in the context of recovering risk-neutral densities implied by option prices. In the authors' opinion, the two processes (recovering the parameters of the model and, therefore, the implied PDFs and using the model to price options with the recovered parameters) cannot be separated. Separating the two processes might give rise to negative PDFs and PDFs that do not integrate to unity, which would lead to option pricing under distributions with absolutely no physical meaning. Jondeau and Rockinger (2000) used the ESE method to recover the PDF of FX over-the-counter options in a study in which various methods of recovering implied distributions were examined and compared. To the authors' knowledge, this was the first time that the method was used to recover PDFs from European-style options. The model captured the phenomenon of large skewness that is present, to a certain extent, in the FX markets and had a good overall performance compared with the alternative methodologies. The speed of the algorithm was also a considerable advantage.

The rest of the article is organized as follows. In the second section, the proposed methodology for recovering the PDF from American options is introduced along with the theoretical considerations for the development of the model. The third section contains the description of the data set used to test the validity of the model. The fourth section discusses the results of the analysis, and the fifth section concludes the article.

METHOD FOR RECOVERING IMPLIED PDFs FROM AMERICAN OPTION PRICES

In this section, a new method for recovering implied distributions from American option prices is presented. The first subsection presents bounds for American option prices in terms of the PDF of the underlying asset. The proposed functional form of the PDF is demonstrated in the second subsection, along with the formulae of option prices estimated with that distribution. The third subsection presents the estimation procedure and, finally, the fourth subsection discusses the constraints that need to be imposed to recover meaningful PDFs.

American Option Prices

It is straightforward to obtain the expected value of an option when early exercise is not optimal. In situations where riskless arbitrage is not possible, Cox and Ross (1976) suggested that the value of the option can be inferred from its expected value at maturity. Hence, the knowledge of the terminal distribution would be an adequate means for pricing options, and, vice versa, the knowledge of option prices would be sufficient to recover the implied PDF.

The early exercise premium that is embedded in American options does not allow the application of such techniques. In general, the American option value depends on the entire stochastic process of the underlying asset, thus making difficult the recovery of implied PDFs.

To deal with this deficit, the value of the option is expressed as a combination of an upper and a lower bound, as by Chaudhury and Wei (1994) and Melick and Thomas (1997), who restricted the value of the option within a very tight range. The use of bounds is a rather limiting methodology because it can only be applied to assets that follow martingale processes. Available evidence[1] indicates that oil futures prices do martingale; therefore, the methodology can be applied to the data set under consideration.

[1]See Dominguez (1989), Kumar (1992), and Deaves and Krinsky (1992).

The equations giving the upper $[\overline{C}_{G,t}, \overline{P}_{G,t}]$ and lower $[\underline{C}_{G,t}, \underline{P}_{G,t}]$ bounds for call and put options, respectively, are the following:[2]

$$\overline{C}_{G,t}(K) = E_t[\max\{0, F_T - K\}] \tag{1a}$$

$$\underline{C}_{G,t}(K) = \max\{E_t[F_T] - K, e^{-r_f T}E_t[\max\{0, F_T - K\}]\} \tag{1b}$$

$$\overline{P}_{G,t}(K) = E_t[\max\{0, K - F_T\}] \tag{1c}$$

$$\underline{P}_{G,t}(K) = \max\{E_t[K] - F_T, e^{-r_f T}E_t[\max\{0, K - F_T\}]\} \tag{1d}$$

where F_T denotes the arbitrary futures price at the expiration of the option, G represents the probability distribution from which the arbitrary futures price at the expiration of the option will be drawn, and K denotes the option's strike price. Time periods are indexed prior to the option's expiration so that T expresses the total time to expiration. $E_t[\cdot]$ represents the expectation taken with respect to the risk-neutral distribution that the asset returns are assumed to satisfy at current time t; r_f is the risk-free interest rate.

The upper and lower bounds in Equations 1a–d differ only in the discounting factor. An immediate implication is that the bounds will be extremely tight in situations in which interest rates are low and maturity is close. Even if interest rates are high, however, and maturity time is long, these equations adequately bound the American option price.

The main advantage of this set of bounds is that they are expressed in terms of the terminal distribution alone. They can be calculated by the integration of the option's payoff function with respect to the underlying asset's risk-neutral distribution. Parametric option-pricing approaches assume a specific functional form for this distribution, which implies that the values of these bounds are easily obtainable within the framework of such a methodology.

Functional Form of the Terminal PDF and European Option Prices

In this article, the use of an ESE is proposed as a method to parameterize the functional form of the risk-neutral density of the underlying asset. The true risk-neutral density is expressed as a series expansion around a known approximating distribution. The coefficients in the expansion have the desirable property of being simple functions of the moments of the true risk-neutral distribution, thus accounting for

[2]The validity of these bounds was discussed by Melick and Thomas (1996). They suggested that in markets in which continuous trading takes place (e.g., NYMEX), option prices can be expressed in terms of these bounds.

differences between the moments of the true and approximating risk-neutral distributions. This is an appealing feature of the ESE methodology because the parameters involved in the model have an intuitive meaning, rather than just being abstract mathematical quantities.

If the approximating distribution of futures returns as $p(F_T)$, the true risk-neutral distribution $g(F_T)$ can be obtained is a series expansion in terms of $p(F_T)$:

$$g(F_T) = p(F_T) + \frac{(\kappa_2(g) - \kappa_2(p))}{2!}\frac{d^2p(F_T)}{dF_T^2} - \frac{(\kappa_3(g) - \kappa_3(p))}{3!}\frac{d^3p(F_T)}{dF_T^3} + \frac{(\kappa_4(g) - \kappa_4(p)) + 3(\kappa_2(g) - \kappa_2(p))^2}{4!}\frac{d^4p(F_T)}{dF_T^4} + \varepsilon(F_T) \tag{2}$$

where $\kappa_1(g) = \kappa_1(p)$. In Equation 2, the terms $\kappa_i(g)$ and $\kappa_i(p)$ are the ith-order cumulants of the true and approximating risk-neutral distributions, respectively. Cumulants are similar to moments.[3] In fact, the first cumulant of a distribution is equal to its mean, and the second is equal to its variance. Jarrow and Rudd (1982) argued that, with a good choice for the approximating distribution, higher order terms in the remainder $[\varepsilon(F_t)]$ are likely to be negligible. They suggested the use of the lognormal distribution for the role of the approximating distribution, the preeminence of which is an undoubted fact in the option-pricing literature.

Using the lognormal distribution yields the following formulae for European-style call and put options:[4]

$$V_{i,t}(F_T; g) = V_{i,t}(F_t; p) - e^{-rT}\frac{(\kappa_3(g) - \kappa_3(p))}{3!}\frac{dp(K)}{dF_T} + e^{-rt}\frac{(\kappa_4(g) - \kappa_4(p))}{4!}\frac{d^2p(K)}{dF_T^2} \tag{3}$$

$$i = \begin{cases} 1, & \text{call options} \\ 2, & \text{put options} \end{cases}$$

[3]The characteristic function of a probability distribution $\phi(t) = E(e^{itx})$ is calculated with the following Taylor series expansion:

$$\phi(t) = \exp\left\{\sum_{r=1}^{\infty}\frac{\kappa_r t^r}{r!}\right\} = \sum_{r=0}^{\infty}\frac{\mu_r' t^r}{r!}$$

where μ_j (moments) is the coefficient of $\frac{(it)j}{j!}$ and k_j is the coefficient of $\frac{(it)j}{j!}$ log $\phi(t)$.

