Board of Governors of the Federal Reserve System

International Finance Discussion Papers

Number 541

February 1996

USING OPTIONS PRICES TO INFER PDF'S FOR ASSET PRICES: AN APPLICATION TO OIL PRICES DURING THE GULF CRISIS

William R. Melick and Charles P. Thomas

ABSTRACT

We develop a general method to infer martingale equivalent probability density functions (PDFs) for asset prices using American options prices. The early exercise feature of American options precludes expressing the option price in terms of the PDF of the price of the underlying asset. We derive tight bounds for the option price in terms of the PDF and demonstrate how these bounds, together with observed option prices, can be used to estimate the parameters of the PDF. We infer the distribution for the price of crude oil during the Persian Gulf crisis and find the distribution differs significantly from that recovered using standard techniques.

Using Options Prices to Infer PDF's for Asset Prices: An Application to Oil Prices During the Gulf Crisis

William R. Melick and Charles P. Thomas*

I. Introduction

Option prices contain information about market participants' perceptions of the distribution of the price of the underlying asset. To recover this information, analysts typically assume a stochastic process for the price of the underlying asset, such as Brownian motion; use observed option prices to recover the parameters of the assumed process; and then integrate the process to recover the distribution of the price of the underlying asset. However, in certain instances, it is more natural to begin with an assumption about the distribution of the price of the underlying asset, rather than the stochastic process by which it evolves, and use option prices to directly recover the parameters of that distribution.[1] This paper develops a method to directly estimate such a distribution from American options and applies it to the crude oil market during the Persian Gulf crisis. We also compare our estimated distributions to those recovered by more standard methods.

We find that the estimated distributions are consistent with the market commentary at the time, in that they imply a significant probability of a major disruption in the oil markets. We also find that if policy makers or analysts had used the standard Black-Scholes model, they would have generally overestimated the market's assessment of the probability of a major disruption and underestimated the impact on prices of such a disruption.

---

* Correspondence should be directed to C. Thomas, Mail Stop #42, Federal Reserve Board, Washington DC 20551; Tel: (202) 452-3698; Email: thomasc@frb.gov. The authors are staff economists in the Division of International Finance, Board of Governors of the Federal Reserve System. This paper represents the views of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or other members of its staff. We would like to especially thank David Bates, Jon Faust, Greg Duffee, Christian Gilles, and Ed Green for their extensive comments. We also thank Allan Brunner, Neil Ericsson, Jeff Fuhrer, Ludger Hentschel, and George Moore, as well as seminar participants at Vanderbilt's Owen School of Management, Ohio State University, the American Finance Association winter 1995 meetings, the Bank for International Settlements, the University of Neuchatel, the Deutsche Bundesbank, the European Monetary Institute, and the Federal Reserve Board International Finance Division. Elizabeth Vrankovich and Dara Akbarian provided valuable research assistance.

[1] Other studies which focus on the asset's terminal pdf include Breeden and Litzenberger (1978), Jarrow and Rudd (1982), Shimko (1991), and Malz (1996).

Recovering information from option prices is complicated by two factors. First, option prices incorporate preferences towards risk as well as beliefs about outcomes. Short of modeling these preferences or assuming that oil prices are unrelated to other determinants of investor wealth, and thus that oil-price risk is unpriced, the estimated parameters of the stochastic processes or the implicit distributions can only represent the risk-neutral (martingale equivalent) parameters rather than the true (actuarial) parameters.[2] Second, the early exercise feature of American options makes it difficult to derive a closed form expression for the price of an option. To surmount this problem, we derive bounds on the price of an American option in terms of the risk-neutral distribution of the price of the underlying asset, allowing us to proceed with the estimation of the parameters of the assumed distribution.

Our methodology requires us to place some structure on the form of the implied distribution. This structure is similar in spirit to the assumption of a particular stochastic process made in other studies (for example, the jump-diffusion assumption of Bates (1990)).[3] Given the wide range of possible outcomes during the Persian Gulf crisis, we assume that market participants expected oil prices to be drawn from a mixture of three lognormal distributions. This assumption is discussed more fully in section IV.

Placing structure on the terminal distribution rather than the stochastic process has both costs and benefits. With regard to costs, the recovered distribution is silent about the evolution of the asset price prior to expiration. This means that the technique will provide no guidance for constructing dynamic hedges or replication strategies for the option. In addition, the technique does not allow the

[2] Evidence that the risk-neutral and true parameters are very different would include a finding that the oil futures price did not follow a martingale process. Available evidence (see Dominguez (1989), Kumar (1992), and Deaves and Krinsky (1992)) indicates that the oil futures price does in fact martingale. This evidence at least allows for the possibility that our recovered parameters are quite close to the actuarial parameters.

[3] To our knowledge, the only attempt to estimate these distributions non-parametrically is that of Aït-Sahalia and Lo (1995) which uses intra-day quotes for European Options.

time series properties of the underlying asset to be used in ex-post evaluation of the model.[4]

The benefits of the technique arise from its flexibility, generality, and directness. A reasonably flexible functional form for the terminal distribution, such as that used here, can easily accommodate a wide variety of shapes for the terminal distribution. Manipulating equally flexible processes can be quite difficult. More importantly, starting with the terminal distribution is a more general approach since a given terminal distribution encompasses many stochastic process, where a given process is consistent with only one terminal distribution.[5] Finally, as shown by our bounds, for a given terminal distribution over the relevant horizons, the prices of American options are determined almost entirely by the terminal distribution regardless of the stochastic process generating that terminal distribution. Thus, most of the information in the American option prices pertains to the terminal distribution rather than the particular process. A feature of our technique is that it allows the options data to speak directly to the shape of the distribution rather than forcing them to pass through the potentially distortionary filter of a misspecified process. This feature is particularly important in situations, such as the Persian Gulf Crisis, where interest is more naturally focussed on possible asset price outcomes rather than on the asset price process.

The rest of the paper is organized as follows: Section II presents bounds for the price of an American option on a futures contract conditional on the distribution of the futures price at the option's expiration. Section III illustrates how the bounds can be used with option prices to recover the distribution for futures prices. Section IV discusses the particulars of an application to the oil market, while Section V presents the results of that application. A summary and concluding remarks are found in Section VI.

---

[4] However, the method as a whole, rather than a particular estimate, can be evaluated ex-post using EDF tests. See Fackler and King (1990) and Silva and Kahl (1993) for examples.

[5] For example, Rubinstein (1994) begins with an estimate of the risk neutral distribution at expiration, using Breeden and Litzenberger (1978), before constructing, with binomial trees, one process consistent with that distribution.

II. Bounds on American Options' Prices

With European style options the relationship between the distribution of futures prices and the option price is very direct. For calls (puts), the value of the option is simply the value of the portion of the distribution above (below) the strike discounted back to the present using an appropriate interest rate.[6] For American style options the relationship between the distribution and the option price is less direct owing to the early exercise premium. In general, the option's value will depend on the entire stochastic process for futures prices, not just the distribution for futures prices at the option's expiration. To deal with this early exercise premium we develop bounds for the maximum and minimum value of an option given that the futures price is taken from a particular distribution at the option's expiration. That is, for all stochastic processes that imply a given distribution for the futures at the option's expiration, there are bounds for the option's value which can be expressed in terms of that distribution alone.[7] In the estimation routine these bounds are weighted to arrive at a predicted value for the option.

The lower bounds are well known and stated below without proof. The upper bounds, however, are new. Some intuition behind them is given in the text and a proof is given in Appendix I. To fix notation, let f_0 denote the (random) price of the underlying asset at the expiration of the option and let X denote the option's strike price. We index time by periods prior to the option's expiration and denote a period's length by δ_t. The total time to expiration is T. The risk-free interest rate, r, is assumed constant and the one-period discount factor in period t is $e^{-r\delta_t}$, while the discount factor to the option's expiration is e^{-rT}. $E_t[]$ denotes expectations taken t periods prior to expiration. The bounds are predicated on the assumption that the underlying asset martingales with respect to the

[6] See Cox and Ross (1976) for a discussion of the risk-neutral valuation technique.

[7] There are many bounds in the option literature, but to our knowledge this is the first set of bounds on the price of American options in terms of the risk-neutral terminal distribution. Perrakis and Ryan (1984) and Levy (1985) derive bounds for European options in discrete time in terms of the true distribution. Grundy (1991) derives bounds for the true distribution from the moments of the risk-neutral distribution. Lo (1987) derives bounds on European options in terms of the first two moments of the risk-neutral distribution.

probabilities used in the $E_t[]$. The upper and lower bounds for American calls and puts are given as follows:

$$C_t^u = \max\left[E_t[f_0] - X,\ e^{-r\delta_t} \cdot E_t\left[\max[0, f_0 - X]\right]\right]; \tag{1}$$

$$C_t^l = \max\left[E_t[f_0] - X,\ e^{-rT} \cdot E_t\left[\max[0, f_0 - X]\right]\right]; \tag{2}$$

$$P_t^u = \max\left[X - E_t[f_0],\ e^{-r\delta_t} \cdot E_t\left[\max[0, X - f_0]\right]\right]; \tag{3}$$

$$P_t^l = \max\left[X - E_t[f_0],\ e^{-rT} \cdot E_t\left[\max[0, X - f_0]\right]\right]; \tag{4}$$

The intuition behind these bounds is straight forward and similar for calls and puts. First, recall that we are constructing bounds on the options prices *conditional on a given distribution for* f_0; moreover, we have assumed that the futures price martingales with respect to the probability measure used in the expectations operator. By the martingale assumption, today's futures price, f_t , equals today's expectation of the futures price at the expiration of the option, $E_t[f_0]$. Since f_t - X is the revenue one could receive today by exercising the call option today, the call option cannot be worth less than $E_t[f_0]$ - X . This is the first item in the max list for both the upper and lower bound. In addition, an American option cannot be worth less than an otherwise identical European option. The second item in the max list for the lower bound is simply the value of the European version of the option.

The upper bound differs from the lower bound only by the discount factor used in the second item in the max list. Instead of discounting the expected value of the option at expiration by the full time to expiration, we discount it by the length of one period. To understand why this is an upper bound, consider the following thought experiment: Take an arbitrary martingale process consistent with a given terminal distribution. The option has some value under this process. Compare this to a second process which is identical to the first except that any uncertainty which was resolved in the last period (instant) under the first process is now resolved in the next-to-last period under the second process. Since the underlying asset does not move during the last period, there is no reason to delay exercising the option after the uncertainty of next-to-last period is resolved--to delay would only postpone the receipt of a certain revenue. Thus the value of the option under the second process is higher than under the first, since it differs only by being discounted by one period less. The experiment can be repeated by starting with the second process and resolving the remaining uncertainty one period sooner and discounting the value by one period less. With each step we have found a process which gives a higher value for the option than at the prior step. In the end we have constructed a martingale process whereby all of the uncertainty is revealed next period. Under this process the options's value is given as in the upper bound.

As shown in the appendix, with continuous trading the upper bound is simply the undiscounted European value of the option. In the estimation we use evening settle prices, implying an overnight carrying period during which there is no trading in the options. For this reason, the upper bounds were implemented with a one day discount factor applied to the European value.[8]

III. Recovering the Distribution

Equations (1) - (4) give bounds for American option prices t periods before expiration in terms of the interest rates, strike prices, and the period t distribution for futures prices in period 0. To recover this distribution from actual option prices we need to construct a point estimate for the option's value from the bounds, impose some structure on the distribution, and clarify exactly what information the option prices contain.

[8] On June 24, 1993, NYMEX began trading on its Access System which allows trading between 5pm and 8am EST. These extended hours may justify using the continuous-trading upper bound for data after June 1993. As a practical matter, however, for interest rates less than 20 percent, the upper bounds for continuous trading and overnight holdings will differ by less than 0.05 percent.

Let $g(f_0 ; \theta_t)$ be a parametric distribution for the period 0 futures price which traders use in forming expectations in period t. The vector θ_t describes this distribution and is what we estimate. Let $\hat{\theta}_t$ represent our estimate of θ_t and let $\hat{E}_t$ denote expectations taken with respect to $g(\cdot ; \hat{\theta}_t)$.

III.1 Point Estimates

To apply standard estimation techniques requires a point estimate for the option price conditional on the estimated distribution. To generate such an estimate we weight the upper and lower bounds computed above. A natural way to interpret where the actual option price falls between the bounds is in terms of how quickly the market expects uncertainty about the future value of the underlying asset to be resolved. If the market expects the uncertainty to be resolved relatively quickly, then the option's value will be close to the upper bound; if traders expect the uncertainty to be resolved later, then the option will be priced nearer the lower bound. The estimated weights capture this expected speed of resolution. Given the nature of the uncertainty about oil prices during the crisis, we designed the weighting scheme so it could capture different speeds of resolution for the upper and lower tails of the distribution. In particular, we chose to estimate two weights. The first weight, w_1, captures the speed of resolution in the lower tail and is used for all call options that are in the money and for all put options that are out of the money, i.e. $(f_t > X)$. The second weight, w_2, captures the speed of resolution in the upper tail and is used for the calls that are out of the money and the puts that are in the money, i.e. $(f_t < X)$.