[4]Longstaff (1995) suggested that this general option-pricing model includes many other option-pricing models as special cases. Examples of models that are nested within this general model include Black and Scholes's (1973) model, Merton's (1973) stochastic interest rate model, and Merton's (1976, Equations 17 and 18) jump diffusion models. In addition, because the risk-neutral density in Equation 2 can match the first four moments of any continuous density, the four parameters of the model can be chosen in such a way that Equation 3 closely approximates most existing option-pricing models.

where

$$V_{1,t}(F_T; p) = e^{-rT}[F_t N(d) - KN(d - \sigma\sqrt{t})] \tag{4a}$$

$$V_{2,t}(F_T; p) = e^{-rT}[KN(\sigma\sqrt{T} - d) - F_t N(-d)] \tag{4b}$$

$$d = \frac{\log\left(\frac{F_t}{K}\right) + \sigma^2 T/2}{\sigma\sqrt{T}}$$

where $N(\cdot)$ is the cumulative standard normal and σ^2 is the variance of the asset. In Equation 3, the first and second cumulants of the true and approximating distribution have been equated, that is, $\kappa_1(g) = \kappa_1(P)$ and $\kappa_2(g) = \kappa_2(P)$, in agreement with Jarrow and Rudd (1982). This argument is also justified on numerical grounds by Corrado and Su (1996), who noticed that, without this condition, there exists a problem of multicollinearity between the second and fourth moments.

Estimation Procedure

This article follows, to a certain extent, Corrado and Su (1996). The option-pricing equation for the European option from Equation 3 can be expressed in a simplified way as

$$V_{i,t}(F_T; g) = V_{i,t}(F_T; p) + \lambda_1 Q_3 + \lambda_2 Q_4 \tag{5}$$

where the terms λ_i, $i = 1$ and 2, and Q_j, $j = i + 2$, are defined as follows:

$$\lambda_i = \gamma_j(g) - \gamma_j(p) \tag{6a}$$

$$Q_3 = -F_t^3(e^{\sigma^2 T} - 1)^{\frac{3}{2}} \frac{e^{-rT}}{3!} \frac{dp(K)}{dF_T} \tag{6b}$$

$$Q_4 = F_t^4(e^{\sigma^2 T} - 1)^2 \frac{e^{-rT}}{4!} \frac{d^2 p(K)}{dF_T^2} \tag{6c}$$

In Equations 6a–c, $\gamma_3(g)$ and $\gamma_3(p)$ are skewness coefficients for the true and approximating distributions, respectively, whereas $\gamma_4(g)$ and $\gamma_4(p)$ are excess kurtosis coefficients. It is known that the relation between skewness and excess kurtosis coefficients and the cumulants of a distribution is given by the following formulae:

$$\gamma_3(g) = \frac{\kappa_3(g)}{\kappa_2(g)^{\frac{3}{2}}} \quad \text{and} \quad \gamma_4(g) = \frac{\kappa_4(g)}{\kappa_2(g)^2} \tag{7}$$

The cumulants of the risk-neutral distribution are given by the following relations:

$$\kappa_1(g) = F_t \tag{8}$$

$$\kappa_2(g) = \kappa_2(p) \tag{9}$$

where F_t is the value of the futures contract today. Equation 8 essentially sets the mean of the implied PDF equal to the recorded settle price of the futures to satisfy the risk-neutrality argument.

Jarrow and Rudd (1982) showed that, when the approximating distribution belongs to the family of lognormal distributions, the volatility parameter σ is the solution to the following equation:

$$\kappa_2(p) = \kappa_1^2(p)(e^{\sigma^2 t} - 1)$$

If the substitution $q^2 = e^{\sigma^2 t} - 1$ is made, the skewness and excess kurtosis coefficients can be written as

$$\gamma_1(p) = 3q + q^3$$
$$\gamma_2(p) = 16q^2 + 15q^4 + 6q^6 + q^8$$

Combining all the formulae and substituting them back in Equation 5, we have remaining three unknown parameters to be estimated for the price of the European option. To express the option price as a function of the bounds, a weight is introduced.[5] The price of any option on a given day is expressed as a linear combination of the two bounds, and the number of the unknown parameters rises to four:

$$C_t[K] = w_t\overline{C}_{G,t}[K; \sigma, \lambda_1, \lambda_2]_t + (1 - w_t)\underline{C}_{G,t}[K; \sigma, \lambda_1, \lambda_2]_t \tag{10}$$

$$P_t[K] = w_t\overline{P}_{G,t}[K; \sigma, \lambda_1, \lambda_2]_t + (1 - w_t)\underline{P}_{G,t}[K; \sigma, \lambda_1, \lambda_2]_t \tag{11}$$

Optimization Constraints

On any given day, the implicit parameters $(\sigma, \lambda_1, \lambda_2, w)_t$ are extracted by minimization of the following sum of squares:

$$\min_{\sigma,\lambda_1,\lambda_2,w} \sum_{K_i} \{(C_t(K_i) - C_{obs,t}(K_i))^2 + (P_t(K_i) - P_{obs,t}(K_i))^2\} \tag{12}$$

[5] The same tests were carried out with the use of two weights, one for in-the-money options and one for out-of-the-money options. The distributions recovered were almost identical, the only difference being a marginal improvement on the sum of squared errors. From an empirical point of view, the obtained accuracy does not compensate for the additional complexity introduced by the extra parameter.

where $C_t(K)$ and $P_t(K)$ are the theoretical call and put prices, respectively, as expressed in Equations 10 and 11. In the objective function, both the call and put options are included, contrary to Corrado and Su (1996, 1997), who used only the call option or only the put option. Under the risk-neutral framework, both puts and calls should be priced off the same distribution.

This is a constrained optimization problem because the four parameters cannot take any possible value. Parameter σ, for example, is the implied volatility and cannot take negative values, whereas parameter w representing the weight of the lower and upper bound can only take values between zero and one. Those two constraints that define the nature of σ, and w are necessary but not sufficient to extract a meaningful set of parameters.

The values recovered through Equation 12 when inserted in Equation 13 should give rise to a properly defined PDF. The truncated PDF is given by

$$g(F_T) = p(F_T) - \frac{(\kappa_3(g) - \kappa_3(p))}{3!}\frac{d^3p(F_T)}{dF_T^3} + \frac{(\kappa_4(g) - \kappa_4(p)) + 3(\kappa_2(g) - \kappa_2(p))^2}{4!}\frac{d^4p(F_T)}{dF_T^4} \tag{13}$$

The estimation procedure followed in this study differs from Corrado and Su (1996) in that it imposes additional constraints on the parameter values recovered from Equation 12 so that Equation 13 is always a valid PDF. Two additional conditions should, therefore, be satisfied for Equation 13 to represent a valid PDF. It should always be positive and have an integral from zero to infinity equal to one. Formally,

$$g(F_T) \geq 0 \qquad \forall\, F_T, F_T \in [0, \infty) \tag{14}$$

$$\int g(F_T) = 1 \tag{15}$$

These two constraints should always be imposed[6] whether the parameters recovered from Equation 12 are used only for pricing purposes as by

[6]Imposing the constraints of the approximating density being positive and integrating to one can make the interpretation of changes in PDFs more difficult if these constraints are binding. The alternative, however, to imposing the constraints would imply that no reliable statistics could be estimated off that distribution. The fact that an implied distribution estimated with the ESE methodology can take on negative values or not integrate to unity could be due in some circumstances to mathematical reasons alone. We, therefore, believe that imposing the constraints can be an advantageous practice.

Corrado and Su[7] or for recovering the risk-neutral PDF, as is the case in this article.

The optimization was performed with the MATLAB® constrained optimization routine that requires the function to be minimized, the upper and lower bounds for the parameters to be estimated, and any other constraints that should apply as inputs. Bounds for the parameters were set so that $0.01 \leq \sigma \leq 0.7$, $-2 \leq \lambda_1 \leq 2$, and $0 \leq \lambda_2 \leq 7$. None of these constraints was binding on any day of the sample period. The additional constraints included Equations 14 and 15, the latter requiring the associated integral to equal unity with a tolerance of ± 0.001. The algorithm focuses on the relevant Kuhn–Tucker equations and attempts to compute directly the Lagrange multipliers, with a quasi-Newton updating procedure. Methods of this nature are commonly called sequential quadratic programming methods and represent the state of the art in nonlinear programming methods.[8]

The optimization procedure was run with and without imposing the constraints, and the results indicated that it is essential that the constraints be imposed. Figures 1 and 2 show how the shape and information conveyed by the risk-neutral density function significantly differ, depending on the way the optimization is carried out.