The actual option prices can be written in terms of estimated parameters of the assumed distribution and an error term as follows:

$$C_t[X] = \hat{w}_{it} \cdot C_t^{u}[X;\hat{\theta}_t] + (1 - \hat{w}_{it}) \cdot C_t^{l}[X;\hat{\theta}_t] + \hat{\varepsilon}_{ct}[X] \qquad \textbf{(5)}$$

$$P_t[X] = \hat{w}_{it} \cdot P_t^{u}[X;\hat{\theta}_t] + (1 - \hat{w}_{it}) \cdot P_t^{l}[X;\hat{\theta}_t] + \hat{\varepsilon}_{pt}[X] \qquad \textbf{(6)}$$

where $i = 1$ if $\hat{E}_t\,[f_0] > X$ and $i = 2$ otherwise.

The error, $\hat{\varepsilon}$, will be the result of any error in estimating the weights or the parameters of the distribution plus any noise in the system, two examples of which are immediately obvious. First, as the same weights are applied across all options for a given contract/day, there will be a pricing error induced by weighting the two bounds. Second, actual option prices are rounded to the nearest penny, also creating errors in the equation.

III.2 Functional Form of Distribution

In choosing a functional form for the estimated distribution we tried to balance flexibility, parsimony, and ease of interpretation. For reasons explained in Section IV, we specify that the futures price at the option's expiration is drawn from a mixture of three lognormal distributions (hereafter MLN).[9] More formally, the distribution function for futures prices, g[], is given by

$$g[f_0] = \pi_1 g_1[f_0] + \pi_2 g_2[f_0] + \pi_3 g_3[f_0] \tag{7}$$

where

$$g_i[f_0] = \left(\frac{1}{\sqrt{2\pi}\,\sigma_i \cdot f_0}\right) \cdot \exp\left[\left(\frac{\ln(f_0) - \mu_i}{\sigma_i}\right)^2 / 2\right] \tag{8}$$

Using equations (5) and (6), a pricing equation can be written for any option in terms of eleven parameters $((\pi_i, \mu_i, \sigma_i, w_1, w_2)\ i = 1,2,3)$ and three observables $(X, e^{-r\delta_t}, e^{-rT})$.

The parameters of the model, exemplified by equations (5) and (6), are estimated by minimizing the sum of squared errors for all options on a given contract/day, imposing the following constraints:

---

[9] This mixture assumption was also used by Ritchey (1990) who derives European options prices when the pdf is a mixture of lognormals.

$$\sum_{i=1}^{3} \pi_i = 1; \quad 0 \leq w_i \leq 1, \; i = 1, 2. \tag{9}$$

The restriction on the sum of the πs reduces the number of parameters from eleven to ten. An additional restriction could be imposed by noting that under the martingale assumption $E_t[f_0] = f_t$. However, as discussed in Section V below, price limits on the futures were binding for about eighty contract/days during the period. On these days, the futures settle price did not reflect its expectation and this parameter restriction would not be valid. In addition, by not imposing this restriction, we were able to use the futures price as a measure of the goodness-of-fit for the estimated distribution. Details on the data and estimation are found in Section IV.

In the previous section, we claimed that for reasonable discount factors, the bounds are quite close, so a weighted average of them is a good approximation to the option's value. Chart One gives a feel for how close together the bounds are using the estimated distribution for a typical day. The two panels plot the distance between the bounds both in absolute terms and as a percent of the actual option price.

For the day plotted, the options had 38 days to expiration and the relevant T-bill rate was about 7 percent. Thus the discount factors in the formulae for the bounds differed by about 0.8 percent. On this day the futures price was about \$29. For the deeply in-the-money calls (with strikes between \$17 and \$20), the lower bound was determined by the value of exercising today and the difference between the bounds was less than 0.8 percent. For the calls with strikes at or above \$20, the difference between the bounds was about 0.8 percent of the option price. For the puts, none was sufficiently in the money for the lower bound to be determined by the value of exercising today and the bounds differ from each other by about 0.8 percent. For the very low priced puts (with strikes below \$20), the estimated upper bound was below the actual option price, thus the width of the bounds falls to less than 0.8 percent of the actual option price.

III.3 Standard Model as Benchmark

In order to gauge the results from our model (MLN), the distribution for futures prices was also recovered using a "standard" option pricing model. The most common assumption in the option pricing literature is that the underlying commodity price follows a geometric Brownian process, which implies that the futures price at expiration will be drawn from a single lognormal distribution. This is the assumption behind Black's (1976) model for pricing European options on futures. Several approximations have been developed to price American options under this "standard" assumption, with the quadratic approximation of Barone-Adesi and Whaley (1987) (hereafter BAW) being the most easily calculated and most commonly used.[10] We constructed a "standard" model by assuming that prices will be drawn from a single lognormal distribution and used the BAW approximation to generate option pricing equations. We recovered the 2 parameters (μ_b and σ_b) of the single lognormal (hereafter SLN) distribution by minimizing the sum of squared deviations of predicted from actual option prices. (See Appendix II for details of the BAW approximation and SLN estimation.)

Thus, there are two differences between the SLN and MLN models. First, the models account for the early exercise premium in different ways-- SLN uses the BAW quadratic approximation to price an American option, while MLN uses the weighted upper and lower bounds for an American option to price an American option. Second, SLN assumes that the futures price at expiration will be drawn from a single lognormal distribution, while MLN assumes that the futures price at expiration will be drawn from a mixture of lognormal distributions. SLN is almost nested within MLN, except that SLN uses a different technique to account for the early exercise premium. This non-nesting will become important when the two models are statistically compared.

[10] Overdahl and Matthews (1988), when studying a more tranquil period in the oil market, also recovered the parameters of a single lognormal, that is the parameters from the standard Black-Scholes option pricing model.

III.4 Data Limitations vis-a-vis the Distribution

Before proceeding to the estimation, it is useful to note how data limitations and the assumed functional form for the distribution interact. The part of the bounds formulae that shed light on the higher moments of the distribution are the terms of the form $E_t[\max[0, f_0 - X]]$ and $E_t[\max[0, X - f_0]]$. These two terms can be written as $(E_t[f_0 \mid f_0 \geq X] - X) \cdot \mathrm{Prob}_t[f_0 \geq X]$ and $(X - E_t[f_0 \mid f_0 < X]) \cdot \mathrm{Prob}_t[f_0 < X]$. It is clear that even if there were no errors in the pricing relations, the fact that strikes are at discrete intervals and, more importantly, that they do not span the entire support of futures prices places an important limitation on what the option prices can reveal about the distribution. The recorded option prices only contain information about the conditional expectation and probability mass in the following segments of the support: 1) the segment below the lowest strike, 2) the segments between each strike, and 3) the segment above the highest strike. In particular, if X_L and X_H are the lowest and highest strikes, then all the information revealed by the options will be in terms of the following:

$$E_t[f_0|f_0<X_L] \cdot \mathrm{Prob}_t[f_0<X_L] \tag{10}$$

$$E_t[f_0|X_i<f_0<X_j] \cdot \mathrm{Prob}_t[X_i<f_0<X_j] \quad X_L \leq X_i < X_j \leq X_H \tag{11}$$

$$E_t[f_0|f_0\geq X_H] \cdot \mathrm{Prob}_t[f_0\geq X_H] \tag{12}$$

Any number of distributions could generate the same results for the conditional expectations and probabilities in (10)-(12). For example, for any given distribution we can construct a second distribution out of a series of non-overlapping uniform densities which will be observationally equivalent to the given distribution relative to the data described by (10)-(12).

Thus, it is clear that any estimated distribution requires careful interpretation, especially in the regions below the lowest strike and above the highest strike. For crude oil, strikes are almost always

$1.00 apart (in a few instances $5.00), allowing a fine demarkation of the distribution within the range of strikes. In the tails beyond the strikes, however, we have information only on the conditional expectations and the probabilities. Thus the shape of the distribution in the tails will depend importantly on the functional form assumed for the distribution. Chart Two illustrates this point with three observationally equivalent distributions. The solid line is a mixture of three lognormals, while the dashed lines replace the upper tail with uniform densities that yield the same results for (12).

IV. Application to the Oil Market

IV.1 Data Sources

Data on settle prices for all crude oil options on futures for all trading days over the period July 2, 1990 through March 30, 1991 were purchased from NYMEX. We used the settle price as the value of the option in equations (5) and (6). The settle price is determined at the end of each day by a settlement committee made up of roughly 20 options 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 market has settled. In the event of a limit move on the futures market[11], the settlement committee relies on options on the unconstrained spot or nearby contract and spread trading.[12] Using settle prices avoids the problems associated with asynchronous quotes inherent in transaction data.

During July 1990 through March of 1991, trading was concentrated in seven contracts. For each contract/day, all option prices that were recorded with no open interest, no volume, and no

[11] There were no price limits in the options market. Over the entire sample there were limits on crude oil futures price changes for all contracts except for the one closest to expiration. In December of 1990 the limits on crude oil futures price movements were widened substantially.

[12] We are grateful to NYMEX Board of Directors member Jim Zamora of ZAHR Trading and former NYMEX employee Brad Horne for their descriptions of the settlement prices.

settlements were excluded from the data set. In addition, trading days within five working days of the contract's expiration were also excluded from the data set.[13] Table 1 lists summary information for each of the contracts after the exclusions.

Table 1									
Contract	Estimation Range	Total Days	Total Options	Number of Options per Day		Range of Strikes per Contract ($)		Range of Futures Prices per Contract ($)	
				Min	Max	Min	Max	Min	Max
Oct 90	7/2/90 - 8/29/90	41	1254	17	42	14	37	17.74	31.93
Nov 90	7/2/90 - 10/4/90	66	2335	15	57	15	45	18.11	40.42
Dec 90	7/2/90 - 11/1/90	84	3150	13	55	15	44	18.38	38.80
Jan 91	7/2/90 - 11/29/90	104	3889	13	48	16	42	18.55	37.30
Feb 91	8/2/90 - 1/3/91	104	3866	11	43	5	51	23.27	35.95
Mar 91	9/10/90 - 1/31/91	99	3588	12	51	10	50	18.99	34.53
Apr 91	8/1/90 - 2/28/91	144	4827	11	44	10	45	17.91	33.21

Daily prices for the seven Treasury bills that matured as close as possible after the options contracts expired were used to calculate the discount factors. For each contract/day there are N (# of options) equations like (5) and (6) that form a constrained, nonlinear minimization problem. Among the seven contracts there are a total of 642 trading days; each trading/contract day was treated separately, therefore, 642 minimizations were performed. Each day yielded two sets of parameter estimates, the set of ten parameters from MLN and the set of two parameters from SLN.

[13] One day's worth of data for the December contract was also excluded due to an obvious error in data entry on the part of NYMEX.

IV.2 Estimation

Throughout the Persian Gulf crisis, market commentary focused on three distinct outcomes: 1) a return to pre-Crisis conditions (e.g. Iraq would peacefully withdraw from Kuwait), 2) a severe disruption to Persian Gulf oil supplies (e.g. damage to Saudi Arabian facilities during a war), and 3) a continuation of unsettled conditions over the relevant horizon (e.g. a prolonged stalemate in which outcome 1 or 2 might eventually occur). Given these three possibilities, we chose a mixture of three lognormals as the form of the distribution to be estimated. If in fact market participants felt that prices were likely to be drawn from a tri-modal distribution, this could be easily captured by the mixture. Moreover, the mixture could also easily accommodate a single lognormal distribution if that would best fit the data (e.g. $\pi_1 = \pi_2 = 0$). Ex-ante, we expected that as news hit the market, the relative weighting of the three lognormals might change, as well as the parameters of each of the three lognormals. For example, news of an Iraqi rocket attack on a Saudi Arabian oil field might increase the weighting on the lognormal distribution with the highest mode, as well as increase the relevant range encompassed by this lognormal distribution. Section V presents estimated distributions for selected events during the Persian Gulf crisis.

Estimation of MLN and SLN was performed with the Numerical Algorithms Group (NAG) FORTRAN algorithm E04UPF on an IBM RS-6000. Bounds for the parameters were set so that $0 < \hat{\mu}_i < \infty$, $0.0001 < \hat{\sigma}_i < \infty$.[14] Analytic derivatives were provided for both estimations. The derivatives were calculated using Mathematica and they were numerically verified within the E04UPF algorithm prior to estimation. (Details of the estimated equations are relegated to appendix II.)

The estimation procedures are illustrated in Chart Three. The top panel plots the estimated density function using both the MLN and SLN models. Given the density from the single lognormal, the BAW formulae give predicted values for the option prices. The triangles in the lower panels plot

[14] In addition, each (μ, σ) pair was restricted such that the probability of the futures price reaching \$150 per barrel was less than 5 percent, under each of the lognormal distributions. These bounds prevented the algorithm from taking unreasonable first steps.

the difference between these SLN predictions and the actual prices.

Given the MLN estimated distribution, we can compute our upper and lower bounds for the option prices. The dashed (dotted) line in the lower panels plots the difference between the upper (lower) bound for the option price and the actual option price. The predicted option price is a weighted average of these bounds, where (as noted above) the weights are determined in the minimization routine. The boxes plot the difference between the option prices predicted from the MLN distribution and the actual prices. We note that this same MLN distribution was used to draw the plots in Chart One.[15]

V. Results

V.1 Summary Measures

The means of the estimated distributions from MLN and SLN were very similar and were extremely close to the actual futures price. (The actual futures price can be viewed as an independent estimate of the mean of the distribution.) The percentage mean absolute difference (PMAD) between the mean from MLN and the mean from SLN was 0.1 percent. The PMAD between the mean from MLN and the actual futures price across all contract/days amounted to 0.45 percent, while the PMAD between the mean from SLN and the actual futures price amounted to 0.51 percent. Even though they are small, these latter measures significantly overstate the difference between the estimated mean and the true mean, because they include days for which the futures price was subject to price limits. This can be seen in Chart Four which plots the difference between the MLN estimated mean and the futures price against the daily movement in the futures price. Note that the contract/days for which there was

---

[15] The bounds and errors plotted in Charts One and Two differ slightly from those used by the minimization routine for the following technical reason: The minimization routine behaves significantly better if the objective function is differentiable and if analytic derivatives are supplied. The derived formulae for the bounds are not differentiable since they include the max operator. In estimation, we used a logit weighting scheme to construct a differentiable approximation to the max operator where the weights on the two items in the max move to zero and one as the items more farther apart. The data plotted in the charts were constructed using the actual formulae for the bounds, rather than the differentiable approximation, together with the estimated distribution and weights.

a substantial discrepancy between MLN's mean and the futures price (points off the horizontal line through zero) were the contract/days for which actual futures prices moved exactly \$1.00, \$1.50, \$2.00, \$3.00, or \$4.00, that is, contract/days for which there was a limit move on the futures contract. As discussed above, there were no limits in the options market, hence the mean from MLN on these days should not be expected to equal the futures price.

The similarity between the MLN and SLN distributions does not carry over to higher order moments. Chart Five depicts representative probability density functions taken from the two models; the top panel uses estimates from the October contract on July 10th (3 weeks before the crisis), and the bottom panel uses estimates from the January contract on October 10th (in the midst of the crisis). Prior to the outbreak of the crisis there is little qualitative difference in the two estimates, while during the crisis the estimates from SLN cannot as easily accommodate the significant probability mass above \$50 per barrel without over-weighting the distribution between \$40 and \$50 per barrel.

Chart Six attempts to shed some light on the differences between the right hand tails of the two estimates, using the April contract as an example. The top panel plots $1.25 \cdot f_t$ along with $E_t[f_0 | f_0 > 1.25 \cdot f_t]$ from MLN and SLN. The bottom panel plots $\text{Prob}_t [f_0 > 1.25 \cdot f_t]$ from the two models. As can be seen from the chart, the conditional expectation from MLN is generally above that of SLN, while the probability that the futures price will rise by 25 percent is generally lower in MLN than in SLN. The reason for this result is visually apparent in the bottom panel of Chart Five. The large σ_b estimated via SLN forces relatively more of the probability mass to the right but, since the distribution must remain unimodal, leaves the bulk of the right-hand mass nearer the unconditional mean. These results hold across all the contracts. For 574 out of the 642 contract/days the conditional expectation from MLN is above that of SLN. For 476 out of the 642 contract/days the probability of being above $1.25 \cdot f_t$ from MLN is below that of SLN. To make this more concrete: If a policy maker or analyst were using the SLN estimates when the MLN were closer to the truth, she would tend to overestimate the market's assessment of the probability of a major disruption while underestimating the impact on prices of such a disruption.

These differences in the right-hand tails of the distributions are also apparent when examining the pricing errors generated by SLN and MLN. The right-hand tail of the distribution will be more important for pricing out-of-the-money calls and in-the-money puts. As shown in Table Two, for out-of-the-money calls, across all contract days, SLN had a mean error (actual - predicted) of \$0.0865 compared to \$0.0005 for MLN. For in-the-money puts SLN had a mean error of \$0.0445 compared to -\$0.0004 for MLN. For these options, SLN, on average, underpredicted the prices, again indicating that SLN did not allocate enough probability mass to the right-hand tail. As might be expected, SLN tended to overpredict the prices for in-the-money calls (mean error of -\$0.0430) and out-of-the-money puts (mean error of - \$0.0388), an overallocation of probability mass to the left-hand tail of the distribution.

Table 2

Mean Pricing Errors
(Actual - Predicted, \$)

	In the Money		Out of the Money	
	Calls ($X<f_t$)	Puts ($X>f_t$)	Calls ($X>f_t$)	Puts ($X<f_t$)
SLN	-0.0430	0.0445	0.0865	-0.0388
MLN	-0.0001	-0.0004	0.0005	0.0013

V.2 Statistical Model Comparison

Although the differences between the estimates from MLN and those from SLN are apparent, it may be the case that these differences are not significant in a statistical sense. This issue is complicated since the SLN model cannot be nested within MLN.[16] Since the models are not nested,

[16] The nesting issue is as follows: We have two competing non-linear models that explain a vector of option prices (y) on any given day. Denote them by MLN: $y = g[\pi_1, \pi_2, \pi_3, \mu_1, \mu_2, \mu_3, \sigma_1, \sigma_2, \sigma_3, w_1, w_2 \mid Z]$ and SLN: $y = h[\mu, \sigma \mid Z]$, where Z is a data matrix containing strikes and interest rates, g[] involves the weighted bounds and h[] involves the BAW approximation. If g[] and h[] did not represent different functional forms, SLN could be nested within MLN with the restrictions ($\pi_1 = \pi_2 = 0$, or $\mu_1 = \mu_2 = \mu_3$ and $\sigma_1 = \sigma_2 = \sigma_3$).

the asymptotic, chi-square assumption cannot be used when forming a likelihood ratio test (or its F-test analog). The standard tests for formally comparing such non-linear, non-nested models are the J and P (or JA and PA) tests.[17] We use the P test, which compares the two models with the following regression

$$\hat{\varepsilon}_{1,j} = \hat{d}_{1,j}\,\beta + \alpha \cdot (\hat{p}_{2,j} - \hat{p}_{1,j}) + \text{residual}, \tag{13}$$

where $\hat{\varepsilon}_{1,j}$ is the pricing error for option "j" from model 1; $\hat{d}_{1,j}$ are the derivatives of model 1 with respect to the parameters of model 1 for option "j" evaluated at the estimates for the Model 1 parameters; $\hat{p}_{1,j}$ and $\hat{p}_{2,j}$ are the fitted values for option "j" from the two models; β is a vector of coefficients with length equal to the number of parameters in model 1; and α is a single coefficient. The P test is simply the t-statistic for α. Intuitively, the difference in the fitted values from the two models should not help explain the errors of model 1. Hence, model 1 is rejected in the event of a significant t-statistic for α. Obviously, SLN and MLN can both serve as either model 1 or model 2. Table Three presents the t-statistics for α, after equation (13) has been pooled across options and trading days for each contract. The first two columns of the table treat SLN as model 1 (with degrees of freedom denoted by DF), while the last two columns treat MLN as model 1.

Every t-statistic in the first column is significant at the 5 percent level, while no t-statistic in the fourth column is significant at the 5 percent level. Clearly, the difference in the fitted values from MLN and SLN help to explain the pricing errors of SLN, while the converse is not true. Thus, available evidence (the data and MLN) can reject SLN, but MLN cannot be rejected.

[17] See MacKinnon (1992) for a discussion of these tests.

Table 3

P Test t-statistics

	SLN as Model 1			MLN as Model 1		
Contract	t-stat	DF	P-value	t-stat	DF	P-value
October	133.98	590	0.00	-0.29	581	0.61
November	172.98	1329	0.00	0.27	1320	0.39
December	130.71	1958	0.00	0.83	1949	0.20
January	122.75	2504	0.00	-0.12	2495	0.55
February	164.54	2504	0.00	-1.28	2495	0.90
March	219.26	2242	0.00	-1.08	2233	0.86
April	260.51	3020	0.00	-0.44	3011	0.67

MacKinnon (1992) notes that the P test can have poor finite sample properties, especially when Model 2 has a large number of parameters. For example, in the limiting situation where a completely over parameterized Model 2 exactly matches the observed option prices, the regression in (13) collapses with an infinite t-statistic for α. Given that MLN has 11 parameters and that the average trading day has roughly 35 option prices, there may be reason to question the results of the P test. To ensure that this is not the case, we ran the following Gauss Newton Regression (GNR) described in MacKinnon (1992) where $\tilde{d}_{2,j}$ are the derivatives of model 2 with respect to the parameters of model 2

$$\hat{\varepsilon}_{1,j} = \hat{d}_{1,j}\beta + \tilde{d}_{2,j}\gamma + \text{residual} \tag{14}$$

for option "j" evaluated at a point chosen independently of the data used to estimate model 1. These derivatives will not be correlated with $\hat{\varepsilon}_1$ in small samples, and if Model 1 is correct, they will not help explain the errors from Model 1. Table 4 presents the F-tests for $\gamma = 0$.

Table 4

GNR F-Tests

	SLN as Model 1			MLN as Model 1		
Contract	F-test	DF (n,d)	P-value	F-test	DF (n,d)	P-value
October	94.77	11, 1031	0.00	0.02	2, 1031	0.98
November	197.32	11, 2045	0.00	0.01	2, 2045	0.99
December	275.64	11, 2872	0.00	0.54	2, 2872	0.58
January	264.02	11, 3638	0.00	0.86	2, 3638	0.42
February	297.17	11, 3638	0.00	0.06	2, 3638	0.94
March	428.72	11, 3321	0.00	2.50	2, 3321	0.08
April	462.05	11, 4594	0.00	0.89	2, 4594	0.41

The conclusion is the same as that of the P-test, SLN is rejected while MLN is not.[18]

V.2 Selected Events

Throughout the Persian Gulf crisis, the oil market often experienced large price movements as "news" hit the market and participants revised their expectations concerning likely outcomes to the crisis. Comparing the estimated PDFs from the two models immediately before and after the receipt of "news" allows us to infer how the market interpreted the news and further highlights the differences between the MLN and SLN models.

[18] When SLN is treated as Model 1 in equation (14), the derivatives of MLN were evaluated at $\sigma_1 = \sigma_2 = \sigma_3 = .2$, $w_1 = w_2 = .5$, $\mu_2 = \ln(f_t) - \sigma_1^2/2$, $\mu_1 = .95 \cdot \mu_2$, $\mu_3 = 1.05 \cdot \mu_2$, $\pi_1 = \pi_3 = .1$, and $\pi_2 = .8$. Given that the futures price for a given trading day was not used in the estimation of the parameters for either MLN or SLN, these derivatives, given SLN is correct, will not be correlated with the errors from SLN. When MLN is treated as Model 1 in equation (14), the derivatives of SLN were evaluated at the estimated values of μ and σ from SLN. This is a conservative approach that increases the likelihood of a significant F-statistic.

On Thursday October 25, 1990 the London Financial Times reported that Iraqi forces had attached explosives to 300 of Kuwait's 1000 oil wells, quoting a senior Kuwaiti engineer who had left Kuwait one week earlier. This revelation pushed oil prices up sharply, with the futures contract nearest to expiration (December) rising \$3.17 per barrel. Chart Seven plots the PDFs from MLN and SLN for October 22 (top panel) and October 25 (bottom panel) using the January contract. On October 22, market expectations for futures prices were centered quite tightly around \$24 per barrel. The news of the mining widened each model's distribution significantly, with MLN allowing for a sizeable probability mass between \$60 and \$70 per barrel.

The largest one-day change in oil prices in NYMEX history occurred on Thursday January 17, 1991 when 1) several governments announced a coordinated release of oil from their emergency inventories and 2) it became clear that the coalition forces had total air supremacy. On that day the settle price for the March contract fell \$9.66 while the settle price for the April contract fell \$7.82. The six panels of Chart Eight trace the evolution of expected PDFs on the days surrounding January 17. Prior to the first air strike (as can be seen in the first two panels), the market was still expecting a fairly significant chance of a major oil market disruption (perhaps Iraqi damage of Saudi Arabian oil facilities) that could push prices to the \$40-\$60 per barrel range. On January 17th these PDFs tightened dramatically, and on ensuing days the PDF generated from MLN moved closer and closer to that from SLN. By January 23, there was little difference between the two PDFs, as the market returned to almost a pre-crisis distribution (compare panel 6 with the top panel of Chart 2).