In Figure 1(a), the parameters are recovered without the constraint that Equation 13 be always positive. This leads to negative values of the PDF for certain values of the underlying asset, which is mathematically not feasible. An apparent practical implication of recovering a PDF with the shape of Figure 1(a) would be the overestimation of the probability that the underlying asset takes values below 15, resulting in the overpricing of put options with strike prices within that range. As is obvious from Figure 1(a), the distribution is more peaked in the former, having greater mass below the value of 15 to compensate for the fact that it becomes partly negative beyond that point while being forced to maintain an integral equal to 1 at all times.

A situation very similar to the one described in the previous paragraph arises if the restriction that the integral of the recovered distribution function be equal to one is not imposed.

In Figure 2(a), the parameters are recovered without the constraint that the integral of the distribution recovered be equal to one. Imposing

[7]In their article, the authors did not mention anything about these constraints, and a simple substitution of the parameter values they used shows that they did not impose any.

[8]Schittowski (1985) implemented and tested a version that outperformed every other tested method in terms of the efficiency, accuracy, and percentage of successful solutions over a large number of test problems. For more information on the relevant procedures, the reader is referred to MATLAB® Optimization Toolbox User's Guide, Version 5.

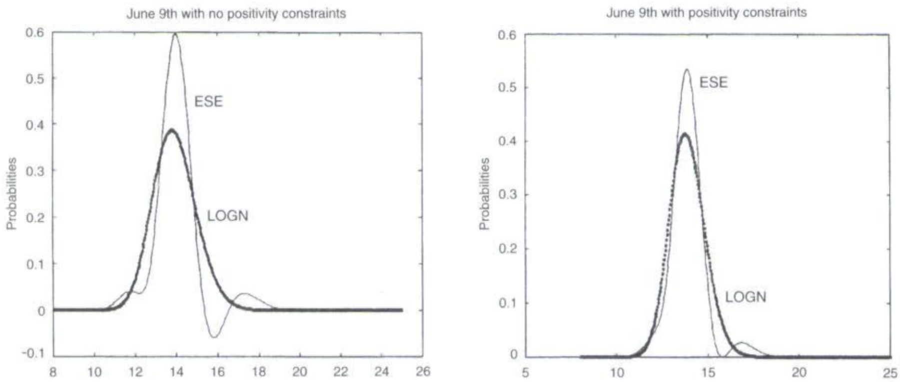

FIGURE 1
Probability distributions recovered for July 1998 contracts on June 9, 1998 without and with positivity constraints.

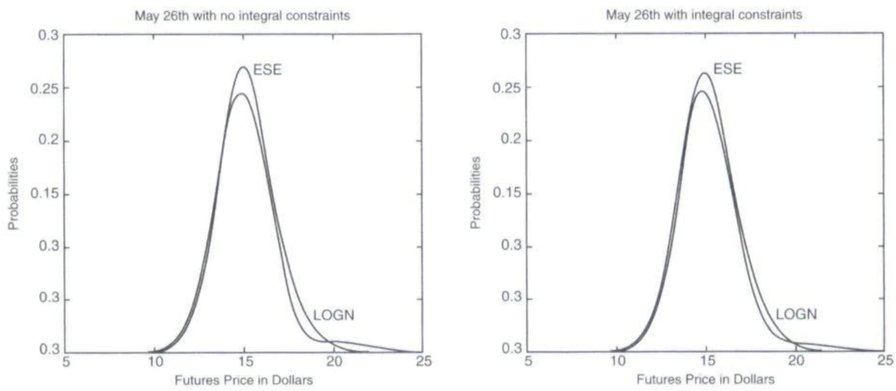

FIGURE 2
Probability distributions recovered for July 1998 contracts on May 26, 1998 without and with integral constraints.

this restriction leads to a significantly different shape for the implied risk-neutral PDF depicted in Figure 2(b). The shape of the distribution in this figure approaches more that of the fitted lognormal. It is less leptokurtic, and its right tail is considerably smoother. For a policy maker, this implies that, relying on a distribution of Figure 2(a) to extract the market sentiment, he would find that there are two possible scenarios in the market, indicating a strong bimodality, whereas the situation would be much more normal than would have been thought.

DATA

The New York Mercantile Exchange (NYMEX) West Texas Intermediate (WTI) futures contract is the world's largest volume futures contract trading on a physical commodity. It offers excellent liquidity and price transparency that makes it appropriate for a pricing benchmark.

Our data set contains information for options on the WTI futures contract traded on NYMEX for the period January 1, 1998 to December 31, 1998 and includes bid and ask prices, settlement prices, trading volumes, and open interest for the options. Data for the underlying futures contracts and risk-free interest rates (Treasury bills with maturity dates close to the options' maturity) were obtained from Datastream.

The data set was refined so that it did not contain meaningless information. Options with trading volume of less than five contracts were excluded because of illiquidity reasons. Also, on a given day contracts with different strike prices that had the same premium were excluded. Finally, the last 3 trading days of the options contracts were not included in the data set because of any undesirable noise information that they might have conveyed.

Settlement prices were used to represent the value of the options. The settlement price is determined at the end of each day by a settlement committee made up of roughly 20 market participants. The committee frequently relies on the average of bid and ask prices, during the last minutes of trading, as starting points for the settlement prices. Heavily traded options are priced first, with put–call parity used to price low-volume options at the same strike when the futures markets has settled (Melick & Thomas, 1997).

Market news data were collected from two individual sources. One was the Financial Times of London newspaper, and the other was the FT Discovery (FT Electronic Publishing), which covers a wide variety of news vendors. All the relevant articles were reviewed, and those of greater importance were used for our analysis.

APPLICATIONS AND RESULTS

This section examines whether the qualitative information recovered from option prices is consistent with the market commentary at the time our data were collected. This will further confirm the validity of the proposed methodology. The model is also compared with a single lognormal model to investigate whether the additional complexity introduced is associated with significantly superior explanatory ability. Finally, the model's ability to recover robust implied PDFs is examined.

A study of implied distributions, however, especially when large potential price changes are being analyzed, has some inherent limitations. For most markets, there are very few, if any, liquid option contracts trading at strike prices in far-from-the money regions. The shape of the recovered PDF in these regions, unfortunately, depends more on the choice of the estimation method or smoothing technique used rather than on the data themselves (see the discussion in the Stability Tests section). The ESE methodology is no exception to this rule, and one should be cautious when making inferences in these outer regions. At best, one should ensure that the recovered PDF is consistent with available data (internally consistent) and also economically sensible.

Because the primary task is to investigate whether the distributions recovered from the options prices were consistent with the market commentary then, it was not considered necessary to examine all existing contracts in the next two sections. Contracts due to expire in a month when oil prices followed an abnormal pattern were selected because it was more likely that implied distributions would deviate from a lognormal during such periods.

Figure 3 suggests that two appropriate candidates for this task are the July and December futures contracts and, consequently, the respective options contracts.