VI Conclusion

This paper develops a method for using option prices to infer the market's probability distribution for commodity prices. The method is quite general, allowing the standard lognormal distribution to be replaced by any from within a wide class of distributions. The particular assumption of a mixture of three lognormal distributions used here was driven by conditions in the oil market during the Persian Gulf crisis. As the focus is only on the probability distribution of the commodity price, minimal structure is placed on the stochastic process governing movements in the commodity price over time. This lack of structure is appealing since we generally have little *a priori* information about the stochastic process that market participants have assumed. Our methodology should be useful to researchers who wish to impose a minimum of structure and are 1) examining other markets during unsettled times, or 2) investigating asset price distributions that are not adequately described by the lognormal distribution (e.g. leptokurtotic distributions). For example, in the foreign exchange market it can be used to shed light on the "Peso problem." Or, in the money market, it can be used to infer the market's assessment of possible changes in monetary policy.[19]

In the application to the oil market we find that the options prices were consistent with the market commentary at the time, in that they reflected a significant probability of a major disruption in oil prices. We find that the estimated price of oil conditional on a major disruption was often in the \$50-\$60 per barrel range, which is also consistent with market commentary. We also find that the standard lognormal assumption did a poorer job of characterizing the data than did our model and the two models have different implications. In particular, compared to our model, the lognormal model implies an overestimate of the market's assessment of the probability of a major disruption and an underestimate of the impact on prices of such a disruption.

Finally, examination of particular days confirmed the large shift in market expectations that occurred when significant crisis-related news reached the oil market.

---

[19] Deutsche Bundesbank (1995) provides examples using related techniques.

References

Aït-Sahalia, Y., and A. W. Lo. "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices." *NBER Working Paper Series*, 5351 (1995).

Barone-Adesi, G., and R. Whaley. "Efficient Analytic Approximation of American Option Values." *Journal of Finance*, 42 (1987), 301-320.

Bates, D. S. "The Crash of '87: Was It Expected? The Evidence From Options Markets." *Journal of Finance,* 46 (July 1990), 1009-1044.

Black, F. "The Pricing of Commodity Contracts." *Journal of Financial Economics*, 3 (1976), 167-179.

Black, F., and M. Scholes. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81 (1976), 637-659.

Breeden, D., and R. Litzenberger. "Prices of State-Contingent Claims Implicit in Options Prices." *Journal of Business*, 51 (1978), 621-651.

Cox, J. C., and S. A. Ross. "The Valuation of Options for Alternative Stochastic Processes." *Journal of Financial Economics*, 3 (1976), 145-166.

Deaves, R., and I. Krinsky. "Risk Premiums and Efficiency in the Market for Crude Oil Futures." *The Energy Journal,* 13 (1992), 93-118.

Deutsche Bundesbank, "The Information Content of Derivatives for Monetary Policy." *Deutsche Bundesbank Monthly Report*, 47 (November 1995), 17-32.

Dominquez, K. "The Volatility and Efficiency of Crude Oil Futures Contracts." In *Oil and Money: Coping With Price Risk Through Financial Markets*, Harvard International Energy Studies, 5 (1989), 48-87.

Fackler, P. L., and R. P. King. "Calibration of Option-Based Probability Assessments in Agricultural Commodity Markets." *American Journal of Agricultural Economics*, 72 (Feb. 1990), 73-83.

Gradshteyn, I.S., and I.M. Ryzhik. *Tables of Integrals, Series, and Products*, Academic Press (1980).

Grundy, B. D. "Option Prices and the Underlying Asset's Return Distribution." *Journal of Finance*, 46 (July 1991), 1045-1069.

Jarrow, R., and A. Rudd. "Approximate Option Valuation for Arbitrary Stochastic Processes." *Journal of Financial Economics*, 10 (1982), 347-369.

Kumar, M. "The Forecasting Accuracy of Crude Oil Futures Prices." IMF Staff Papers, 39 (June 1992).

Lo, A. W. "Semi-parametric Upper Bounds for Option Prices and Expected Payoffs." *Journal of Financial Economics*, 19 (1987), 373-387.

Levy, H. "Upper and Lower Bounds of Put and Call Option Value: Stochastic Dominance Approach." *Journal of Finance*, 40 (1985), 1197-1217.

MacKinnon, J. G. "Model Specification Tests and Artificial Regressions." *Journal of Economic Literature*, 30 (March 1992), 102-146.

Malz, A. M. "Recovering the Probability Distribution of Future Exchange Rates from Option Prices." Federal Reserve Bank of New York, Mimeo (1996).

Melick, W. R. and C. P. Thomas "War and Peace: Recovering the Market's Probability Distribution of Crude Oil Futures Prices During the Gulf Crisis." Board of Governors of the Federal Reserve System International Finance Discussion Papers, No. 437 (October 1992).

Overdahl, J. A., and H. L. Matthews. "The Use of NYMEX Options to Forecast Crude Oil Prices." *The Energy Journal*, 9 (1988), 135-147.

Perrakis, S., and P. J. Ryan. "Option Pricing Bounds in Discrete Time." *Journal of Finance*, 39 (1984), 519-525.

Ritchey, R. J. "Call Option Valuation for Discrete Normal Mixtures." *Journal of Financial Research*, 13 (Winter 1990), 285-296.

Rubinstein, M. "Implied Binomial Trees." *Journal of Finance*, 49 (1994), 771-818.

Shimko, D. C. "Beyond Implied Volatility: Probability Distributions and Hedge Ratios Implied by Option Prices." Mimeo, Department of Finance and Business Economics, University of Southern California, (November 1991).

Silva, E. M. S., and K. H. Kahl. "Reliability of Soybean and Corn Option-Based Probability Assessments." *Journal of Futures Markets*, 13 (1993), 765-799.

Appendix I

Derivation of bounds

Theorem One: If the asset, f_t , underlying an American option is a martingale, then in discreet time an upper bound for the call option's value, in terms of the terminal distribution for the underlying asset, is given as follows: $C_t^u = \max\left[E_t[f_0] - X,\ e^{-r\delta_t} \cdot E_t\left[\max\left[0,\ f_0 - X\right]\right]\right]$, where f_0 is the asset price at expiration of the option; X is the option's strike price; r is the (constant) risk free rate of interest; δ_t is the minimal time interval between trades/exercise at period t; and E_t denotes expectations taken t periods prior to expiry. If trade is taking place continuously at time t, then the upper bound is $C_t^u = E_t\left[\max\left[0,\ f_0 - X\right]\right]$ which is the undiscounted European value. Similarly, an upper bound for the American put option is as follows: $P_t^u = \max\left[X - E_t[f_0],\ e^{-r\delta_t} \cdot E_t\left[\max\left[0,\ X - f_0\right]\right]\right]$.

Proof: Discrete time: Let **J** = $\{\tau_0, \tau_1, \ldots \tau_t \ldots\}$ be a countable set of points in the interval [0, T] representing the remaining life of the option, with $\tau_0 = 0$, and $(i > j) \Rightarrow (\tau_i \geq \tau_j)$. Each element of **J** represents a point at which the option can be traded/exercised. Define $\delta_t \equiv \tau_t - \tau_{t-1}$ and note $\delta_t \geq 0$ $\forall t \geq 1$.

The proof is by induction. Let $C_t^u \equiv \max\left[E_t[f_0] - X,\ e^{-r\delta_t} \cdot E_t\left[\max\left[0,\ f_0 - X\right]\right]\right]$. Step One shows $C_1 = C_1^u$, implying $C_1 \leq C_1^u$. Step Two shows that $\left(C_t \leq C_t^u\right) \Rightarrow \left(C_{t+1} \leq C_{t+1}^u\right)$. Thus by induction $C_t \leq C_t^u$ $\forall t \geq 1$.

We index time by 'periods' before expiry of the option. The option expires in period 0; the period prior to expiry is period 1, etc. Let $y_t \equiv f_t - X$. By assumption, f_t is a martingale, thus $y_t = E_t[y_{t-1}] = E_t[y_0] = E_t[f_0] - X$. Define $m(y) \equiv \max[0, y]$. Since m() is convex, by Jensen's inequality we have $E_t[m(y_0)] \geq m(E_t[y_0])$.

Step One: The value of an American call option at expiry is

$$C_0 = \max\left[0,\ f_0 - X\right] = m(y_0)\ . \tag{A1.1}$$

Its value one period prior to expiry is

$$C_1 = \max\left[f_1 - X,\ e^{-r\delta_1} \cdot E_1[C_0]\right] = \max\left[y_1,\ e^{-r\delta_1} \cdot E_1[m(y_0)]\right] = \max\left[E_1[f_0] - X,\ e^{-r\delta_1} \cdot E_1\left[\max[0,\ f_0 - X]\right]\right] = C_1^u \tag{A1.2}$$

which completes Step One.

Step Two: We want to show that $(C_t \leq C_t^u) \Rightarrow (C_{t+1} \leq C_{t+1}^u)$. The value of an American option can be written recursively as follows:

$$C_{t+1} = \max\left[y_{t+1},\ e^{-r\delta_{t+1}} \cdot E_{t+1}[C_t]\right]\ \forall t \geq 0 \text{ and } C_0 = m(y_0) \tag{A1.3}$$

By assumption, $C_t \leq C_t^u$, thus

$$C_{t+1} \leq \max\left[y_{t+1},\ e^{-r\delta_{t+1}} \cdot E_{t+1}[C_t^u]\right] \tag{A1.4}$$

Substituting for C_t^u,

$$C_{t+1} \leq \max\left[y_{t+1},\ e^{-r\delta_{t+1}} \cdot E_{t+1}\left[\max\left[y_t,\ e^{-r\delta_t} \cdot E_t[m(y_0)]\right]\right]\right]. \tag{A1.5}$$

Since $m(y) \geq y$, replacing y_t with $m(y_t)$ can only raise the RHS, thus

$$C_{t+1} \leq \max\left[y_{t+1},\ e^{-r\delta_{t+1}} \cdot E_{t+1}\left[\max\left[m(y_t),\ e^{-r\delta_t} \cdot E_t[m(y_0)]\right]\right]\right]. \tag{A1.6}$$

Since y martingales, we can replace y_t with $E_t[y_0]$:

$$C_{t+1} \leq \max\left[y_{t+1},\ e^{-r\delta_{t+1}} \cdot E_{t+1}\left[\max\left[m(E_t[y_0]),\ e^{-r\delta_t} \cdot E_t[m(y_0)]\right]\right]\right]. \tag{A1.7}$$

Since $E_t[m(y_0)] \geq m(E_t[y_0])$ replacing $m(E_t[y_0])$ with $E_t[m(y_0)]$ can only raise the RHS, thus

$$C_{t+1} \leq \max\left[y_{t+1},\ e^{-r\delta_{t+1}} \cdot E_{t+1}\left[\max\left[E_t[m(y_0)],\ e^{-r\delta_t} \cdot E_t[m(y_0)]\right]\right]\right]. \tag{A1.8}$$

Since $e^{-r\delta_t} \leq 1$, $e^{-r\delta_t} \cdot E_t[m(y_0)] \leq E_t[m(y_0)]$ and $e^{-r\delta_t} \cdot E_t[m(y_0)]$ can be removed from the interior max without changing its value.

$$C_{t+1} \leq \max\left[y_{t+1},\ e^{-r\delta_{t+1}} \cdot E_{t+1}\left[E_t[m(y_0)]\right]\right] . \tag{A1.9}$$

By iterated expectations we have

$$\begin{gathered} C_{t+1} \leq \max\left[y_{t+1},\ e^{-r\delta_{t+1}} \cdot E_{t+1}[m(y_0)]\right] = \\ \max\left[E_{t+1}[f_0] - X,\ e^{-r\delta_{t+1}} \cdot E_{t+1}\left[\max[0,\ f_0 - X]\right]\right] \equiv C^u_{t+1} \end{gathered} \tag{A1.10}$$

This completes Step Two and thus the proof.

Continuous Trading: Note that

$$E_t[f_0] - X \leq E_t\left[\max[0,\ f_0 - X]\right] . \tag{A1.11}$$

First take the case in which the period t distribution for f_0 has some mass below X, so that the inequality in **(A1.11)** is strict. There is a $\bar{\delta}_t$ s.t. $\forall \delta_t < \bar{\delta}_t$, $e^{-r\delta_t} \cdot E_t\left[\max\left[0,\ f_0 - X\right]\right] > E_t[f_0] - X$, and the upper bound can be written as

$$C^u_t = e^{-r\delta_t} \cdot E_t\left[\max[0,\ f_0 - X]\right] \text{ for all } \delta_t < \bar{\delta}_t \tag{A1.12}$$

Taking the limit as the interval between trades/exercise goes to zero yields:

$$\lim_{\delta_t \to 0} C^u_t = E_t\left[\max[0,\ f_0 - X]\right] \tag{A1.13}$$

which is the undiscounted European value of the option.

If **(A1.11)** holds with equality, then

$$E_t[f_0] - X = E_t\left[\max[0,\ f_0 - X]\right] > e^{-r\delta_t} \cdot E_t\left[\max[0,\ f_0 - X]\right] \tag{A1.14}$$

which again leaves the upper bound equal to the undiscounted European value.

The proof for puts is identical to the above with $y_t \equiv X - f_t$ replacing $y_t \equiv f_t - X$.

Note: Although this proof is direct, it obscures the basic intuition behind the result. Melick and Thomas (1992) provide an alternative, less general, proof which directly mirrors the intuition given in Section II of the text.

Appendix II

Estimation Details

MLN Model

Equations (5) and (6) in the text give the pricing equations for puts and calls in terms of the bounds derived in Appendix I. They are repeated here for the reader's convenience

$$C_t[X] = \hat{w}_{it} \cdot C_t^{u}[X;\hat{\theta}_t] + (1-\hat{w}_{it}) \cdot C_t^{l}[X;\hat{\theta}_t] + \hat{\varepsilon}_{ct}[X] \qquad \textbf{(A2.1)}$$

$$P_t[X] = \hat{w}_{it} \cdot P_t^{u}[X;\hat{\theta}_t] + (1-\hat{w}_{it}) \cdot P_t^{l}[X;\hat{\theta}_t] + \hat{\varepsilon}_{pt}[X] \qquad \textbf{(A2.2)}$$

Substituting for the bounds using equations (1)-(4) and noting that

$E\big[\max[0, z]\big] = E[z \mid z \geq 0] \cdot \Pr[z \geq 0]$ yields pricing equations written in terms of the terminal distribution.

$$\begin{aligned} C_t[X] = {} & \hat{w}_{it} \cdot \max\Big[\hat{E}_t[f_0] - X \,,\ e^{-\tau\delta_t} \cdot (\hat{E}_t[f_0 \mid f_0 \geq X] - X) \cdot \hat{\Pr}[f_0 \geq X]\Big] + \\ & (1-\hat{w}_{it}) \cdot \max\Big[\hat{E}_t[f_0] - X \,,\ e^{-\tau T} \cdot (\hat{E}_t[f_0 \mid f_0 \geq X] - X) \cdot \hat{\Pr}[f_0 \geq X]\Big] + \hat{\varepsilon}_{ct}[X] \end{aligned} \qquad \textbf{(A2.3)}$$

$$\begin{aligned} P_t[X] = {} & w_{it} \cdot \max\Big[X - \hat{E}_t[f_0],\ e^{-\tau\delta_t} \cdot \big(X - \hat{E}_t[f_0 \mid f_0 \leq X]\big) \cdot \hat{\Pr}[f_0 \leq X]\Big] + \\ & (1-w_{it}) \cdot \max\Big[X - E_t[f_0],\ e^{-\tau T} \cdot \big(X - \hat{E}_t[f_0 \mid f_0 \leq X]\big) \cdot \hat{\Pr}[f_0 \leq X]\Big] + \hat{\varepsilon}_{pt}[X] \end{aligned} \qquad \textbf{(A2.4)}$$

It is assumed that the distribution for futures prices (f_0) is a mixture of lognormals i.e.

$$g[f_0] = \pi_1 g_1[f_0] + \pi_2 g_2[f_0] + \pi_3 g_3[f_0] \qquad \textbf{(A2.5)}$$

where

$$g_i[f_0] = \left(\frac{1}{\sqrt{2\pi}\,\sigma_i \cdot f_0}\right) \cdot \exp\left[\left(\frac{\ln(f_0) - \mu_i}{\sigma_i}\right)^2 / 2\right]$$

Using the properties of the lognormal distribution, the terms involving the expectations operator in equations (A2.3) and (A2.4) may be substituted for as follows:

$$E[f_0] = \sum_{i=1}^{3} \pi_i \cdot \exp\left[\mu_i + \frac{\sigma_i^2}{2}\right] \quad \textbf{(A2.6)}$$

$$\Pr[f_0 \geq X] = \sum_{i=1}^{3} \pi_i \cdot \left(1 - \Phi\left[\frac{\ln[X] - \mu_i}{\sigma_i}\right]\right) \quad \textbf{(A2.7)}$$

$$\Pr[f_0 \leq X] = \sum_{i=1}^{3} \pi_i \cdot \Phi\left[\frac{\ln[X] - \mu_i}{\sigma_i}\right] \quad \textbf{(A2.8)}$$

$$E[f_0 \,|\, f_0 \geq X] = \sum_{i=1}^{3} \pi_i \cdot \frac{\exp\left[\frac{\sigma_i^2 + 2\mu_i}{2}\right]\left(\Phi\left[\frac{\ln[X] - \mu_i - \sigma_i^2}{\sigma_i}\right] - \frac{1}{2}\right)}{\Pr[f_0 \geq X]} \quad \textbf{(A2.9)}$$

$$E[f_0 \,|\, f_0 \leq X] = E[f_0] - E[f_0 \,|\, f_0 \geq X] \quad \textbf{(A2.10)}$$

where Φ represents the cumulative normal distribution function.[20] The pricing equations are now written in terms of the parameters of the model. A final difficulty is presented by the non-

[20]The mean of the lognormal distribution is $\exp[u+\sigma^2/2\}$. Calculation of (A2.9) and (A2.10) used integral 3.322 from Gradshteyn and Ryzhik (1980).

differentiability of the max operator in equations (A2.3) and (A2.4). Constrained, nonlinear minimization is greatly enhanced when analytic derivatives are provided. Therefore, we replaced the max operator with a logisitic approximation as follows:

$$\text{logitmax}[x,y] \equiv \frac{1}{1+\exp[-5\cdot(x-y)]} \qquad \textbf{(A2.11)}$$

$$\max[x,y] \approx \text{logitmax}[x,y]\cdot x + (1-\text{logitmax}[x,y])\cdot y \qquad \textbf{(A2.12)}$$

SLN Model

The SLN Model uses the approximation of Barone-Adesi and Whaley (1987) (BAW) to express the option price in terms of the parameters of a single lognormal (u_b and σ_b). In our notation, the BAW approximation for the price of an American call option is written as

$$C_t[X] = c\big[E_t[f_0],\sigma_b,X\big] + A_2[f^*,\sigma_b,X]\left(\frac{E_t[f_0]}{f^*}\right)^{q_2[\sigma_b]} \quad \text{when} \quad E_t[f_0]<f^* \qquad \textbf{(A2.13)}$$

$$C_t[X] = E_t[f_0] - X \quad \text{when} \quad E_t[f_0]\geq f^* \qquad \textbf{(A2.14)}$$

where

$$c[E_t[f_0],\sigma_b,X] \equiv E_t[f_0]\cdot e^{-rT}\cdot\Phi\big[d_1[E_t[f_0],\sigma_b,X]\big] - X\cdot e^{-rT}\cdot\Phi\big[d_2[E_t[f_0],\sigma_b,X]\big]$$

$$A_2[f^*,\sigma_b,X] \equiv \frac{\sigma_b}{q_2[\sigma_b]}\cdot\Big(1-e^{-rT}\cdot\Phi\big[d_1[f^*,\sigma_b,X]\big]\Big)$$

$$q_2[\sigma_b] \equiv \frac{1 + \sqrt{1 + \frac{8\rho_t}{\sigma_b^2 \cdot (1 - \exp[-rT])}}}{2}$$

$$d_1\left[E_t[f_0],\sigma_b,X\right] \equiv \frac{\ln\left[\frac{E_t[f_0]}{X}\right] + \frac{\sigma_b^2 \cdot T}{2}}{\sigma_b \cdot \sqrt{T}}$$

$$d_2\left[E_t[f_0],\sigma_b,X\right] = d_1\left[E_t[f_0],\sigma_b,X\right] - \sigma_b \cdot \sqrt{T}$$

$$E_t[f_0] = \exp\left[\mu_b + \frac{\sigma_b^2}{2}\right]$$

The term f^* is the "critical commodity price," and is solved for implicitly according to

$$f^* - X = c[f^*,\sigma_b,X] + \frac{\left(1 - e^{-rT} \cdot \Phi\left[d_1[f^*,\sigma_b,X]\right]\right) \cdot f^*}{q_2[\sigma_b]} \qquad \textbf{(A2.15)}$$

The formula for the price of a put is similar to that for the call, see BAW for details. NAG algorithm C05AJF was used to solve for f^*. To improve the estimation of μ_b and σ_b, derivatives were provided to the NAG minimization algorithm E04UPF. Given the option pricing formula (equations (A2.13) and (A2.14)) is not differentiable, the logitmax operator described above was used as follows.

$$\begin{aligned} C_t[X] &= \text{logitmax}\left[f^*,E_t[f_0]\right] \cdot (\text{RHS of (A2.13)}) + \\ &\left(1 - \text{logitmax}\left[f^*,E_t[f_0]\right]\right) \cdot (\text{RHS of (A2.14)}) \end{aligned} \qquad \textbf{(A2.16)}$$