The same contracts are used in the Stability Tests section, in which the robustness of the implied PDFs is examined. This was driven by the argument that, if the recovered PDFs proved to be robust in turbulent

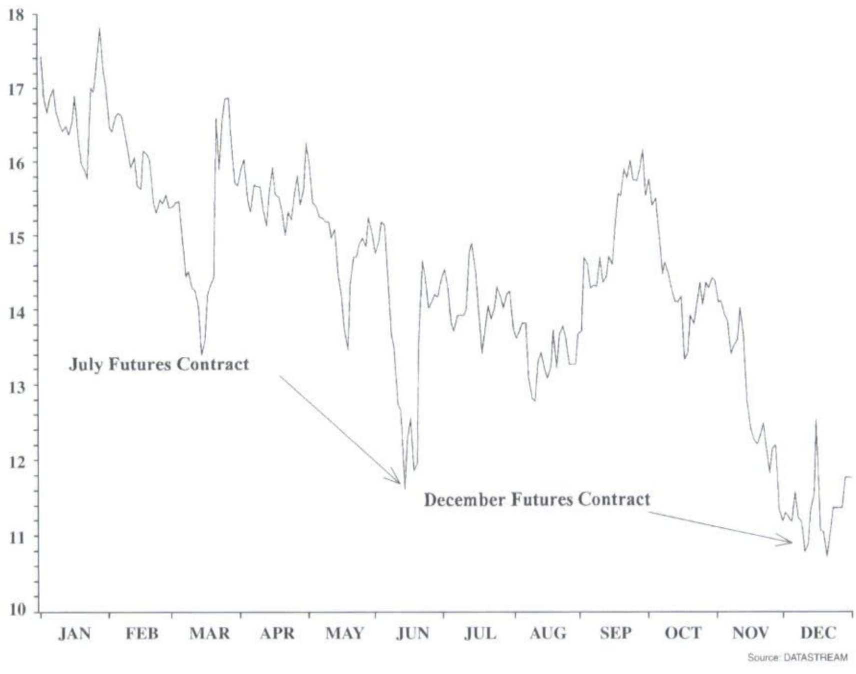

FIGURE 3

Near-month WTI futures contract in dollars per barrel from January 2, 1998 to December 31, 1998.

periods, there would not be any reason they were not during more settled ones.

In contrast to that, for the model's statistical comparison in the Statistical Comparison section, it is essential that we use all available information. The whole of the period, January 1, 1998 to December 31, 1998, was, therefore, used.

Multimodal Scenaria

The data for the July futures options span a time period of around 40 days, starting from the May 1 until the June 10, 1998. During this period, oil prices were sliding because of increased supply in the oil market, whereas, at the same time, demand was at relatively low levels. Demand in Asia had fallen because of the Asian crisis. The International Energy Agency expected demand for oil for the whole year to increase only by 1.7%, whereas the percentage increase the previous year was 2.7%. In a meeting that took place in Riyadh in March, OPEC and non-OPEC countries agreed to make a total cut in production of 1.5 million barrels per day to stabilize the oil market.

During May, there was growing concern as to whether it was feasible for oil prices to sustain the then current levels without any further action on behalf of the oil-producing countries. Oil prices were expected to drop even more unless further cuts in production were made that would stabilize or even raise them. This bimodality in traders' beliefs obviously could not be detected simply by observation of the price of oil futures contracts. However, the use of an implied distribution is assumed to offer greater insight into investors' available information. If the distributions recovered by methods described in this article are to be used to recover the market sentiment in turbulent periods, they should be capable of revealing the bimodal scenario that was dominant in investors' expectations at the time.

As can be seen in Figure 4, the distribution recovered with the ESE does, indeed, capture the bimodal scenario. It assigns a greater mass than a fitted lognormal does in values around the mean (thus being more leptokurtic); at the same time, it assigns larger probability mass at the right tail, consistent with the scenario of further anticipated future cuts. This bimodal pattern is a standard feature of most of the PDFs recovered for that period.

Despite OPEC's decision on June 24 for further cuts in oil production, the market consensus did not dramatically change in the subsequent months. Implied distributions recovered for the period

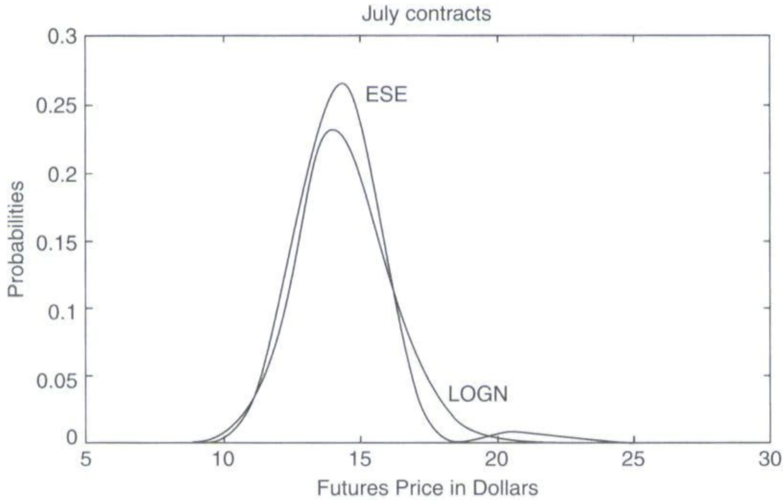

FIGURE 4
Probability distribution implied by July 1998 options contracts on May 12, 1998.

October–November 1998 exhibited very similar characteristics to those implied by the July contracts. Autumn 1998 was a period of great uncertainty for oil prices. Prices declined, mainly because of increasing supplies because non-OPEC members were not willing to comply with the organization pumping limits. In addition, Venezuela, one of the major oil-exporting countries, announced that the previously decided exploration cuts would cease to apply. The International Monetary Fund lowered its predictions for oil prices for the current and following years, and the American Petroleum Institute reported a surge in U.S. crude oil stocks. Iraq's confrontation with the United Nations arms inspectors and the forecast for a cold northern winter gave rise to speculation for higher oil prices. A bimodal pattern that was present in most of the recovered distributions from December contracts, similar to that observed in the July contracts' distributions, further confirms the consistency of our model.

Isolated Events

Market sentiment reflects investors' expectations about future values of the underlying asset, which are formed on the basis of available information. However, news arrives randomly in the market on an every day basis, forcing investors to constantly revise their expectations in the light of this newly arrived information. This section investigates whether implied distributions recovered in this article are consistent with this

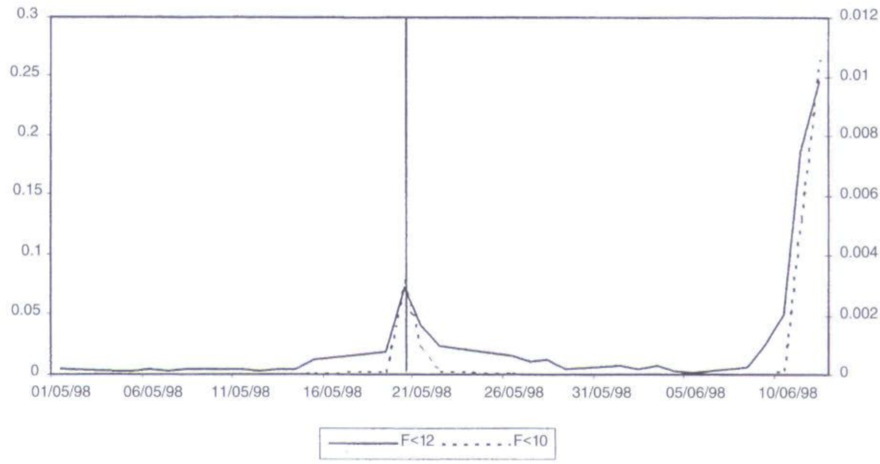

FIGURE 5
Implied probabilities of the WTI futures price being lower than \$12 or \$10 per barrel for July contracts from May 1, 1998 to June 12, 1998.

feature and are able to reflect the extent to which important news arrival has an impact on market participants' reactions. To this end, a series of isolated events and the information one could infer by recovering the implied distributions of option prices were examined.

On May 20, 1998, the American Petroleum Institute published data showing an 8.79 million barrel increase in U.S. crude stocks when the market expected a fall of 2.5 million. Stocks went up to 353.13 million barrels, the biggest stock of crude oil since August 1993, and the news was reflected in the implied distribution of WTI. In a graph showing the probability of WTI being less than \$12 or less than \$10 per barrel (Fig. 5), a peak was observed on that day, with the former probability being 3 times higher than it had been the previous day and the latter being 20 times larger than its corresponding value the previous day.

In an examination of the December options contracts, an exactly similar situation arises on October 20, 1998, when Venezuela's energy minister announced that his country's oil output would return to normal levels in the second half of 1999, after the two production cuts in 1998. This was in contrast to traders' hopes for a sustained reduced oil production throughout the whole of 1999. This is highlighted in Figure 6, where probabilities exhibit a jump around the October 20.

On Thursday, June 4, an unexpected meeting took place in Amsterdam of Saudi Arabia, Venezuela, and Mexico. The three oil-exporting countries announced that they would reduce production by a

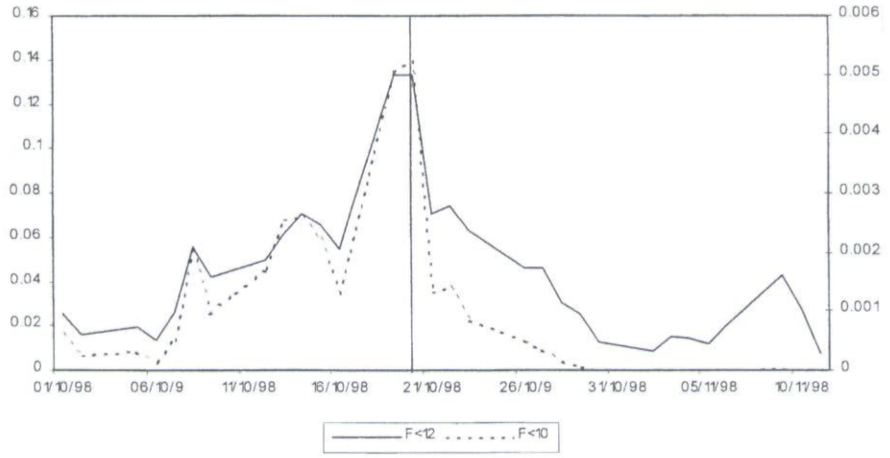

FIGURE 6
Implied probabilities of the WTI futures price being lower than \$12 or \$10 per barrel for December contracts from October 1, 1998 to November 11, 1998.

total of half a million barrels per day. The initial reaction of the markets was positive; prices rose on that Friday. During the weekend that followed, however, investors realized that the cuts were not to the extent that they had expected and that the possibility of further cuts seemed more distant now than before.[9]

Observing the distributions from that day onward until the last day the July contract was traded, we find a new feature present. Both the right and left tails have larger mass than that of the fitted lognormal, whereas only the right tail had this property before. The central part of the distribution remained more peaked. The heavier left tails of the distribution reflect the growing market belief that it was much more likely now, than before, that oil prices would slide. Distributions are now more peaked, as shown in Figure 7, because the maturity day of the contract (June 17) is approaching and the uncertainty of the level of the futures price at expiration is, to a great extent, resolved.

This feature is clearly evident in Figure 8, which plots the interquartile ranges[10] of the recovered distributions across time and which shows that the distributions recovered are consistent with the fact that uncertainty tends to fall as expiration approaches (Melick & Thomas, 1999).

[9]Hopes for further cuts were placed at the June 24th meeting of OPEC.

[10]The interquartile range of a distribution is defined as the difference between the 75 and 25% quartiles. It is the interval that contains the central 50% of the distribution.

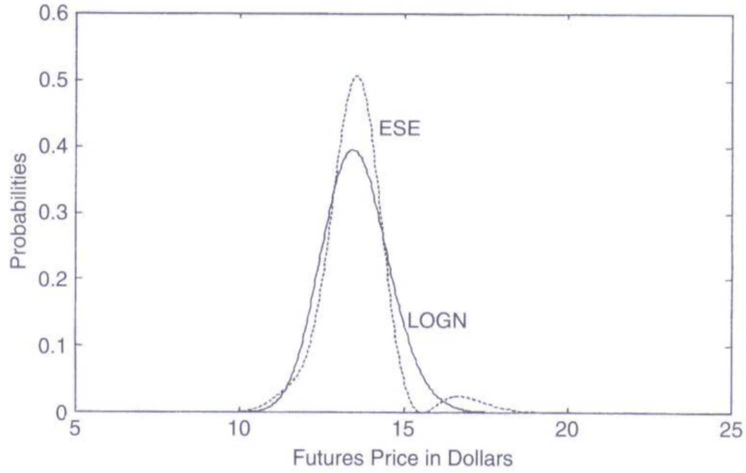

FIGURE 7
Probability distribution implied by July 1998 options contracts on June 10, 1998.

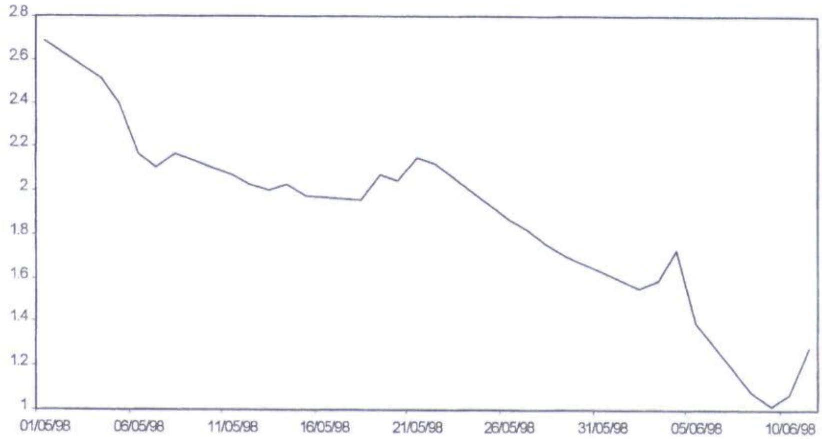

FIGURE 8
Interquartile range for July 1998 contracts from May 1, 1998 to June 12, 1998.

Statistical Comparison

This section examines the goodness of fit of the proposed model versus that of the fitted single lognormal one. Even though there is enormous literature documenting the superiority of alternative parametric and nonparametric models over the Black–Scholes model, to the authors' knowledge, existing studies (the only exemption being Melick & Thomas,

1997) make comparisons in terms of the mean-squared errors and offer no discussion on the statistical significance of the accuracy that is achieved with a model with parameters additional to those of the Black–Scholes model. Our concern was to examine whether the use of the ESE model would offer a statistically significant better fit, not just a better fit.

To this end, a single lognormal model was estimated with the same optimization procedure that was used for the series expansion model. Option prices were expressed as a combination of the bounds described in Equations 1a–d, but this time, the expectation was taken over a single lognormal distribution. This implies that the optimization procedure involved the estimation of two parameters, namely, the implied standard deviation and the weighting factor. The model can, therefore, be seen as a nested one within the ESE model.

The ESE model is expected to provide a better fit to observed option prices than the nested single lognormal one because it involves the estimation of additional parameters. If the single lognormal distribution could adequately describe the distribution of the underlying asset returns, then the use of terms correcting for skewness and kurtosis in the proposed method would not significantly improve the fit to observed option prices.

To test for the significance of these terms, we need to employ a test that compares two competing nonlinear models. Let $f[\sigma, \lambda_1, \lambda_2, w/\Omega]$ denote the option valuation formula for a call or put option that results from the ESE method (Equations 10 and 11). Ω represents an information matrix that contains strike prices, interest rates, and times to maturity. If $g[\sigma, w/\Omega]$ is the corresponding formula consistent with the Black–Scholes assumptions, it is obvious that the latter formula arises from the former by imposing the restrictions that λ_1 and λ_2 equal zero.

Because the two option-pricing formulae do not represent different functional forms, $g[\cdot]$ is nested within $f[\cdot]$; a likelihood-ratio (LR) test, testing the hypothesis that the two coefficients are jointly zero, can be applied. It will effectively test whether the ESE method provides superfluous information for the description of the underlying asset's distribution or whether the additional complexity introduced by our model is, indeed, useful in explaining distributional characteristics of option prices. The test had to be carried out separately for call and put options because they have different functional forms.