Chart 1
Width of Bounds on Option Prices
Upper Bound - Lower Bound

Calls

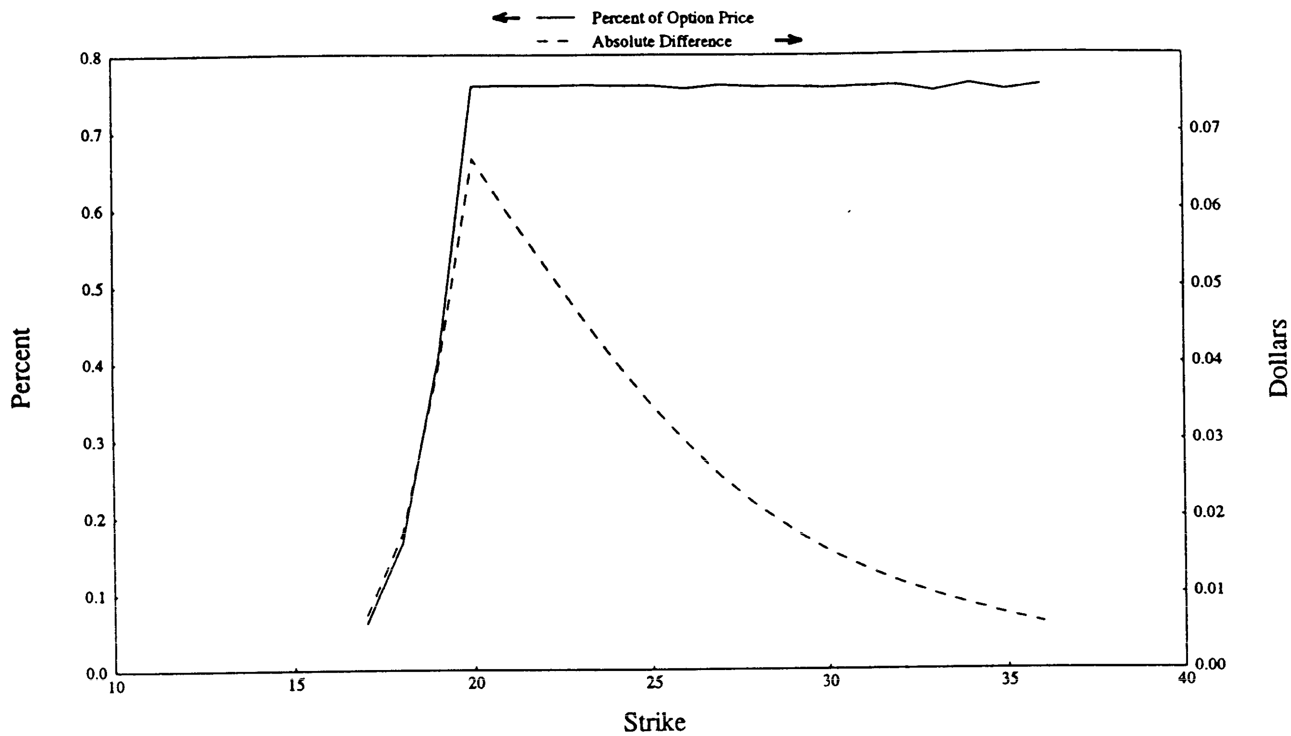

Puts

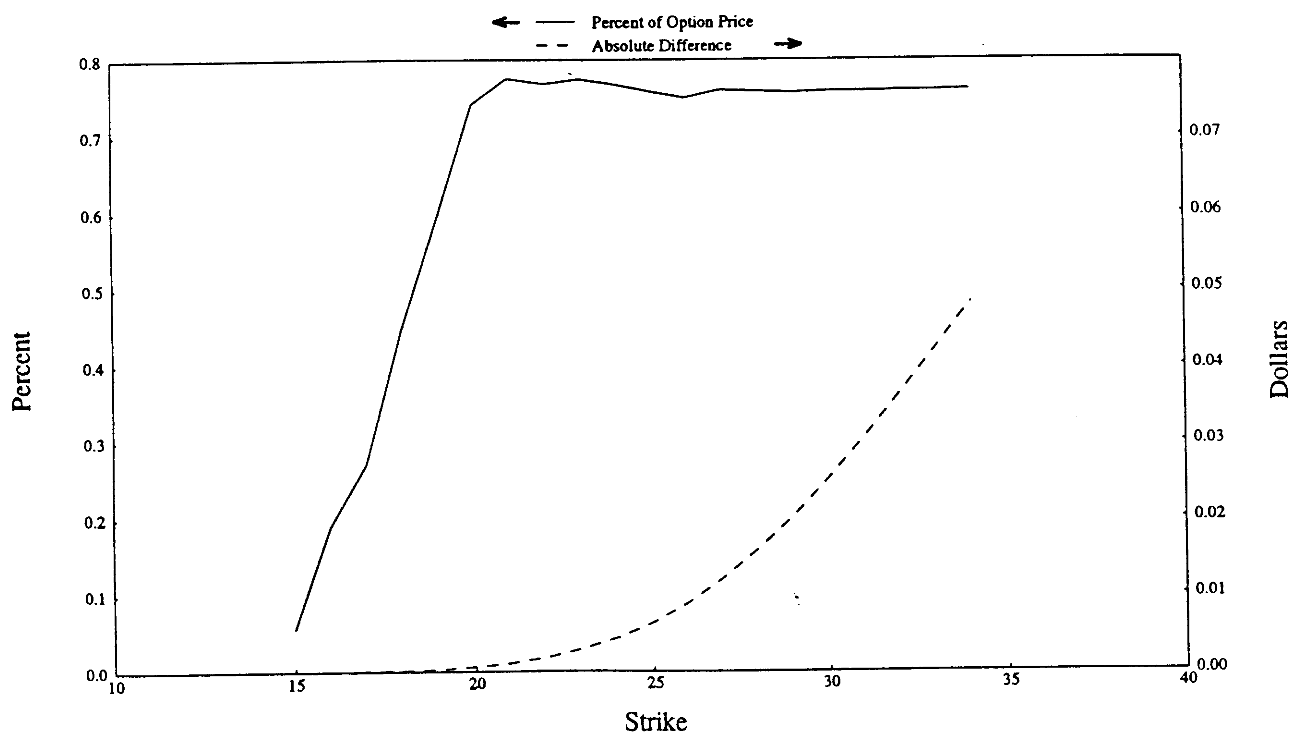

Chart 2

Observationally Equivalent Density Functions

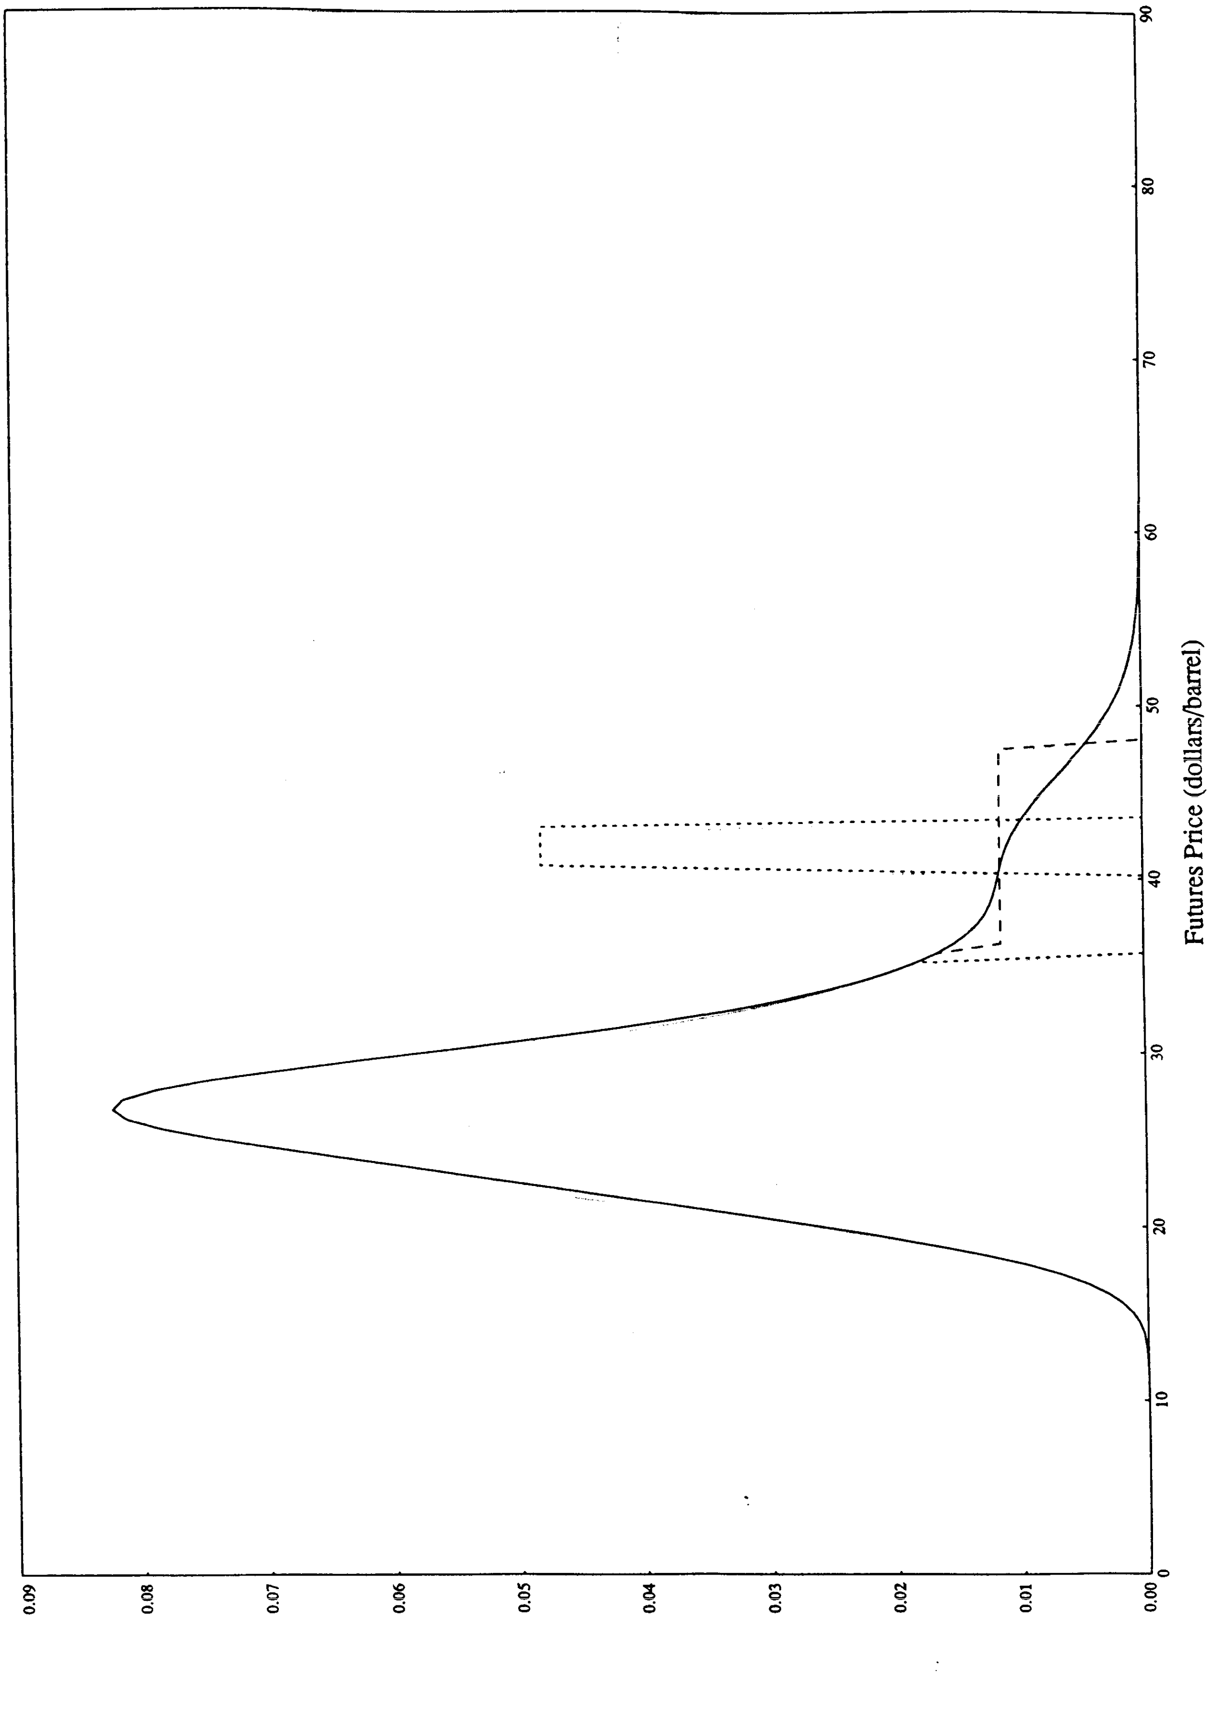

Chart 3
Representative Results for a Single Day

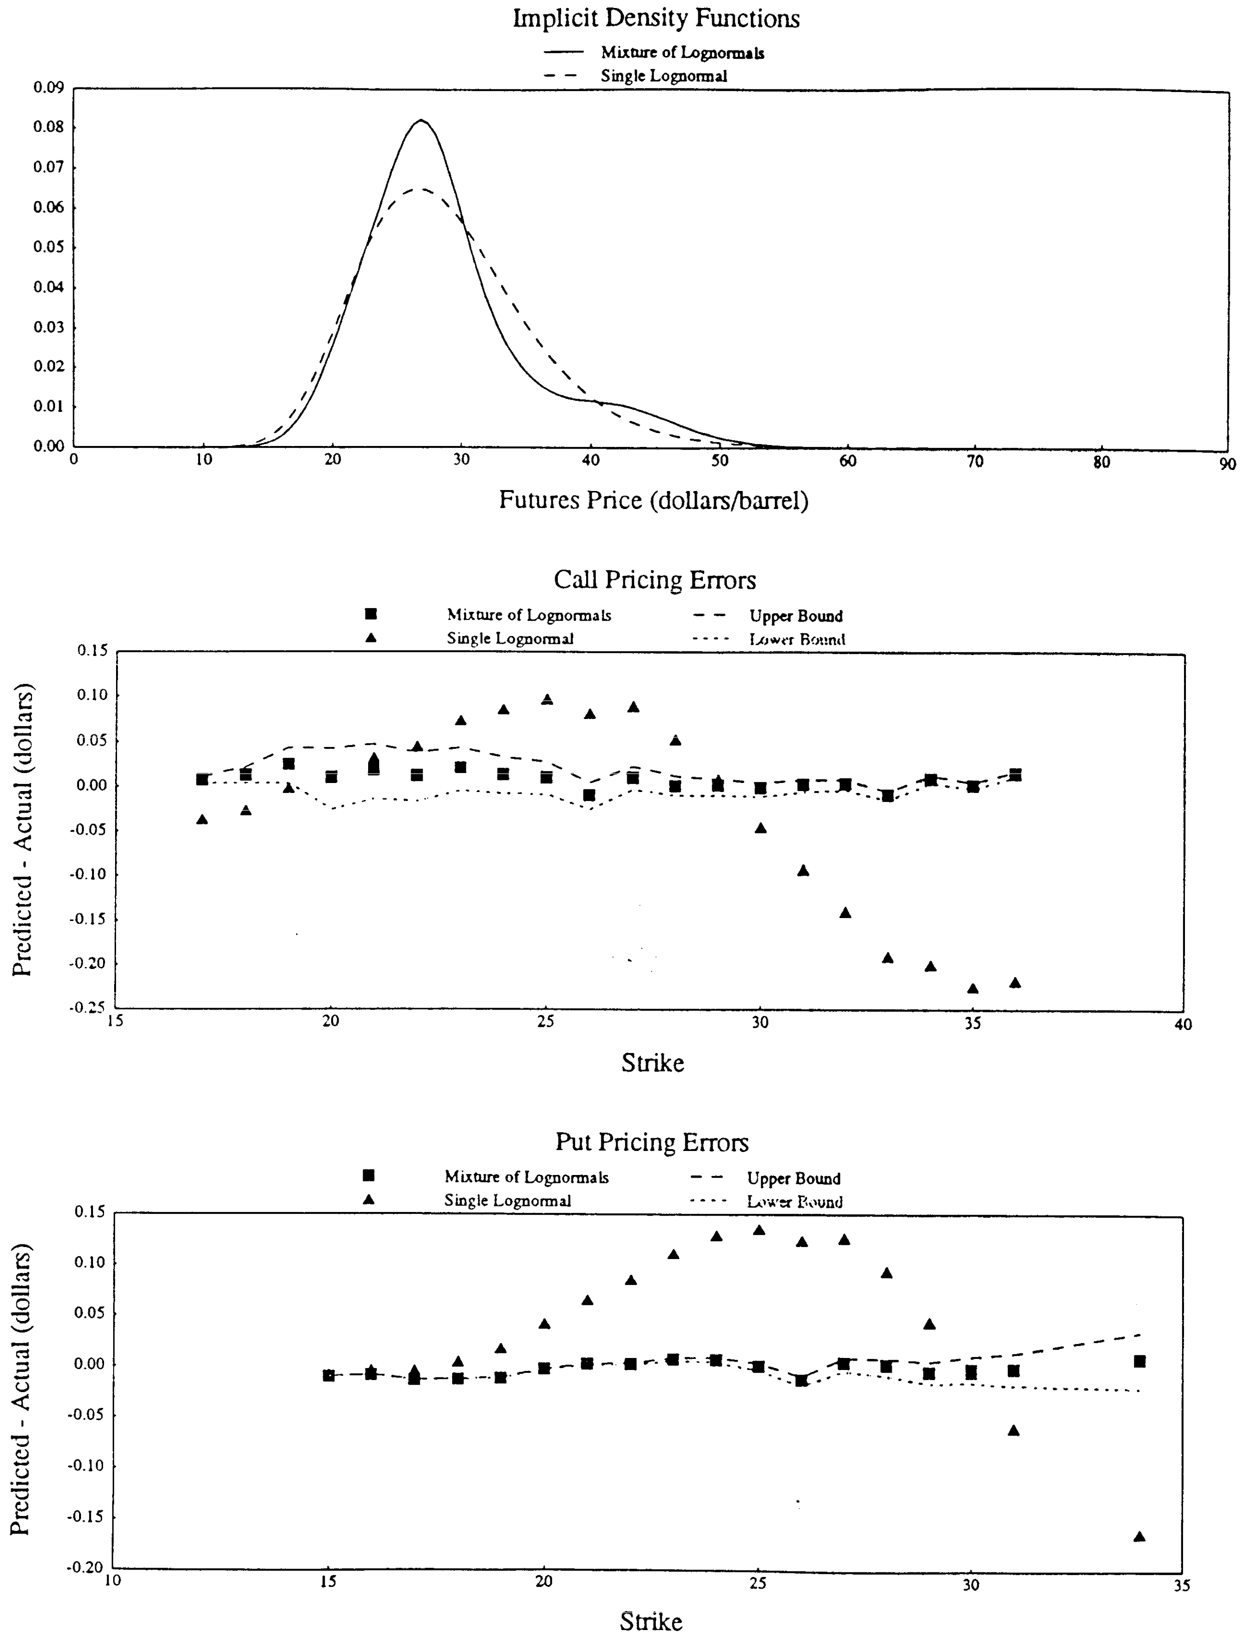

Chart 4