Table I reports the results of the LR test carried out for the entire data set. The statistic $-2(L_R - L_{UR})$ follows an asymptotic $\chi^2(m)$, where

TABLE I

Likelihood-Ratio Tests for Call and Put Options for January 12, 1998 to December 31, 1998

Call Options

$$\text{Restricted: } C_{obs} = \longrightarrow C_{th}(\sigma, 0, 0, w|\Omega)$$

$$\text{Unrestricted: } C_{obs} = \longrightarrow C_{th}(\sigma, \lambda_1, \lambda_2, w|\Omega)$$

Contracts	*Observed*	*Log-Likelihood Restricted* (L_R)	*Log-Likelihood Unrestricted* (L_{UR})	$-2(L_R - L_{UR})$
May	1,051	967.74	1,211.82	488.16
Jun	1,014	919.54	1,126.39	413.70
Jul	961	841.32	1,120.88	559.12
Aug	837	1,002.71	1,107.35	209.28
Sep	837	540.01	782.49	484.96
Oct	694	735.16	946.36	422.40
Nov	590	655.64	869.23	427.18
Dec	1,789	526.50	1,196.79	1,340.58

Put Options

$$\text{Restricted: } P_{obs} = \longrightarrow P_{th}(\sigma, 0, 0, w|\Omega)$$

$$\text{Unrestricted: } P_{obs} = \longrightarrow P_{th}(\sigma, \lambda_1, \lambda_2, w|\Omega)$$

Contracts	*Observed*	*Log-Likelihood Restricted* (L_R)	*Log-Likelihood Unrestricted* (L_{UR})	$-2(L_R - L_{UR})$
May	916	700.06	788.04	175.97
Jun	643	536.20	631.55	190.70
Jul	638	452.68	587.96	270.56
Aug	588	394.17	517.49	246.64
Sep	597	322.43	552.21	459.56
Oct	456	507.24	572.46	130.46
Nov	438	483.78	653.78	339.99
Dec	935	422.35	534.76	224.82

m is the number of parameter restrictions. The restricted model (the Black–Scholes model) imposes two restrictions: the skewness of the distribution is 0 and the kurtosis of the distribution is known. The number of parameter restrictions are, therefore, two ($m = 2$), and the critical values for $\chi^2(2)$ at the 95 and 99% levels are 5.99 and 9.21, respectively. A statistic greater than 5.99 implies that the null hypothesis that the restrictions apply can be rejected.

The test indicated that, in both cases, the coefficients that correct the risk-neutral distribution assumed within the Black–Scholes model for non-zero skewness and excess kurtosis are jointly statistically significant at the 5 and 1% levels for each of the contracts examined. This implies that the proposed functional form for the distribution of the underlying asset's returns serves as a more realistic representation of the true distribution than the standard single lognormal assumption. The fact that the derived results are independent of the data set used further enhances the consistent superiority of the proposed model.

Melick and Thomas (1997) performed similar tests under the assumption that the underlying asset returns distribution is a mixture of three lognormal distributions. In their tests, however, they compared their model with Barone-Adesi and Whaley's (1987) model, arguing that it was the most commonly used model among practitioners. Available evidence,[11] however, examining the relative performance of alternative American option-pricing models highlights that, even though Barone-Adesi and Whaley's model is the fastest such model, it is also the least accurate one.

The specification of the single lognormal model, as suggested in this article, makes it consistently comparable with the model against which it is being tested. Both models use the same bounds for the American option price, being different only in the assumed distribution's functional form. Intuitively, this seems to be a relatively more valid way to assess the distributional differences through the examination of option prices.

Stability Tests

The previous section examined the statistical significance of the improvement of the fit achieved with an ESE model versus that achieved with a fitted single lognormal one. The goodness of fit, however, should not be the sole criterion to verify the validity or superiority of the ESE model. As recent studies highlight,[12] issues arising from the quality of the data used as inputs for estimating implied PDFs and the effect of possible errors on the estimation of the PDFs should also be addressed.

This section examines the robustness of the ESE methodology versus that of a mixture of two lognormals distribution model (MLN).[13] The methodology employed is thoroughly described in Bliss and Panigirtzoglou (2000) and consists of two main steps: first, the recorded options prices are perturbed by a random quantity between $-1/2$ and

[11]See Broadie and Detemple (1996).

[12]See Bliss and Panigirtzoglou (2000) and Melick and Thomas (1998).

[13]The mixture of two lognormal specifications was chosen because it is a very commonly used parametric model, also studied in terms of stability by Bliss and Panigirtzoglou (2000), Cooper (1999), and Melick and Thomas (1998).

$+1/2$ a tick, independently generated with a continuous uniform distribution; and second, implied PDFs are repeatedly estimated from the artificial options' cross sections.

For every trading day, a set of 100 observationally equivalent options' cross sections were produced. The magnitude of the perturbation was chosen so that it reflected the possible measurement error. The size of the tick for the WTI option contract is 0.01.

The second step included the estimation of implied PDFs from the artificially produced options' cross sections, with the ESE and MLN methods.

As it is literally impossible to examine the stability of the actual PDFs, the stability of their descriptive statistics was assessed instead. This study considered an extended set of what is called a *standard set of results* by Melick (1999), which also included the skewness and kurtosis coefficients. The full set included the following descriptive statistics of the implied PDFs:

$\hat{\sigma}$: The standard deviation or second central moment of the distribution.[14]

$\widehat{Sk}$: The skewness coefficient, defined as the third central moment normalized by the cube of the standard deviation.

$$\widehat{Sk} = \frac{\hat{m}^3}{\hat{\sigma}^3}$$

$\widehat{Kurt}$: The kurtosis coefficient, defined as the fourth central moment normalized by the square of the variance

$$\widehat{Kurt} = \frac{\hat{m}^4}{\hat{\sigma}^4}$$

The percentiles $\hat{X}_{005}$, $\hat{X}_{01}$, $\hat{X}_{05}$, $\hat{X}_{10}$, $\hat{X}_{25}$, $\hat{X}_{50}$, $\hat{X}_{75}$, $\hat{X}_{90}$, $\hat{X}_{95}$, $\hat{X}_{99}$, and $\hat{X}_{995}$. $\hat{X}_n$ is the solution to the equation $n = F(\hat{X})$, with interpolation where necessary.

The estimation procedure, as well as the computation of the statistics listed previously, involved the use of numerical techniques that may, in some cases, not converge to a solution within the expected range, thus generating outlying observations. To further proceed with our analysis, we excluded all values outside the range defined by the 0.5 and

[14]The mean of the distribution was not considered because, in both methodologies, it was set equal to the risk-neutral mean, the futures price.

99.5 percentiles of each of the daily distributions of the 100 observations, thus placing 99% confidence in our results. In addition, if a recovered distribution within a daily series resulted in at least one outlier summary statistic, it was excluded from the series.

This filtering process resulted in the removal of 294 observations (10.5% of the sample) for the parameters recovered with the ESE method for the July WTI contracts and 287 observations (10.3% of the sample) for the parameters recovered with the MLN method. The respective numbers of observations removed from the December contracts are 296 (9.9% of the sample) and 351 (11.7% of the sample).

The robustness of the two estimation techniques was then assessed in terms of the stability of the convergence to the original solution and the stability at the solution. The former refers to the ability of the model to estimate, on average, PDFs whose summary statistics converge to the respective statistics calculated from the original data. To quantify this, we constructed the absolute percentage deviation (APD) measure, which is given by the following equation:

$$\mathrm{APD}_{Z,t} = \left| \frac{Z_{\text{unperturbed},t} - \bar{Z}_{\text{perturbed},t}}{Z_{\text{unperturbed},t}} \right| \tag{16}$$

Equation 16 expresses the discrepancy between the summary statistic Z at date t as computed from the implied PDF of the unperturbed data and the trimmed mean of the respective statistic taken over 100 observations resulting from the perturbed option prices, also at date t, as a percentage of Z. This measure is less sensitive to the presence of outliers, as 0.5% on each side of the distribution of $Z_{\text{perturbed},t}$ is excluded.

Because our objective is to measure the magnitude of this deviation, we consider the absolute operator.