Actual - Predicted Futures Price vs. Change in Actual Futures Price

All Contracts

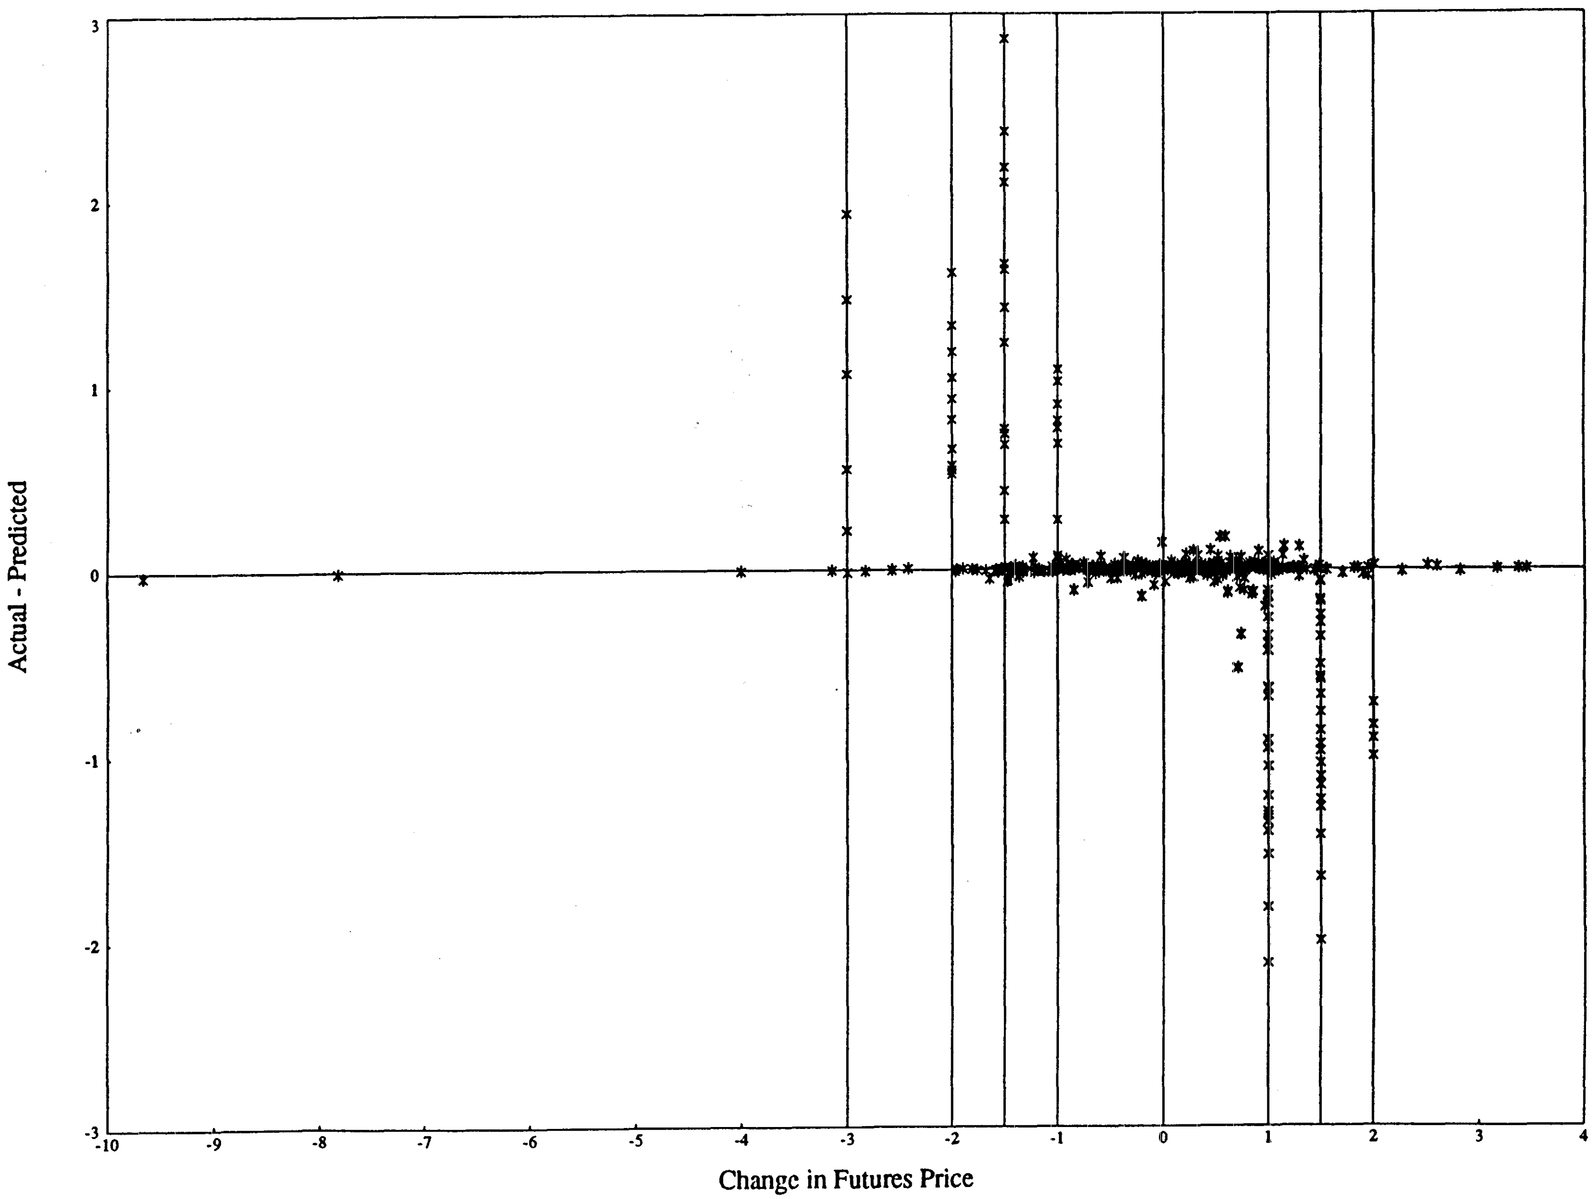

Chart 5
Implicit Density Functions

July 10, October Contract

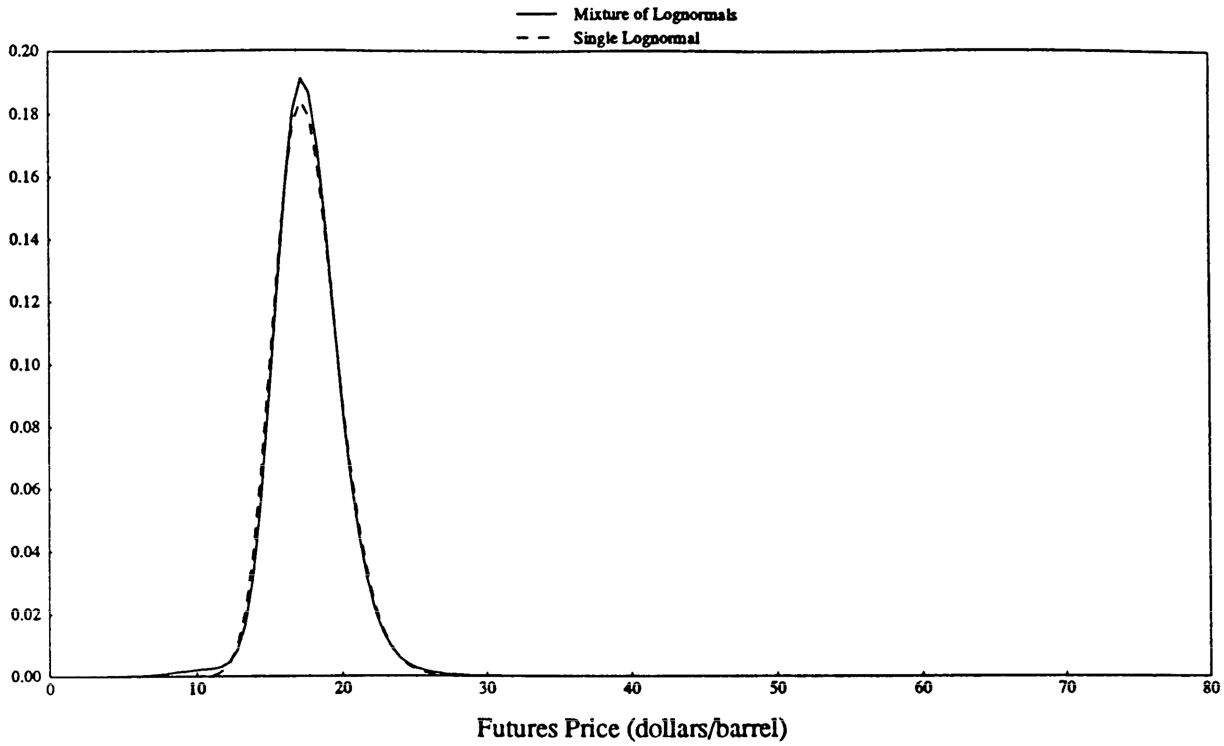

October 18, January Contract

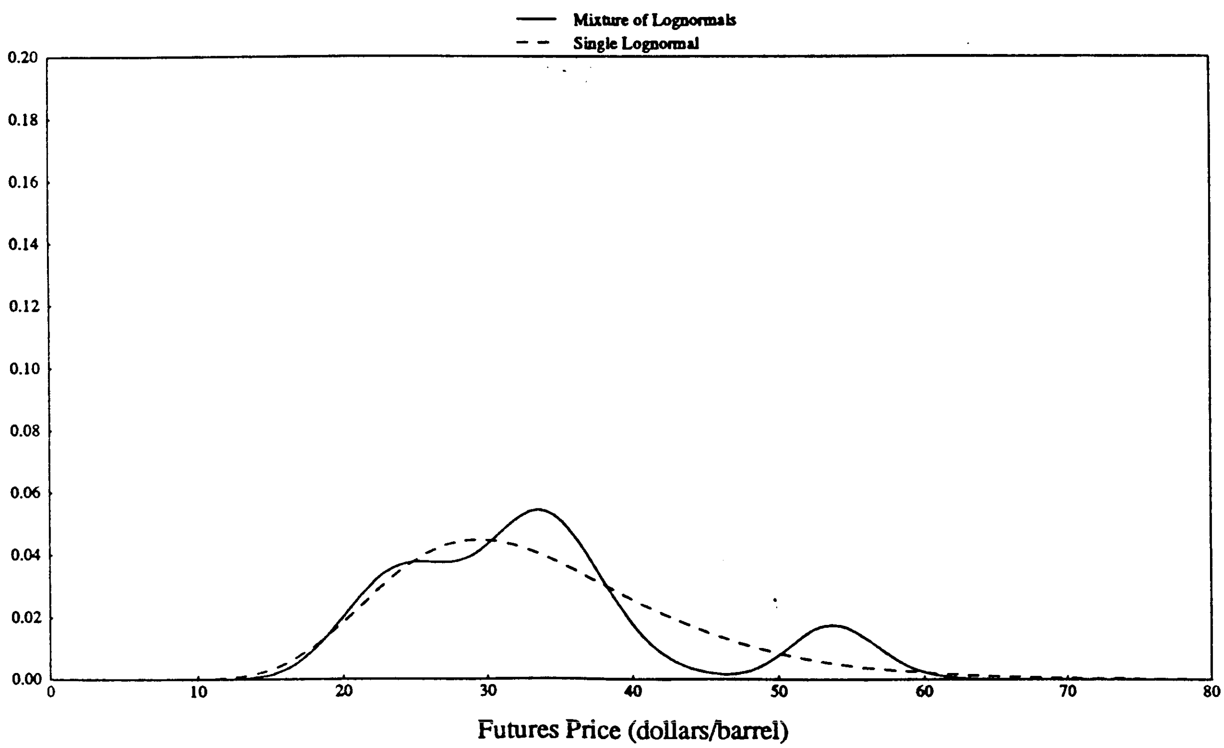

Chart 6
Right-Hand Tail Characteristics: Implicit Density Functions
April Contract