For stability at the solution, the objective is to examine the stability of the estimates themselves by an examination of the dispersion of a particular statistic within a daily series. Typical measures of dispersion (i.e., the standard deviation) are expressed in terms of units of the variate, thus making it difficult to compare dispersions in different populations. To overcome this deficit, we used Karl Pearson's coefficient of variation, which is defined as follows:

$$CV = \frac{\sigma}{\mu'_1}$$

where σ is the standard deviation of the population and μ'_1 is the first centralized moment. To account for potential biases due to outliers in the sample, we replaced σ with a less sensitive dispersion measure, namely, the Pearson and Tukey (1965) robust measure of dispersion,

$$\sigma_{\mathrm{PT}} = \frac{X_{95} - X_{05}}{3.25}$$

and μ'_1 with the trimmed mean $\overline{Z}_{\mathrm{perturbed}}$.

The robust coefficient of variation (RCV) is consequently defined by

$$\mathrm{RCV}_{Z,t} = \left| \frac{\sigma_{\mathrm{PT}_{Z,t}}}{\overline{Z}_{\mathrm{perturbed},t}} \right| \tag{17}$$

The quantities $\sigma_{\mathrm{PTZ},t}$ and $\overline{Z}_{\mathrm{perturbed},t}$ now refer to the distribution of the summary statistics alone (100 observations per day). The RCV measures the dispersion of the estimated statistic and is free of any possible biases associated with the magnitude of the statistic under examination. By construction, the lower the RCV is, the less dispersed the calculated summary statistics are. The absolute value is taken for reasons mentioned earlier in the text.

The measures defined by Equations 16 and 17 are expressed in percentage terms and are free of any measurement unit. This provides great flexibility to our analysis for two reasons: first, it allows us to assess the performance of the models, regardless of how close the summary statistics of the implied distributions are with the summary statistics of the market distribution or, in other words, of how well the assumed distribution approximates the true distribution, and second, it makes it possible to compare the summary statistics among themselves and identify those on which we can place more confidence.

Tables II and III report average APDs and average RCVs for the summary statistics $\hat{\sigma}$, $\widehat{Sk}$, $\widehat{Kurt}$, and $\hat{X}$ of the implied distributions recovered from the July and December WTI contracts, respectively. Reporting only the average APDs and RCVs, however, may lead to an inaccurate interpretation of the results, as the mean is quite sensitive to the outlying parts of the frequency distribution, especially for an empirical one. To deal with this, we also computed the median APDs and RCVs; Tables IV and V report the median APDs and median RCVs for the summary statistics $\hat{\sigma}$, $\widehat{Sk}$, $\widehat{Kurt}$, and $\hat{X}$ of the implied distributions recovered from the July and December WTI contracts, respectively.

The evidence presented in Tables II–V are supportive of the superior performance of the ESE model. Across the summary statistics, the ESE method appears more stable than the MLN method, consistently leading to lower average APDs and average RCVs. The ESE method appears to

TABLE II

Average Absolute Percentage Deviations (APDs) and Robust Coefficients of Variation (RCVs) of July WTI Futures Options Implied Probability Density Function Summary Statistics

Descriptive Statistic	APD (%) ESE	APD (%) MLN	RCV (%) ESE	RCV (%) MLN
$\hat{\sigma}$	0.300	0.463	0.333	0.714
$\widehat{Sk}$	5.933	4.353	8.978	9.493
$\widehat{Kurt}$	1.694	4.835	1.715	7.503
$\hat{X}_{005}$	0.740	0.647	0.454	1.059
$\hat{X}_{01}$	0.896	0.500	0.556	0.772
$\hat{X}_{05}$	0.122	0.136	0.123	0.243
$\hat{X}_{10}$	0.109	0.046	0.089	0.124
$\hat{X}_{25}$	0.054	0.090	0.045	0.164
$\hat{X}_{50}$	0.010	0.027	0.027	0.065
$\hat{X}_{75}$	0.035	0.080	0.036	0.148
$\hat{X}_{90}$	0.022	0.087	0.033	0.145
$\hat{X}_{95}$	0.034	0.062	0.062	0.178
$\hat{X}_{99}$	0.130	0.390	0.147	0.598
$\hat{X}_{995}$	0.156	0.574	0.167	0.872
Observed	2,506	2,513		

TABLE III

Average Absolute Percentage Deviations (APDs) and Robust Coefficients of Variation (RCVs) of December WTI Futures Options Implied Probability Density Function Summary Statistics

Descriptive Statistic	APD (%) ESE	APD (%) MLN	RCV (%) ESE	RCV (%) MLN
$\hat{\sigma}$	0.041	0.652	0.272	1.128
$\widehat{Sk}$	12.019	9.529	30.980	19.997
$\widehat{Kurt}$	0.131	8.841	1.222	10.515
$\hat{X}_{005}$	0.042	1.244	0.298	1.661
$\hat{X}_{01}$	0.047	0.913	0.361	1.302
$\hat{X}_{05}$	0.013	0.168	0.111	0.347
$\hat{X}_{10}$	0.009	0.083	0.066	0.139
$\hat{X}_{25}$	0.006	0.137	0.036	0.237
$\hat{X}_{50}$	0.004	0.042	0.031	0.175
$\hat{X}_{75}$	0.004	10.209	0.030	6.908
$\hat{X}_{90}$	0.005	0.138	0.033	0.217
$\hat{X}_{95}$	0.008	0.089	0.062	0.178
$\hat{X}_{99}$	0.014	0.827	0.127	0.797
$\hat{X}_{995}$	0.016	0.909	0.142	1.154
Observed	2,704	2,649		

TABLE IV

Median Absolute Percentage Deviations (APDs) and Robust Coefficients of Variation (RCVs) of July WTI Futures Options Implied Probability Density Function Summary Statistics

	APD (%)		*RCV (%)*	
Descriptive Statistic	*ESE*	*MLN*	*ESE*	*MLN*
$\hat{\sigma}$	0.080	0.186	0.332	0.689
$\widehat{Sk}$	0.876	1.654	7.739	8.402
$\widehat{Kurt}$	0.316	1.675	1.895	7.519
$\hat{X}_{005}$	0.066	0.269	0.520	1.105
$\hat{X}_{01}$	0.075	0.200	0.638	0.847
$\hat{X}_{05}$	0.032	0.052	0.117	0.201
$\hat{X}_{10}$	0.021	0.021	0.089	0.101
$\hat{X}_{25}$	0.007	0.040	0.046	0.156
$\hat{X}_{50}$	0.003	0.008	0.028	0.045
$\hat{X}_{75}$	0.007	0.036	0.037	0.126
$\hat{X}_{90}$	0.011	0.030	0.031	0.153
$\hat{X}_{95}$	0.010	0.045	0.061	0.162
$\hat{X}_{99}$	0.026	0.187	0.154	0.612
$\hat{X}_{995}$	0.029	0.241	0.180	0.975
Observed	2,506	2,513		

TABLE V

Median Absolute Percentage Deviations (APDs) and Robust Coefficients of Variation (RCVs) of December WTI Futures Options Implied Probability Density Function Summary Statistics

	APD (%)		*RCV (%)*	
Descriptive Statistic	*ESE*	*MLN*	*ESE*	*MLN*
$\hat{\sigma}$	0.034	0.389	0.218	0.909
$\widehat{Sk}$	0.638	3.349	6.080	9.538
$\widehat{Kurt}$	0.093	4.295	0.527	10.340
$\hat{X}_{005}$	0.019	0.551	0.205	1.721
$\hat{X}_{01}$	0.032	0.500	0.263	1.349
$\hat{X}_{05}$	0.008	0.100	0.111	0.317
$\hat{X}_{10}$	0.008	0.030	0.054	0.114
$\hat{X}_{25}$	0.005	0.085	0.033	0.204
$\hat{X}_{50}$	0.003	0.016	0.030	0.066
$\hat{X}_{75}$	0.003	0.058	0.027	0.179
$\hat{X}_{90}$	0.004	0.101	0.030	0.189
$\hat{X}_{95}$	0.006	0.049	0.059	0.148
$\hat{X}_{99}$	0.010	0.492	0.111	0.852
$\hat{X}_{995}$	0.012	0.463	0.125	1.241
Observed	2,704	2,649		

recover not only distributions that are not greatly affected by the presence of measurement errors, leading to average distributions close to the ones recovered with the original data, but also distributions that, despite the presence of such errors, are very close to each other. Exceptions to these are the average APDs of $\widehat{Sk}$, $\hat{X}_{01}$, and $\hat{X}_{10}$ for the July contracts and of $\widehat{Sk}$ for the December contracts and the average RCVs for the December contract. This inconsistency can be attributed to the presence of outlying daily computed APDs and RCVs because this feature is not present in Tables IV and V.