Expected Price Conditional on Exceeding the Futures by 25 Percent

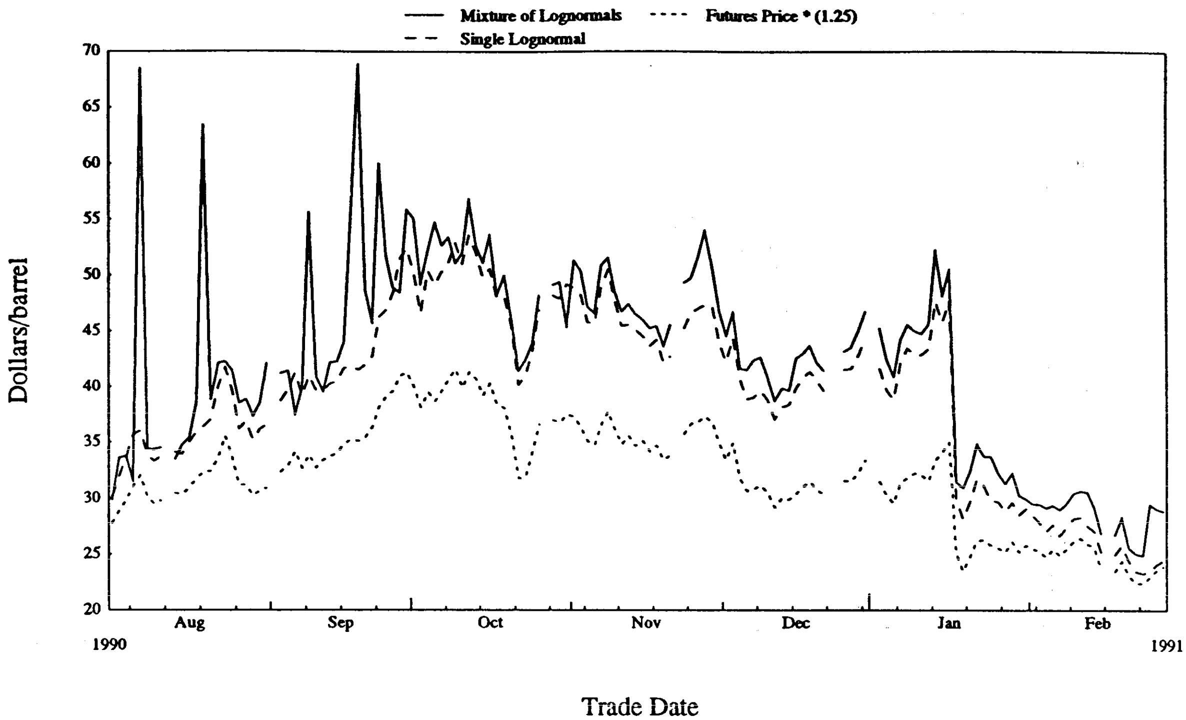

Probability that Price Will Exceed the Futures by 25 Percent

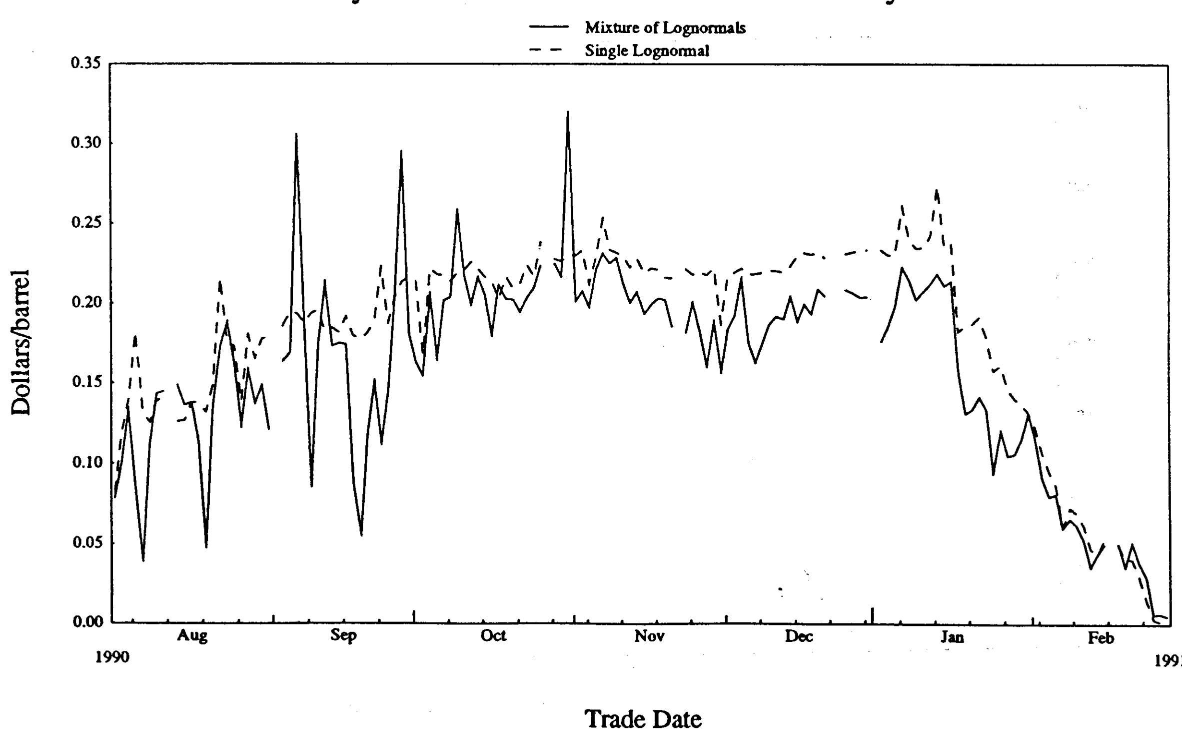

Chart 7
Mining of Wells: Implicit Density Functions

October 22, January Contract

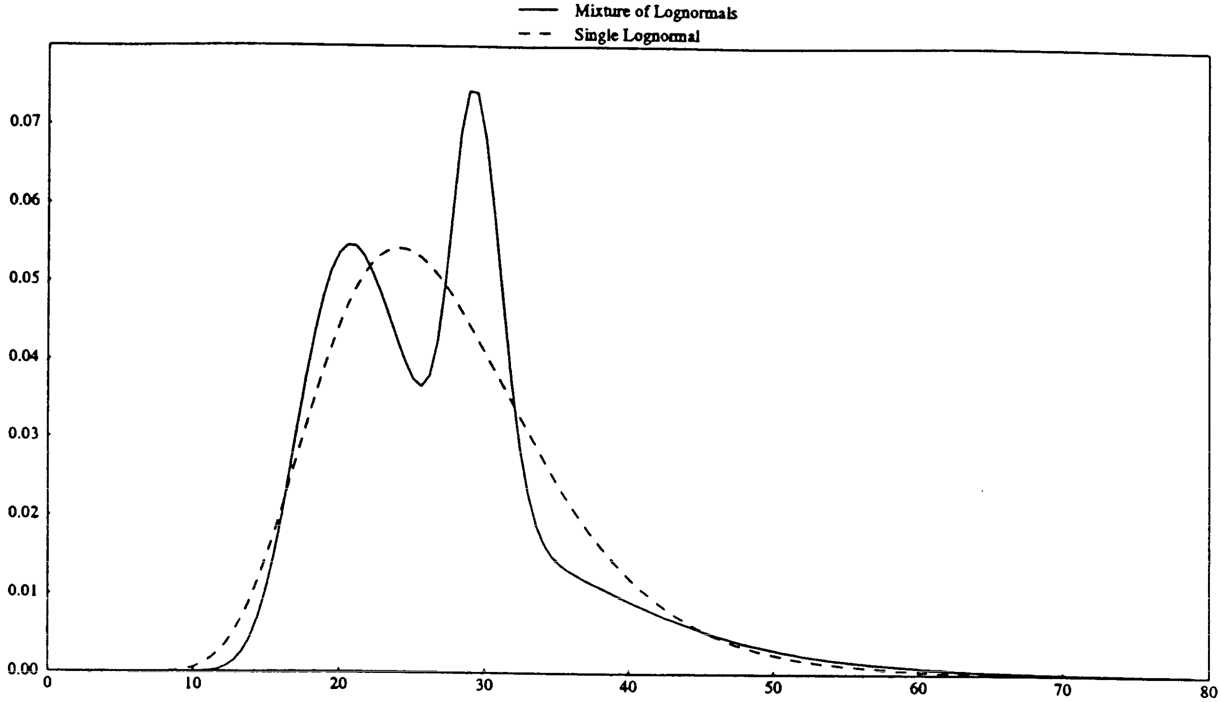

Futures Price (dollars/barrel)

October 25, January Contract

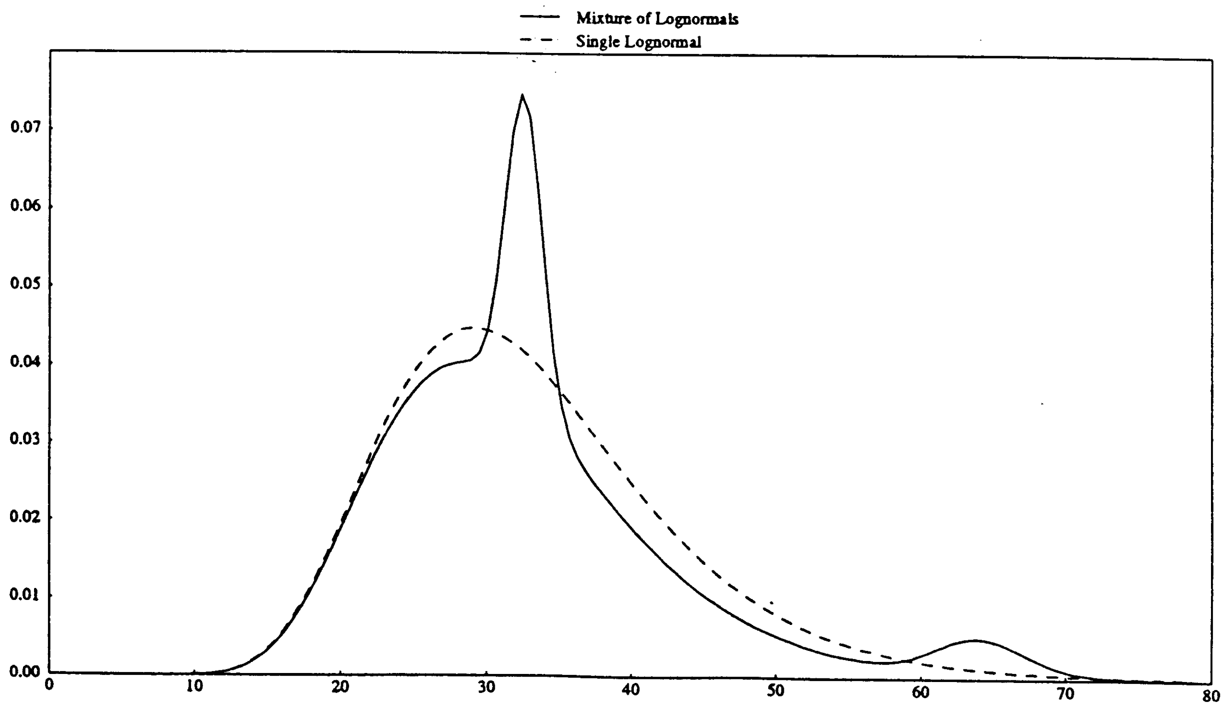

Futures Price (dollars/barrel)

Chart 8

Start of Air War: Implicit Density Functions

Mixture of Lognormals (____) Single Lognormal (_ _ _), April Contract