As far as the stability of the individual statistics is concerned, this study confirms existing evidence that refers to the percentiles as a safer means for implied PDF analysis and interpretation. Not only is there a great degree of stability of the convergence to the original solution, as expressed by a low average APD and an even lower median APD being present, but also the dispersions around the means of the statistic distributions are very low, regardless of the model assumed. This also complements existing evidence, reported in Campa, Chang, and Reider (1998), Coutant, Jondeau, and Rockinger (in press) and Melick (1999), that for certain regions representing a large percentage of total probability (e.g., between $\hat{X}_{10}$ and $\hat{X}_{90}$), implied PDFs are relatively free of mathematical priors imposed by a specific economic model or structure and contribute the finding that the stability of these regions is also relatively independent of the model used. The percentiles $\hat{X}_{005}$ to $\hat{X}_{05}$ and also $\hat{X}_{95}$ to $\hat{X}_{995}$ appear with higher average and median APDs and RCVs, with the ESE model still providing the lower ones. This also supports the argument of Melick (1999) that the tail percentiles are sensitive to the choice of the estimation technique (in Tables II–V, see the rise in the average and median APDs and RCVs of both methodologies as we move further in the tails).

CONCLUSION

This article introduces a new methodology for recovering implied distributions from American option prices. The restricting assumption of normally distributed returns is relaxed, and a more realistic form for the distribution of the underlying asset is proposed. This is the first time that the ESE method has been used in an American options' framework.

The method can generate distributions exhibiting bimodal patterns, a typical characteristic of distributions belonging to the family of two-component mixtures of lognormal distributions. The proposed model, however, is more parsimonious, requiring the estimation of four parameters only under the risk-neutral assumption, with the respective number being five for a mixture of two lognormals model. This advantage can prove significant in markets in which the number of available actively traded strike prices is relatively small. Hence, possible situations of over-parameterization of the data are avoided.

The information conveyed by the implied probability distributions, recovered by the proposed method, proved to be consistent with the market commentary at the time. The implied distributions were shown to be able to capture the general market sentiment and also incorporate isolated events with a significant impact on the market.

In addition, the model's explanatory ability was tested against that of the corresponding nested single lognormal. It was found to offer a statistically significant better fit to observed option prices and could, therefore, be considered a superior means of extracting information implied in option prices.

Finally, the robustness of the ESE model was investigated, and the model was found to be capable of estimating on average PDFs whose summary statistics converge to the original solution and also capable of estimating statistics with a relatively low dispersion with respect to an alternative model that assumes a mixture of two lognormals distribution.

Implied risk-neutral distributions, however, can be used by policy makers or traders only to obtain qualitative information about the true distribution of the underlying asset. They do not represent the true distributions, so additional care should be taken when information is read from them.

BIBLIOGRAPHY

Backus, D., Foresi, S., Li, K., & Wu, L. (1997). Accounting for biases in Black–Scholes. New York: Stern School of Business, New York University.

Barone-Adesi, G., & Whaley, R. E. (1987). Efficient analytic approximation of American option values. Journal of Finance, 42, 301–320.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. Journal of Political Economy, 81, 637–659.

Bliss, R., & Panigirtzoglou, N. (2000). Testing the stability of implied probability density functions. Proceedings of the Chicago Board of Trade 13th Annual European Futures Symposium, Glasgow, Scotland.

Broadie, M., & Detemple, J. (1996). American option valuation. New bounds approximations and a comparison of existing methods. Review of Financial Studies, 9, 1211–1250.

Campa, J., Chang, K., & Reider, R. (1998). Implied exchange rate distributions: Evidence from OTC options markets. Journal of International Money and Finance, 17, 117–160.

Chaudhury, M. M., & Wei, J. (1994). Upper bounds for American futures options: A note. Journal of Futures Markets, 14, 111–116.

Cooper, N. (1999). Testing techniques for estimating implied RNDs from the prices of European and American-style options. Proceedings of the Bank for International Settlements Workshop on Estimating and Interpreting Probability Density Functions, Basel, Switzerland.

Corrado, C. J., & Su, T. (1996). S&P 500 Index option tests of Jarrow and Rudd's approximate option valuation formula. Journal of Futures Markets, 16, 611–629.

Corrado, C. J., & Su, T. (1997, Summer). Implied volatility skews and stock index skewness and kurtosis implied by S&P 500 Index option prices. Journal of Derivatives, 4, 8–19.

Coutant, S., Jondeau, E., & Rockinger, M. (in press). Reading interest rate and bond futures options smiles around the 1997 French snap election. Journal of Banking and Finance.

Cox, J. C., & Ross, S. A. (1976). The valuation of options for alternative stochastic processes. Journal of Financial Economics, 3, 145–166.

Deaves, R., & Krinsky, I. (1992). Risk premiums and efficiency in the market for crude oil futures. Energy Journal, 13, 93–118.

Dominguez, K. (1989). The volatility and efficiency of crude oil futures contracts. Oil and money: Coping with price risk through financial markets. Harvard International Energy Studies, 5, 48–87.

Jarrow, R., & Rudd, A. (1982). Approximate option valuation for arbitrary stochastic processes. Journal of Financial Economics, 10, 347–369.

Jarrow, R., & Rudd, A. (1983). Tests of an approximate option-valuation formula. In M. Brenner (Ed.), Option pricing theory and applications (pp. 81–100). Lexington, MA: Lexington Books.

Jondeau, E., & Rockinger, M. (2000, December). Reading the smile: The message conveyed by methods which infer risk neutral densities. Journal of International Money and Finance, 19(6), 885–915.

Kumar, M. (1992). The forecasting accuracy of crude oil futures prices. International Monetary Fund Staff Papers, 39, 432–464.

Longstaff, F. (1995). Option pricing and the martingale restriction. Review of Financial Studies, 8, 1091–1124.

Melick, W. R. (1999). Results of the estimation of implied PDFs from a common data set. Proceedings of the Bank for International Settlements Workshop on Estimating and Interpreting Probability Density Functions, Basel, Switzerland.

Melick, W. R., & Thomas, C. P. (1996). Using option prices to infer PDFs for asset prices: An application to oil prices during the gulf war crisis (International Finance Discussion Paper No. 541). Board of Governors of the Federal Reserve System, Washington, DC.

Melick, W. R., & 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.

Melick, W. R., & Thomas, C. P. (1998). Confidence intervals and constant-maturity series for probability measures extracted from option prices. Bank of Canada Conference on Information in Financial Asset Prices.

Merton, R. (1973, Spring). Theory of rational option pricing. Bell Journal of Economics and Management Science, 4, 141–193.

Merton, R. (1976). Option pricing when underlying stock returns are discontinuous. Journal of Financial Economics, 3, 125–144.

Pearson, E. S., & Tukey, J. W. (1965). Approximate means and standard deviations based on distances between percentage points of frequency curves. Biometrica, 52, 533.

Schittowski, K. (1985). NLQPL: A FOTRAN-subroutine solving constrained nonlinear programming problems. Annals of Operations Research, 5, 485–500.

Copyright of Journal of Futures Markets is the property of John Wiley & Sons, Inc. / Business and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.